"Digital Signal Processing With Applications:" A New and Successful Approach to Undergraduate DSP Education [1] [2]

Michael D. Zoltowski, Jan P. Allebach, and Charles A. Bouman

School of Electrical Engineering
1285 Electrical Engineering Building
Purdue University
West Lafayette, IN 47907 USA
e-mail: mikedz@ecn.purdue.edu
Phone: 317-494-3512
FAX: 317-494-6440

Abstract

The new approach to undergraduate DSP education at Purdue is based on a simple idea: emphasize applications. Students are assumed to have a significant exposure to sampling and discrete-time signals, systems, and transforms at the junior level. In the senior course, the traditional DSP topics of digital filter design, the DFT, radix-2 FFT's, and quantization are covered in the first five weeks of the semester. Coverage of these topics is augmented by treatment in the laboratory component of the course using diverse software tools and by Matlab based homework assignments. The remainder of the course is devoted to treating the topics of speech processing and image processing in substantial depth and involves a design project. The course has been very successful in terms of increasing enrollment and outstanding student evaluations.

Submitted to IEEE Transactions on Education

Special Issue on Undergraduate Digital Signal Processing Education

[1]This research was supported by Purdue University, including an Undergraduate Curriculum Enrichment Travel Grant.

[2]*IEEE Trans. on Education*, vol. 39, no. 2, pp. 120-126.

1 Introduction

Our undergraduate DSP education approach is a departure from the conventional route of teaching sampling and discrete-time signals, systems, and transforms to juniors, and then following up with a senior course reviewing these fundamental topics and treating FIR and IIR digital filter design, the Discrete Fourier Transform (DFT), radix-2 FFT implementations, quantization, and classical spectrum estimation. For the past five years, we have taught a senior undergraduate DSP course, entitled *Digital Signal Processing with Applications*, with heavy emphasis on applications, primarily speech processing, image processing, and array signal processing. Each of these application areas is treated in substantial depth with coverage of a number of attendant topics that are reserved for graduate school at other schools of electrical engineering. For example, we know of no other school that teaches the principles of image reconstruction from projections, the basis for computed tomography, at the undergraduate level.

Yet, despite the advanced level of the course, it has been very successful by a number of different measures including steadily increasing enrollment, outstanding evaluations from the students, and strong positive feedback from those students who have gone on to graduate school or industry regarding their preparedness for graduate work or research and development in DSP. If there is anything to be learned from our experience, it is that there is no reason to reserve the topics of speech and image processing for graduate level courses. The success of our course has proven that it is possible to treat these application areas in substantial depth at an undergraduate level.

2 Motivation

In March of 1988, during the formative process of this course, an effort was made to assess what was being taught in undergraduate DSP courses at other schools of electrical engineering in the USA. At that time, the following conclusions were drawn:

1. The content of these courses tended to be dominated by a relatively small set of traditional topics.

2. The courses primarily treated DSP in an abstract setting without any substantial discussion of applications.

3. The courses were taught without a laboratory.

With regard to the last point, we should point out that a growing number of schools have a DSP laboratory course separate from the lecture course. These laboratory courses typically concentrate on the implementation of DSP algorithms on DSP chips.

3 Innovations of New Approach to DSP Education

3.1 In-Depth Treatment of Applications

The major innovation is the extensive and in-depth treatment of applications of DSP. Only five weeks of the course are devoted to traditional DSP topics. During the first two years of its offering, the remaining 9 weeks of the course were evenly divided among three different application areas: speech processing, image processing, and array signal processing. Since the institution of a design project in Fall of 1991, the amount of time and coverage spent on these topics now vary from semester to semester depending on which application area is emphasized in the design project. Brief descriptions of the topics covered in these three application areas are provided below.

- *Speech processing:* vocal tract models and characteristics of the speech waveform; short-time spectral analysis and synthesis; linear predictive coding.

- *Image Processing:* two dimensional signals, systems, and spectral analysis; image enhancement; image coding; and image reconstruction.

