120th ASEE Annual Conference & Exposition
FRANKLY, WE DO GIVE A D*MN
June 23-26, 2013

Paper ID #6953

A Longitudinal Study of Student Performance in an Elective Applied Digital Signal Processing Course

Dr. Roger A. Green, North Dakota State University

Roger Green received the B.S. degree in electrical and computer engineering and the M.S. and Ph.D. degrees in electrical engineering from the University of Wyoming in 1992, 1994, and 1998, respectively. During his Ph.D. studies, he also obtained a graduate minor in statistics.

He is currently an Associate Professor with the Electrical and Computer Engineering department at North Dakota State University, where he teaches courses in signals and systems, digital signal processing, random processes, communications, controls, embedded systems, and others. His main research interests include digital and statistical signal processing, time series analysis, spectral and time-frequency analysis, array processing, real-time systems, and data adaptive techniques.

A Longitudinal Study of Student Performance in an Elective Applied Digital Signal Processing Course

Abstract:

In this paper, we describe multiple components of a longitudinal study of student performance in an elective applied digital signal processing course. First, we overview course organization and learning objectives. Next, we present data and analysis on student performance. We assess student preparedness over the period from 2002 to 2012 using, as a course pretest, the discrete-time signals and systems concept inventory (DT SSCI) developed by Wage and Buck. Pretest performance establishes a baseline for various student performance data, including hardware-based project grades over the same eleven-year period. Final exam performance data is presented from 2006, the year the course adopted a standardized final exam that is similar in structure to the DT SSCI but custom designed for this particular course. Particular attention is paid to the time period beginning in 2008, when students began field-testing a digital signal processing text coauthored by the course instructor. Since this is an elective course, we include select department-wide Fundamentals of Engineering exam data to provide an improved context for course data and analysis. Data is interpreted based on a combination of statistical analysis, student feedback, and instructor observations. Successes and difficulties in assessment methodology and process are discussed, including recommendations for improvement.

Introduction:

In 2002, the department of electrical and computer engineering (ECE) at North Dakota State University (NDSU) introduced a new three-credit applied digital signal processing (DSP) and filtering course. Available every fall semester as a senior-level undergraduate or introductory graduate-level elective, this course balances digital signal processing theory with practical application and implementation. On the theoretical side, the course covers:

- continuous-time (CT) signals and systems, including the Fourier series (FS), the Fourier transform (FT), and the Laplace transform (review material);
- design of CT infinite impulse response (IIR) filters, including standard frequency transformations and filter families;
- sampling and reconstruction theory, including a comprehensive treatment of aliasing and discussion of bandpass sampling, resampling, multi-rate processing, and quantization;
- discrete-time (DT) signals and systems with an emphasis on the discrete-time Fourier transform (DTFT), digital processing of analog signals, frequency response, and the z-transform;
- design of DT IIR filters through the techniques of impulse invariance and bilinear transformation with prewarping;
- design of DT finite impulse response (FIR) filters through window design method, frequency sampling, and frequency-weighted least squares, with emphasis on linear-phase FIR filters and an introduction to equiripple FIR filters;
- discrete Fourier transform (DFT), properties, and applications, including zero-padding, DFT-based linear convolution via the overlap-and-add and overlap-and-save methods, Goertzel's algorithm, the fast Fourier transform (FFT), radix-2 and mixed-radix FFTs, and decimation-in-time (DIT) and decimation-in-frequency (DIF) FFT structures.

On the practical side, the course covers:

- digital signal processing hardware overview, including digital signal processors (DSPs), field-programmable gate arrays (FPGAs), application-specific integrated circuits (ASICs), graphics processing units (GPUs), and others;
- hardware realization based on the Texas Instruments (TI) TMS320C6xxx DSK development kits (evolving from the ‘6211 to the ‘6711 to the current ‘6713 DSK), including overview of processor architecture and peripherals, fundamental use of the Code-Composer Studio (CCS) integrated development environment (IDE), and practical advice on C-language programming techniques;
- practical data conversion including introduction to analog-to-digital converter (ADC) technologies, real-world ADC errors, and ADC selection;
- real-time processing, event-driven (interrupt) programming, timing constraints, code and compiler optimizations, and system debugging.

