ILLINOIS INSTITUTE OF TECHNOLOGY
ECE 565 MULTIDIMENSIONAL DIGITAL SPEECH PROCESSING (FALL'98)
Project 1: Fourier Representation of Images

Due Date: November 10, 1998

Objectives

The work items of this computer simulation project are:

1. Implement 2-D DFT using a 1-D FFT algorithm.
2. Implement 2-D IDFT using the 2-D DFT program.
3. Visualize the magnitude spectrum of an image.
4. Synthesize an image from its Fourier components.
5. (Bonus) Perform 2-D FIR filtering.

Ground Rules

You must write your own programs, run your own simulations, and analyse and interpret your simulation results. Submit a report on paper. Document any theoretical/analytical preparation you did (or simply refer to specific material in the textbook) as part of your design of your computer programs. Document your simulation results and plots, any extra effort in addition to the required work detailed below, and your observations and conclusions. Grading is based on the quality of your report and simulation results and any "extra mile" you put into the work. Present your results compactly, clearly, and selectively. A report "stuffed" with redundant material will not qualify for a high grade. Email all the programs you have written for the project, in one *plain-text* file, to chan@ece.iit.edu. Do not include your program printouts with your report.

Your programs may be written in any computer language of your choice. The "student edition" of Matlab may *not* have sufficient capability to handle the project tasks. Your programming style is a factor to consider. It is ultimately your responsibility to choose the language and tools.

The only preprogramed library function you are permitted to use to perform computation for this project is any suitable 1-D FFT routine. Matlab and many other tools supply this, and you will also find sample programs in the book *Numerical Recipes in C*. You may use other preprogrammed functions (e.g. 2-D DFT) to verify your simulation results, but you must write your own programs to generate your results, in accordance with the prescriptions given below, and using no library function other than 1-D FFT.

You may discuss with your classmates about the project, but your effort must be *essentially independent*. Learning is most effective when you work through the problems on your own, and at this stage of learning, the quality of the learning process is more important than the results generated. The minimum penalty for academic dishonesty is a failure grade in the course.

Project Tasks

1. Write a routine to compute a 64×64 2-D DFT. The routine repeatedly calls a library-routine to perform multiple 64-point 1-D FFTs.

2. Write a program to compute a 128×128 2-D DFT. For each 128×128 2-D DFT, the program from task (1) is called several times. Additional computation is needed to combine the 64×64 2-D DFT results.

3. Use the program written for task (2) above, compute the DFT of a 128×128 grey scale image of Lena.

4. Plot the magnitude spectrum of Lena as a grey scale image. Choose a suitable amplitude scale for the spectrum so you can see as much of the texture details of the magnitude spectrum as possible. The origin of the frequency plane should be at the center of the plot.

5. Write a program that performs 2-D 128×128 IDFT by calling the 128×128 2-D DFT routine written for task (2) above. Suitable preprocessing of the input DFT data and postprocessing of the data produced by the 2-D DFT routine may (not?) be necessary in order to perform IDFT using a DFT routine.

6. Synthesize an image using only the phase spectrum of Lena. Plot the synthesized image. Suppose a complex number Z is represented using rectangular coordinates. The magnitude of Z can be set to an arbitrary constant $C > 0$ while the phase of Z is preserved, by multiplying the real and imaginary parts of Z by $C/|Z|$, where $|Z|$ is the modulus if Z. If Z is represented in polar form, then simply set the modulus of Z to C.

7. Synthesize an image using only the magnitude spectrum of Lena. Plot the synthesized image. To zero the phase of a complex number Z represented in rectangular coordinates, replace the real part with $|Z|$ and the imaginary part with zero.

8. Synthesize an image using the magnitude spectrum of Lena augmented by the sign of the spectral amplitude. The nonnegative magnitude spectrum is allowed to become negative by taking the sign of the real part. The operation in task (7) is changed to: replace the real part of Z by $\text{sign}(\text{Re}(Z))|Z|$ and the imaginary part with zero. Equivalently, the phase of Z is quantized to either 0 or π radian. Plot the synthesized image. Of the three synthesized images, which one is most visually informative (relative to the original image)?

Pragmatics

1. The amplitude of each pixel of a grey scale image is represented as an unsigned integer ranging from 0 to 255, with 0 mapped to "dark" and 255 mapped to "white." Before plotting a computed image, its amplitude may need to be scaled and its numeric format properly converted. This must be done with care, so that the pixel numeric range is broadly exercised while no clipping is incurred.

2. The numerical precision used for computation can impact the speed of processing and accuracy of your results.

3. Some languages may not implement indexing of 2-D arrays efficiently.

4. There are a number of ways you can test your DFT programs. For instance, by setting a (properly selected) pair of DFT coefficients to 1 and the rest to zero, you should be able to get a 2-D (sampled) sinusoid in the spatial domain. Conversely, a 2-D sinusoid with suitably selected spatial frequencies would produce a DFT with only two nonzero coefficients.

5. Beware that in order to synthesize a real image, you need certain kind of symmetry.

Work for Bonus Points

You may use preprogrammed Matlab functions to perform these tasks. Only serious and well-thought-out (even if incomplete) attempts will bring you extra points.

1. You can get a feel of the robustness of phase information by replacing the magnitude spectrum with random numbers and observing how much the synthesized image differs from the one synthesized with a constant magnitude spectrum.

2. Suppose you are asked to filter the Lena image with a real FIR filter whose impulse response has size $N \times N$ where N is an odd number. The desired output of the filter is an image as close to that obtained by the phase-only synthesis as possible. How would you design such a filter? Try constructing such a filter for reasonably small N and perform filtering in the spatial domain. How big an N is necessary for the filtered image to look like the phase-only synthesized image?