- *Array Processing:* basic radar principles; representation of propagating waves; delay-and-sum beamformer; array pattern.

It is important to note that aside from exposing the students to these application areas, our approach also provides a vehicle for introducing a number of important analytical tools and concepts that are not usually seen by undergraduates. These include short-time Fourier analysis, linear prediction, two-dimensional signals and systems, the Radon Transform, and spatial filtering using an array of sensors.

3.2 Design Project

Another innovation relative to undergraduate DSP education is the incorporation of a course driven design project. We have found a synergism in the tight coupling of project and course material i.e. project experience reinforces lecture theory, and techniques learned in lecture often motivate project strategies.

The projects are executed by small groups of students with each group attacking the same design problem. The design problems are open ended and allow for a wide variety of possible solutions involving engineering trade-offs between performance and cost. Each student team is required to provide deliverables in the form of a project proposal, progress report, final report, and final oral presentation. The design project topics for the past eight semesters are listed below.

- *Fall '91:* Still Image Coding.
- *Spring '92:* Isolated Word Recognition.
- *Fall '92:* Motion Compensated Video Coding.
- *Spring '93:* Image Scaling and Rotation.
- *Fall '93:* Audio Coding.
- *Spring '94:* Image Deblurring.
- *Fall '94:* Optical Character Recognition.
- *Spring '95:* Compression of Half-Tone Images.

As discussed previously and as indicated in the outline for the course delineated in Appendix A, extra lecture time is devoted to the application area in which the design project falls. For example, during the Spring '92 semester additional speech topics were covered including cepstral analysis, template matching, distance measures, dynamic time warping, and Hidden Markov Models.

3.2.1 Design Project Example

As an example of the type and scope of design project assigned in the course, we briefly describe the speech processing based design project assigned in the Spring '92 semester. The goal was to design and implement a speaker independent word recognition system for a limited vocabulary of four words: yes, no, stop, enter. The words were chosen to be quite phonetically different to allow for a wide variety of creative solutions, and to make completion of the project within the second half of the semester feasible. The projects were judged on the following criteria:

1. The performance of the word recognizer.
2. The computational complexity (cpu run time) of the word recognizer.
3. The creativity, level of sophistication, and/or generality of the approach.
4. The quality of the final report.
5. The quality of the final presentation.

For training and/or testing their algorithms, the students were given access to data files containing ten utterances of each of the four words by each of four different male speakers. The students were directed to include the following items in the Final Report:

1. An overview of the word recognizer. Include references to any sources that were used to develop the method.
2. A summary of the word recognizer's performance on the speech data files provided.
3. A brief description of the software modules that were developed.
4. An analysis of the computational complexity of the overall algorithm.
5. A set of instructions on how to execute the program on a speech data file. The students were directed to structure their program so that execution on a single data file could be executed by a single (simple or shell) command, and should require no user interaction once execution begins.

6. A detailed description of the specific contributions of each team member.

7. Recommendations for further development of the method.

8. (Appendix) Complete listing of source code for the method. The students were directed to document the code at a reasonable level.

Relative to item 1 in the judging criteria, the students' algorithms were tested with data files to which the students did not have access. These files contained utterances of the same four words from the same four speakers, and other speakers as well to verify speaker independence.

To illustrate the creativity and/or level of sophistication of the approaches taken by the different design teams, we briefly describe three approaches that achieved 100% accuracy on 16 test words not supplied to the students.

- *Neural network approach.* The inputs to the neural network included LPC coefficients, prediction error, number of zero crossings, and peak amplitude for each frame.

- *Template Matching approach.* Cepstral coefficients were computed for each frame, and distances to stored templates were computed via dynamic time warping.

- *Ad-Hoc approach.* The word recognizer exploited the location of the unvoiced phoneme in each of the words: STOP - beginning, ENTER - middle, YES - end, NO – no unvoiced phoneme. The voiced/unvoiced decision for each frame was determined from LPC analysis.