Course grading is based on six-to-eight homework assignments (45%), a midterm (15%), individual projects (20%), and a comprehensive final examination (20%). Projects are individually selected by each student and must be implemented in hardware. In addition to traditional pen-and-paper problems, each homework assignment includes a substantial design problem that students are required to realize in hardware. In one homework assignment, for example, students must design and implement the highest possible order IIR filter to meet a given set of design specifications. In another assignment, students must design and implement an FIR equalizer to follow an equal-loudness contour for human hearing. In other assignments, students have implemented automatic gain control for audio signals, acoustic direction finders, non-linear quantization of speech, and other open-ended design problems. Hardware tasks, whether from homework or projects, must be individually presented by each student to the course instructor. Each student is allotted approximately 30 minutes for oral questioning during homework demonstrations and at least 60 minutes for project demonstrations, resulting in four or more hours of beyond-lecture contact time between the instructor and each individual student. The instructor uses one-page questionnaires during each homework check-off to provide consistency between students in structure and evaluation. Due to the individual nature of projects, standardized questions are not feasible during project presentation.

Due to the complexity of projects and homework problems, MATLAB is a necessary coursework tool; faced with the design of a triple-order multiband inverse-Chebyshev filter or a 1000th-order frequency-weighted least squares linear-phase FIR hearing equalizer, students quickly learn that MATLAB is highly preferable to hand- or calculator-only approaches. Numerous examples, both textbook and instructor derived, are provided to students throughout the semester to develop their MATLAB programming skills.

Initially, the course required students to purchase a textbook by Chassaing, first [1] and then [2]. Later, the course adopted a textbook by Welch, Wright, and Morrow [3]. The primary consideration in adopting these selections is that each of these books emphasizes hardware implementation using TI TMS320C6xxx products. Using an implementation-focused class text, students self-learn the DSK and TI’s CCS IDE much more rapidly and with much less boredom than the one-size-fits-all lecture alternative. Unfortunately, all three of these practice-oriented books do not provide the desired level of theoretical topic development for this course. Although TI data sheets and supplemental texts (such as [4], [5], and [6]) are recommended and available, students, especially between 2002 and 2007, obtained most theoretical material solely through class lectures. Based on student feedback and instructor observation, it became clear that students would benefit from additional (required) textbook resources that focus on the course's more theoretical topics. Somewhat coincidentally, the course instructor began coauthoring a DSP textbook in 2007 [7]; this manuscript, in ever-increasing stages of completion, was utilized as a required supplement for this course beginning in 2008. Drafts of all chapters were complete by 2011 and book production began in late 2012 following final negotiations with the publisher.

The course has two prerequisites: an undergraduate, typically junior-level signals and systems course that treats both CT and DT topics and a basic C-language programming course. Given the diverse backgrounds of students taking this elective course, the discrete-time signals and systems concept inventory (DT SSCI) developed by Wage and Buck [8] has been administered as a diagnostic pretest every semester this course has been offered. By reviewing student DT SSCI exam performance, the instructor can identify areas where student prerequisite knowledge is lacking. Now that the pretest has been administered for 11 consecutive years to over 120 students, pretest and other data provide an opportunity to longitudinally study student performance in our applied DSP course. It is this investigation of longitudinal data that is the central topic of this paper.

Analysis and Discussion:

To begin our analysis and discussion, consider student performance on the DT SSCI, which is administered as a course pretest during the first week of class. Since the DT SSCI is designed as a concept inventory, pretest performance provides a good measure of student preparedness. Furthermore, since the same 25 question exam is used yearly, longitudinal data trends are possible to investigate. Test booklets and answer sheets are tightly controlled to ensure that exam integrity is maintained from semester to semester.

Figure 1 displays pretest performance data for 122 students from 2002 to 2012. The dark gray bars show the class average for each particular year, with the sample size indicated at the base of the bar. Data variability is highlighted using light gray bars that extend one standard deviation in either direction about the mean. Solid x's locate individual student pretest scores beginning in 2005, when recording such data began. Two fitted regression lines (dashed) are also shown using data from 2002 to 2012 and from 2008 to 2012, respectively.

Several comments regarding Figure 1 are in order. Entering this class, students are not expected to know all material covered on the DT SSCI, so the somewhat low overall exam averages are not of any particular concern. Since the pretest was initially envisioned to provide anonymous diagnostic information, individual scores were unfortunately not recorded for the early years of 2002 to 2004. Furthermore, year 2006 data was also inadvertently lost or misplaced. More detailed performance data was recorded in later years when the potential for longitudinal analysis was more clearly recognized.

Despite the pretest data deficiencies, Figure 1 does provide interesting and useful information. The first thing to recognize is that student preparedness over the full 2002 to 2012 interval is essentially flat or perhaps slightly increasing. Most students take the same prerequisite signals and systems course, which is almost always taught by the same instructor as the DSP course.

The data in Figure 1 suggests that this prerequisite course is rather consistent in its preparation of students. With a fixed exam such at the DT SSCI, there is a potential danger than an instructor could artificially inflate scores by "teaching to the test" (something the course instructor consciously tries to avoid); however, the near-flat pretest data suggests this is not the case. At least for this class, the data does not support the instructor's and colleagues' occasional feelings that today's students are entering the course less prepared than a decade ago; this is a worthwhile observation for any seasoned instructor to keep in mind. Preparedness appears to trend more upward from 2008 to 2012, which could be due to better preparation, regressing over a shorter interval, natural yearly variations in student skills, improved teaching, or any combination of these factors. We revisit this apparent upward trend later when we investigate final exam performance data. Finally, while the 122 total number of students tested itself is quite satisfactory, sample sizes for individual years can be small to moderate, which can reduce confidence in some of the conclusions that are drawn from the data.