Note that out of 11 groups no two groups took exactly the same approach. One other group took a neural network approach, several groups took template matching approaches, several groups took ad-hoc approaches, and two group used Hidden Markov Models.

3.3 Laboratory Component

Another innovation is a laboratory component to the course. The lab is very useful for a number of purposes including (i) exposing students to a variety of software tools for designing digital filters, analyzing speech waveforms, analyzing images, etc., (ii) helping students visualize phenomena such as aliasing (in both 1D and 2D signals), quantization effects, beamforming, etc., and (iii) augmenting coverage of topics treated briefly in lecture due to the desire to get to the speech and image processing applications. Relative to the last point, two labs are devoted exclusively to FIR and IIR digital filter design. Given the ability of the students to access digital filter design packages on workstations in the DSP laboratory, it makes sense to give more time to these topics in the laboratory than in lecture.

To illustrate the type, scope, and complexity of experiments the students are required to carry out in the laboratory, brief descriptions of the experiments are provided below. Note that the laboratory is primarily software driven, although interested students can make use of Motorola DSP chips in the laboratory.

1. *Introduction to the Laboratory.* Introduction to the Hewlett-Packard Workstations and X-windows environment, use of the Entropic software, communication with other ECN machines at Purdue.

2. *Linear Systems.* Demonstrations of properties of linear systems.

3. *Sampling and Quantization.* Demonstration of sampling principles, design of optimal quantizers, design trade-offs in D/A conversion.

4. *Design of Finite Impulse Response (FIR) Filters.* Design of FIR filters using the Kaiser window, the Parks-McClellan algorithm, and a minimum weighted mean-squared error criterion.

5. *Design of Infinite Impulse Response (IIR) Filters.* Design of Butterworth and Chebyshev Type I IIR filters, comparison of the performance of FIR and IIR filters.

6. *Analysis of Climatic and Stock Market Data.* Design of digital signal processing algorithms for the analysis of temperature and precipitation records, and the generation of buy/sell indicators from stock market data.

7. *Digital Processing of EKG Signals.* Design of linear and nonlinear (morphological) filters for removal of sinusoidal interference and baseline drift in EKG signals.

8. *Time and Frequency Domain Analysis of Speech.* Examination of temporal characteristics of speech, analysis of speech via spectrograms, relation between short-time Fourier transform and the spectrogram.

9. *Speech Vocoder.* Design of a speech coding system based on linear prediction with voiced/unvoiced classification, pitch extraction, and gain computation.

10. *Image Coding.* Design of image coding algorithms via three different approaches: block truncation coding, transform coding using the discrete cosine transform (DCT), and predictive coding.

11. *Image Enhancement and Restoration.* Design of gray-scale transformations for contrast enhancement, quantization, and feature selection. Design of spatial filters for image sharpening, blur removal, and descreening.

12. *Computed Tomography.* Algebraic reconstruction technique (ART), demonstration of reconstruction based on Fourier slice theorem, sampling effects.

13. *Beamforming.* Demonstration of beamforming process, design of receiver array configuration and receiver delays to achieve desired beamforming objective.

3.4 Novel Use of A/V Material

Coverage of each of the three primary application areas is augmented through the use of in-class audio-visual material, as well as through extensive treatment in the laboratory using such software tools as *Entropic.* During the coverage on speech processing, a videotape is shown demonstrating some of the capabilities of the equipment in the *Speech Processing Laboratory* at Purdue University run by Prof. Leah Jamieson. For example, the videotape shows a researcher training a computer to recognize key (spoken) words for file management in a Unix based operating system.

As part of the coverage on image processing, a set of slides have been developed to illustrate basic two-dimensional (2D) Fourier Transform pairs, aliasing effects occurring with two dimensional sampling, and properties of the 2D Fourier Transform. As an example, Figure 1 shows the effects of undersampling a 2D sinewave. For each gray-scale plot image in Figure 1, we also display the magnitude of the corresponding 2D Fourier Transform via an attendant gray-scale image. In the upper left hand corner of Figure 1(a), the small and large intensities (the light and dark bands) represent the valleys and crests (peaks) of the 2D sinewave. The magnitude of the corresponding 2D Fourier Transform plotted in the upper right hand corner of Figure 1(a) reveals the symmetry property of 2D Fourier Transform with a real-valued 2D signal and also reveals that the energy of the 2D sinewave is concentrated at its spatial frequency (as well as at DC due to the average gray level).