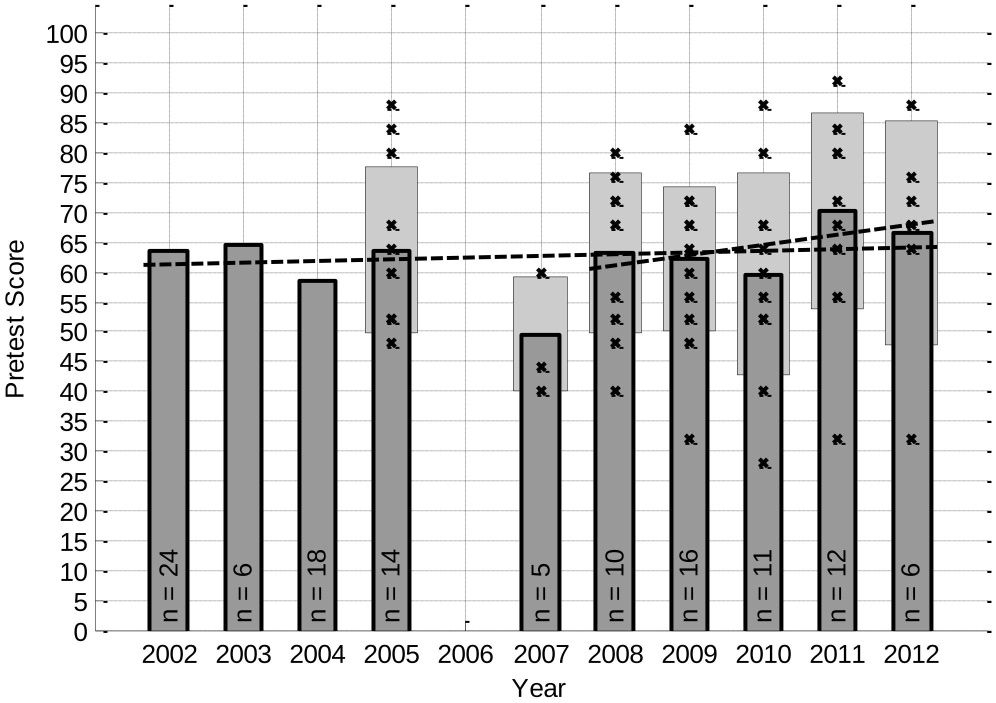

Figure 1: Pretest Performance Data

Motivated in part by the DT SSCI, a standardized multiple-choice final exam was developed for the course in 2006. Students are allowed a formula sheet, calculator, and two hours to complete this comprehensive exam, which contains 35 questions that reflect both theoretical and practical course topics, with an emphasis on the former. In addition to tightly controlling the final exam booklet and answer sheets, the ordering of final exam questions and answers are regularly changed to help ensure exam integrity from semester to semester. Table 1 provides a sample of six questions taken from the exam: questions 1 and 2 reflect more practical- and hardware-oriented topics, while questions 3 through 6 reflect more theoretical-oriented topics. As one can see from Table 1, question level and difficulty varies. For paper clarity, correct answers are highlighted in gray.

1) The TI TMS320C6713 is
a) an 8-bit processor	b) a 16-bit processor	c) a 32-bit processor	d) a 64-bit processor
e) none of the above

2) A large number of random voltages (uniformly distributed between $\pm V_{\text{ref}}$ are applied to a 3-bit ADC. Digital outputs [0, 1, 2, 3, 4, 5, 6, 7] are observed with respective relative frequencies $[\frac{3}{16}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{16}]$. The most likely cause for this anomalous output is
a) ADC gain error	b) ADC missing codes	c) ADC nonlinearity error	d) ADC offset error
e) none of the above

3) To sample $x(t) = \cos(2\pi 1000t)\sin(2\pi 2000t)$ without aliasing, the minimum sampling rate should be
a) 1000 Hz	b) 2000 Hz	c) 3000 Hz	d) 4000 Hz
e) none of the above

4) A channel has some magnitude response $|H(f)|$ and some phase response $\angle H(f)$. FWLS is used to design a digital FIR equalizer. It is most desirable to set the desired phase response equal to
a) 0	b) $\angle H(f)$	c) $-\angle H(f)$
d) $\angle H(f)$ plus a linear phase component	e) $-\angle H(f)$ plus a linear phase component

5) The circular convolution between [3, -2, 1] and [2, 0, -1] is
a) [8, -5, -1]	b) [8, -1, -5]	c) [5, 4, -7]	d) 6 stages
e) none of the above

6) An (N=60)-point mixed-radix FFT has a maximum of
a) 3 stages	b) 4 stages	c) 5 stages	d) 6 stages
e) none of the above

Table 1: Sample Questions from Standardized DSP Course Final Exam

Figure 2 displays final exam performance data for 51 students beginning in 2006, when the exam was developed, to 2012, the most recent year of data collection. As with Figure 1, the dark gray bars of Figure 2 show the class average for each particular year, with the sample size indicated at the base of the bar. Data variability is highlighted using light gray bars that extend one standard deviation in either direction about the mean. Solid x's indicate individual student performance data. Two fitted regression lines (dashed) are also shown using data from 2006 to 2012 and from 2008 to 2012, respectively. The 2008 to 2012 interval corresponds to the time period when students "field-tested" a DSP manuscript coauthored by the course instructor.

As with the pretest data, the final exam data of Figure 2 has some deficiencies. First, since a standardized final exam was not adopted until 2006, there is no meaningful way to include final exam data for the years 2002 to 2005. The data gap in 2007 occurs since only two students completed the class during that particular semester and the final exam was canceled to allow them additional time to work on class projects (no A's were awarded in the class that semester). Additionally, yearly sample sizes are each smaller by one to three students for final exam data compared with pretest data, a consequence of students dropping the course some time during the semester; these somewhat smaller sample sizes slightly reduce confidence in calculated statistics (mean and standard deviation) as well as regression lines. The multiple choice nature of the final exam likely contributes to the low final exam averages (no partial credit is possible), although low averages themselves do not affect the overall conclusions drawn from the data.

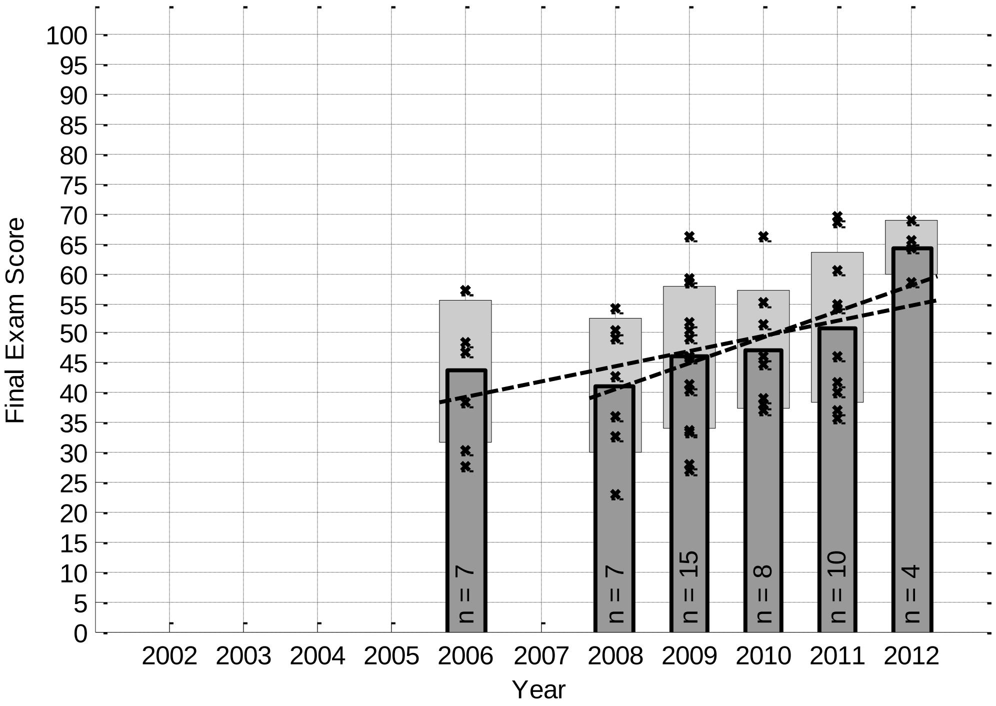

Figure 2: Final Exam Performance Data

Despite these deficiencies, useful information is again present in Figure 2. In particular, we notice that there appears to be a more definite trend of increasing final exam performance, particularly over the years 2008 to 2012. This increase is more pronounced than the pretest trends seen over the same period (see Figure 1), which indicates performance improvements are not just the result of having a better group of students. These yearly improvements in final exam performance are somewhat expected since during this time the course instructor utilized the supplemental DSP manuscript with particular hopes to improve student performance in the more theoretical areas. Improvements in the final exam roughly correlate with the number of completed chapters in the DSP manuscript available to students during each particular year. In other words, the increasing focus and resources on theoretical topics between 2008 and 2012 correlates in a natural and expected way with improving student performance on the final exam, which emphasizes theory more heavily than practice.

While not longitudinal, Figure 3 plots pretest performance versus final exam performance using data from 2008 to 2012. Unsurprisingly, there is a clear positive correlation between pretest and final exam performance. In other words, students who enter the class better prepared tend to do better in the class.