The lower left hand corner of Figure 1(a) (which is repeated in the upper left hand corner of Figure 1(b)) displays an undersampled version of the 2D sinewave. The attendant 2D Fourier Transform magnitude plotted in the lower right hand corner of Figure 1(a) (and repeated in the upper right hand corner of Figure 1(b)) shows how 2D sampling over a rectangular lattice yields replicas of the 2D Fourier Transform for the corresponding image. Undersampling manifests itself in terms of overlap amongst the spectral replicas. The lower left hand corner of Figure 1(b) shows that using a 2D lowpass filter, with cut-off at half the spatial sampling rate along either major axis, to reconstruct the image produces a 2D sinewave propagating in a completely different direction due to aliasing. These two slides lead to another set of slides showing Moire patterns in scenic images due to aliasing.

At this point, it is worthwhile to note that in addition to introducing undergraduate students to a rapidly emerging area of technology with broad application, coverage of image processing and the attendant theory of 2D signals, systems, and transforms reinforces the knowledge of 1D signals, systems, and transforms that the undergraduate student gains through the sophomore and junior years. In addition, some students relate to 2D better because they can visualize as well as see images of scenes and/or objects that they can relate to as opposed to a plot of voltage versus time for a one-dimensional signal, for example.

Another set of slides uses gray-scale images to illustrate a variety of image enhancement techniques including simple 2D linear filtering and histogram modification via gray-scale transformations. Figure 2 shows how a simple 9 tap 2D FIR filter can be used to sharpen an image. The original and sharpened images are pictured in the upper and lower left hand corners, respectively. The magnitude and phase of the 2D Fourier Transform of the original and sharpened images are plotted in the upper and lower right hand corners, respectively

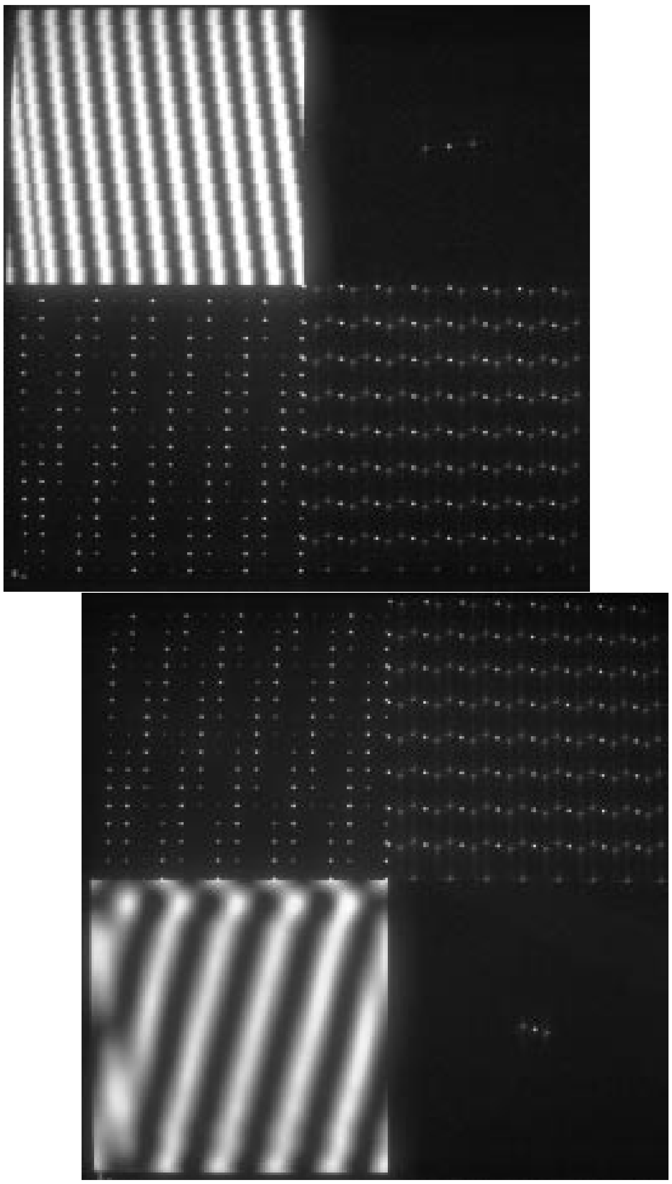

a | b

Figure 1: Sampling of a 2D sinewave and illustration of aliasing effects of undersampling

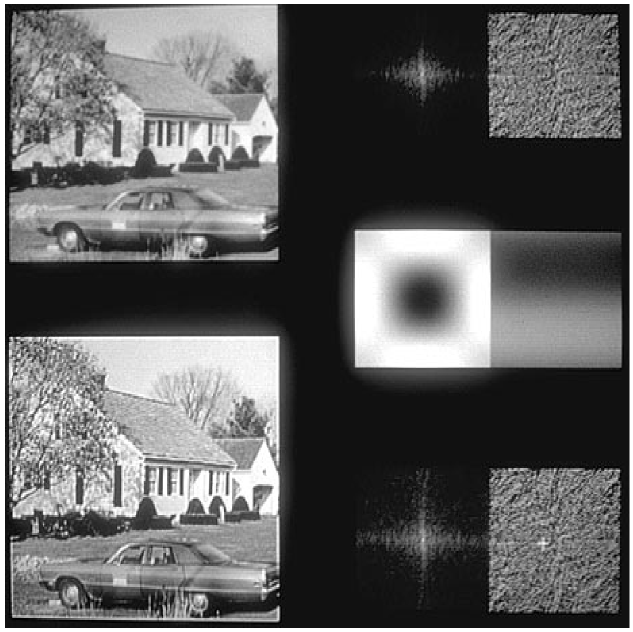

Figure 2: Illustration of image sharpening via simple 2D linear filtering.

(the phase plots are in the far right side corners.) The magnitude and phase of the 2D filter implementing an "unsharp mask" are plotted in the middle right side and it is indeed observed to be a highpass filter.

Figure 3(a) shows how simple modification of the histogram of pixel intensities can brighten up an image. The original and brightened image are shown in the lower right and upper left hand corners of Figure 3(a), respectively. The binary image in the upper right hand corner shows the gray-scale histogram before and after the gray-scale transformation which is depicted as well. Figure 3(b) shows how contours can be generated via a simple gray-scale transformation. The binary image in the upper right hand corner of Figure 3(b) depicts the gray-scale transformation wherein every pixel is mapped to one of two different intensities.

With regard to the coverage on computed tomography, the slides comprising Figure 4 are used to illustrate image reconstruction from projections using the filtered back-projection method. This sequence of slides illustrates how the original object, the letter "O", is progressively formed via the superposition of filtered back-projections over a continuum of projection angles. In each figure, the upper left hand corner depicts the filtered back-projection at a

a | b

Figure 3: Image brightening and contour generation via gray-scale transformation s.

particular angle, while the upper right hand corner shows the sum of the filtered back-projections up to but not including the adjacent filtered back-projection. The upper right hand corner of Figure 4(d) shows the final reconstructed image which may be compared to the original image pictured in the lower left hand corner in each of Figures 4(a)-(d). At the end of the coverage on computed tomography, we show slides of actual CAT scans donated by GE Medical Systems.

As part of the coverage on array signal processing, slides are shown illustrating basic principles of phased array radar. Many of these slides were culled from a tutorial article by Eli Brookner call "Phased Array Radar" that appeared in the February 1985 issue of *Scientific American*.