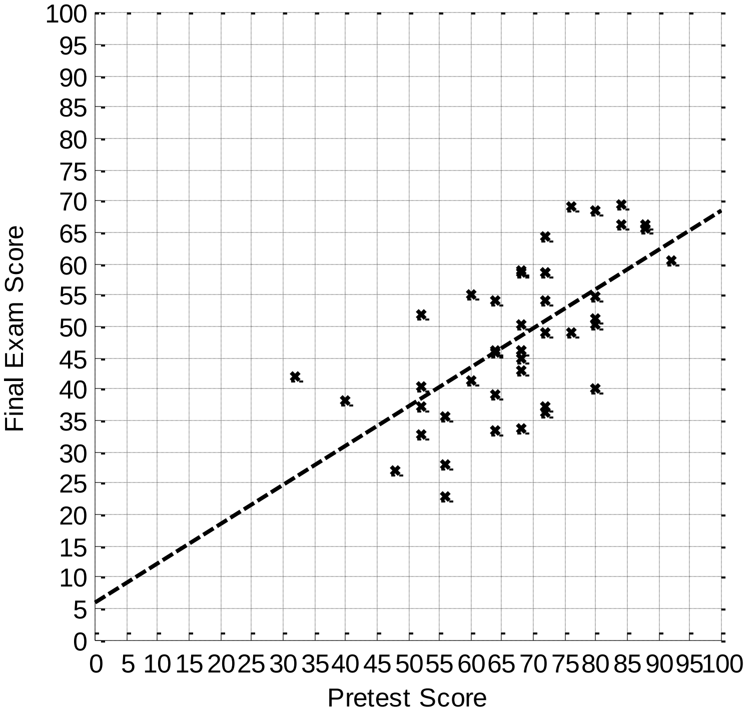

Figure 3: Pretest Performance Versus Final Exam Performance

While some of the improvement trends evident in Figure 2 are likely due to students having additional and improved resources available (the DSP manuscript), there is an additional factor worth consideration. Errors are bound to be present in any manuscript, particularly those in the early stages of preparation. In an effort to encourage students to not only read but to also critically evaluate the manuscript, the instructor awarded students with modest-level extra credit for each manuscript error found. Further, the instructor emphasized that student-identified errors would be corrected, thereby improving the text for later students. This latter point seems to be critically important. The instructor has provided extra-credit incentives in other courses to locate book errors (and errors there were!) with relatively low success. In this case, perhaps since student effort caused actionable change and improvement (that is, their efforts matter in a larger context), many students approached the manuscript review process with enthusiasm and together located 215 errors over the period 2008 to 2012. The left-hand plot of Figure 4 shows the distribution of student-located errors over 2008 to 2012. Relatively few errors were found in 2008, likely since not many chapters were available to students. In the middle years, a good number of errors were found, including one student in 2011 who found 55 errors. The final year saw a reduction in the total number of errors found, which is not very surprising given that the manuscript was at that point completed and fully reviewed, both internally and externally. Perhaps more interesting is the right-hand plot of Figure 4, which shows the relationship between the number of book errors found and final exam performance. While finding book errors is not a necessary condition to do well on the final, students who identified more book errors tended to do better on the final exam, as indicated by the regression line. This is not particularly surprising since students who are spending more time critically reading and evaluating the text should naturally gain deeper topical understanding. The trend is even more pronounced if the outlier data point in the 4-8 error bin is removed.