4 Miscellaneous Issues

4.1 Required Background for Students

The course assumes that at the junior level the student obtains a solid foundation in 1-D sampling theory, discrete-time signals, LTI discrete-time systems, the Discrete-Time Fourier Transform (DTFT), and the Z Transform. At Purdue University, these topics are taught in a junior level course entitled *Signals and Systems* using the text **Signals and Systems** by Oppenheim, Willsky, and Young (Prentice-Hall, 1983). Thus, these topics are not taught "from scratch" at the senior level. Rather, they are quickly reviewed and this foundation is built upon through coverage of FIR and IIR digital filter design, the Discrete Fourier Transform (DFT) and its properties including circular convolution and leakage effects, radix-2 FFT implementations, and quantization. As indicated in the course outline delineated in Appendix A, this *Foundations* unit of the course is completed in the first six weeks of the semester. Again, we are able to cover these topics faster than would normally be the case due to attendant coverage in the laboratory.

It should be pointed out that even at the junior level, an effort is underway to emphasize applications of DSP and thereby make DSP more tangible and concrete to undergraduate engineering students. As a primary example, in the junior level *Signals and Systems* course, CD technology is invoked throughout the course as an enormously successful working example

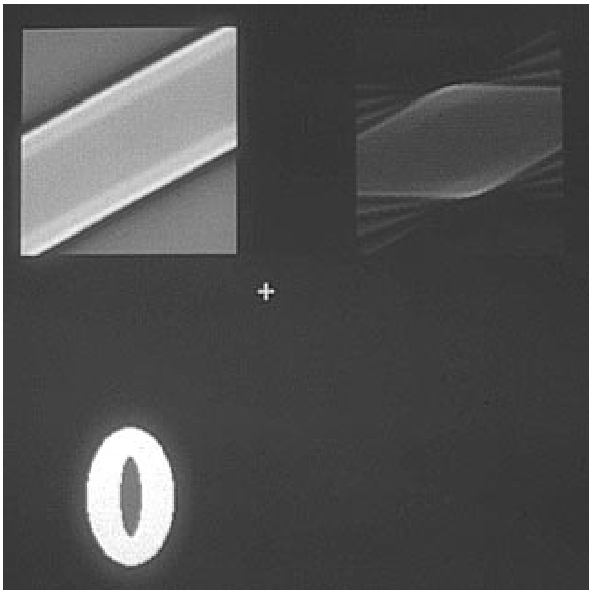

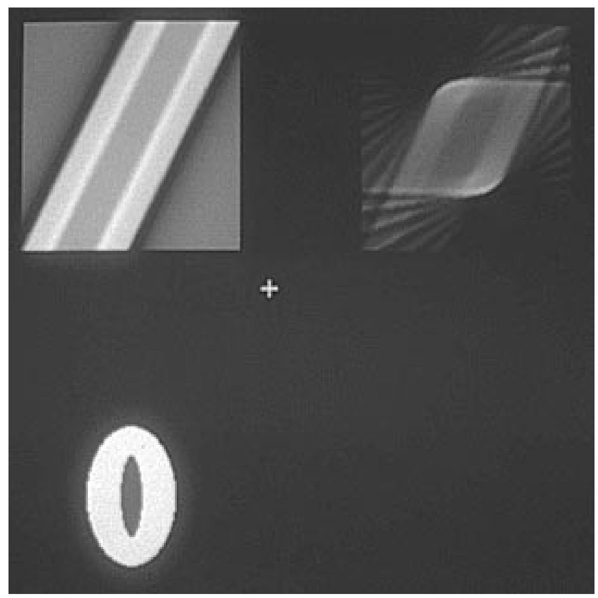

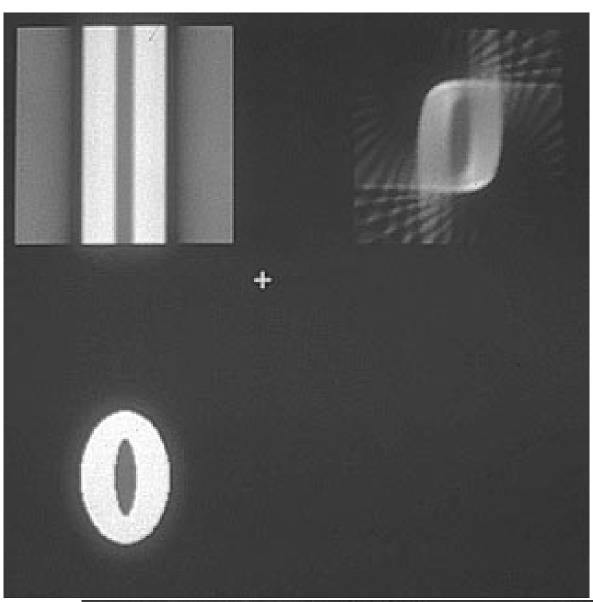

of DSP concepts applied in practice. We discuss how the bits are stored on the disk, how the CD player reads the bits off the disk, what the much touted features of 8 or 4 times oversampling and single bit D/A conversion are all about, where digital filters implemented as difference equations come into play, etc. A CD player is brought into the class on several occasions to illustrate operating principles of the CD player, e.g., to illustrate the error correcting capabilities. It should be noted that this is also part of an effort to integrate system level optics concepts into the lower level undergraduate EE curriculum.

4.2 Use of Matlab in Homework Assignments

Students are expected to have access to Matlab. Homework problems requiring the use of Matlab are assigned on a regular basis and are not treated as anything special. This viewpoint stems from the fact that at Purdue even the junior signals and systems course requires the students to use Matlab and Simulab regularly in homework assignments.

4.3 Measures of Success

As one measure of success, enrollment has increased from 33 students in the first offering of the course in Fall '88 to 63 students in the Spring '95 offering. Note that the course is offered in both the Fall and Spring semesters each year. A plot of the enrollment per semester for each semester since the course's inception is plotted in Figure 5. This rise has come during a time when our undergraduate EE enrollment (sophomores, juniors, and seniors) has dropped precipitously from 1300 students to 875 students. As another measure of success, the teaching of this course has directly resulted in two teaching awards: one voted upon by all seniors and one conferred by the Purdue Chapter of the HKN electrical engineering honorary society.

4.4 Textbook Issues

Presently, there is no suitable textbook for the course. **Discrete-Time Signal Processing** by Oppenheim and Schafer (Prentice-Hall, 1989) is used for the Foundations Unit of the course. For the treatment of applications, we are currently contracting a publishing company to reproduce and bind together under one cover the following supplementary reference

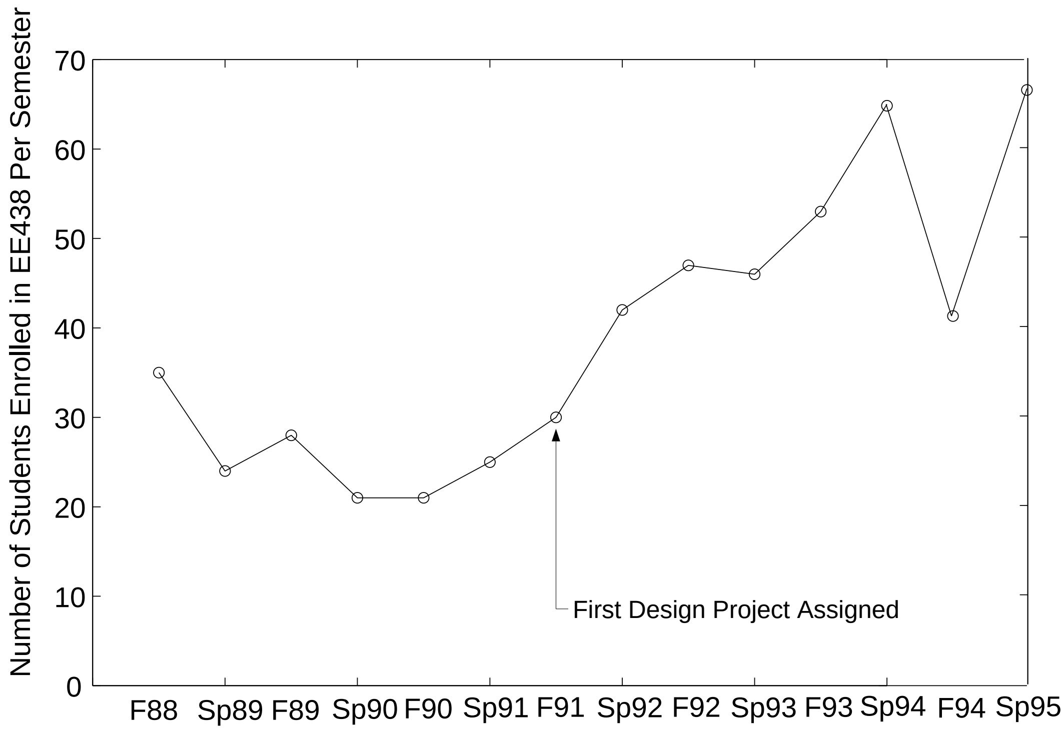

Figure 5: Enrollment versus time.

materials:

- Section 1: Image Fundamentals. Chapter 2, pp 21-40, Digital Image Processing, by R. Gonzalez and R. Woods Addison-Wesley Publishing, Reading MA, 1992.
- Section 2: Image Transforms. Chapter 3, pp 81-119, Digital Image Processing, by R. Gonzalez and R. Woods Addison-Wesley Publishing, Reading MA, 1992.
- Section 3: Image Enhancement. Chapter 4, pp 161-213, Digital Image Processing, by R. Gonzalez and R. Woods Addison-Wesley Publishing, Reading MA, 1992.
- Section 4: Image Reconstruction. Chapter 8, pp 353-430, Digital Picture Processing, by A. Rosenfeld and A. Kak Academic Press, San Diego CA, 1982.
- Section 5: Speech Processing. Chapters 2,3, pp 11-140, Fundamentals of Speech Recognition, by L. Rabiner and B. Juang Prentice Hall, Englewood Cliffs NJ, 1993.

4.5 Exposure to Hardware Implementation

No lecture time is devoted to topics related to implementing DSP algorithms on DSP chips. Also, at present the laboratory is primarily software driven. Although plans are underway to incorporate the use of the Motorola DSP56000 chip into the laboratory, the students are strongly encouraged to take an advanced microprocessor course which covers digital signal processors.

5 Conclusion

The use of novel A/V material in lecture, the use of image and speech processing software tools in the laboratory, and an open-ended design project may be used as vehicles for successfully treating the application areas of speech and image processing in substantial depth at the senior undergraduate level.

Appendix A: Current Course Outline.

1. **Foundations** (6 weeks)
   - (a) Continuous-time and discrete-time signals and spectral analysis (CTFT and DTFT)
   - (b) Continuous-time and discrete-time systems
   - (c) Sampling
   - (d) Decimation and interpolation
   - (e) Z Transform
   - (f) Discrete Fourier Transform (DFT) and Fast Fourier Transform (FFT)
   - (g) Probabilistic methods in digital signal processing
   - (h) Quantization

   Two Units from topics 2.0, 3.0 and 4.0

2. **Speech and Audio Processing** (3 weeks)
   - (a) Speech models and characteristics
   - (b) Short-time Fourier analysis and synthesis
   - (c) Differential Pulse Code Modulation (DPCM)
   - (d) Linear predictive coding

3. **Image Processing** (3 weeks)

(a) 2-D signals and systems
(b) Image coding
(c) Image enhancement
(d) Computed tomography

4. **Array Signal Processing** (3 weeks)

    (a) Basic radar principles
    (b) Representation of propagating waves
    (c) Beamforming

5. **Project Preparation** (2 weeks)

6. **Examinations** (1 week)