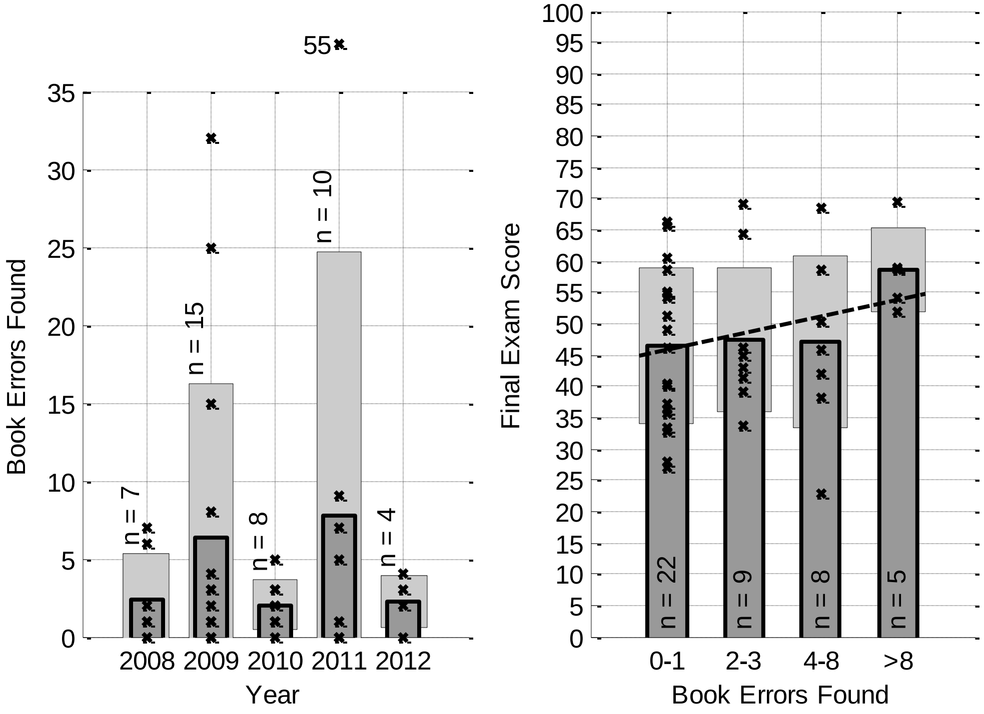

Figure 4: Book Error Data

Next, we turn our attention to course project performance. As described in the introduction, students are required in this course to complete hardware-based projects, which count 20% toward their final grade. Figure 5 presents project performance data for 101 students across the entire 11 year period, 2002 to 2012. In all cases project scores are based on a written report (approximately 25% weight), oral questioning (approximately 25% weight), and project difficulty, design, and functionality (approximately 50% weight). Since project topic is individually selected and uniform evaluation rubrics were not strictly employed over the entire 2002 to 2012 interval, trends for this data may be less reliable than earlier cases where identical evaluation instruments were employed across time. Keeping this in mind, we note that project performance appears relatively flat or slightly decreasing; trends appear more flat if the outlier data points in years 2009 and 2010 are removed (a good case can be made to remove the 2010 project score of zero, in particular, since that student stopped attending class and did not turn in a project; the 2005 project score of zero, on the other hand, should probably remain since the student submitted a plagiarized project). These trends suggest that students maintain relatively strong hardware performance over the years or that performance may be in slight decline, perhaps due to moderately increasing focus on theoretical over practical topics, lower student performance, increasing instructor demands, or some combination thereof. Additionally, at least for this class, the data (combined with pretest data) does not support the instructor’s occasional feelings that standards are more relaxed for today’s students compared to a decade ago, another worthwhile observation for a seasoned instructor. Additionally, the flat-to-slightly-decreasing performance trends in Figure 5 can support that the improvements in final exam performance noted earlier are indeed real, rather than a consequence of instructor bias (unconscious or otherwise) to artificially improve performance. The regression line over the 2008 to 2012 time period almost exactly matches the regression line for the entire 11 year period, suggesting that adoption of a new course manuscript did not particularly impact project performance. Even if some decrease in project score is related to the new manuscript and associated focus on theoretical topics, the payoff would seem justified: over the 2008 to 2012 time period, project scores decrease approximately 5% while final exam scores increase approximately 20%.

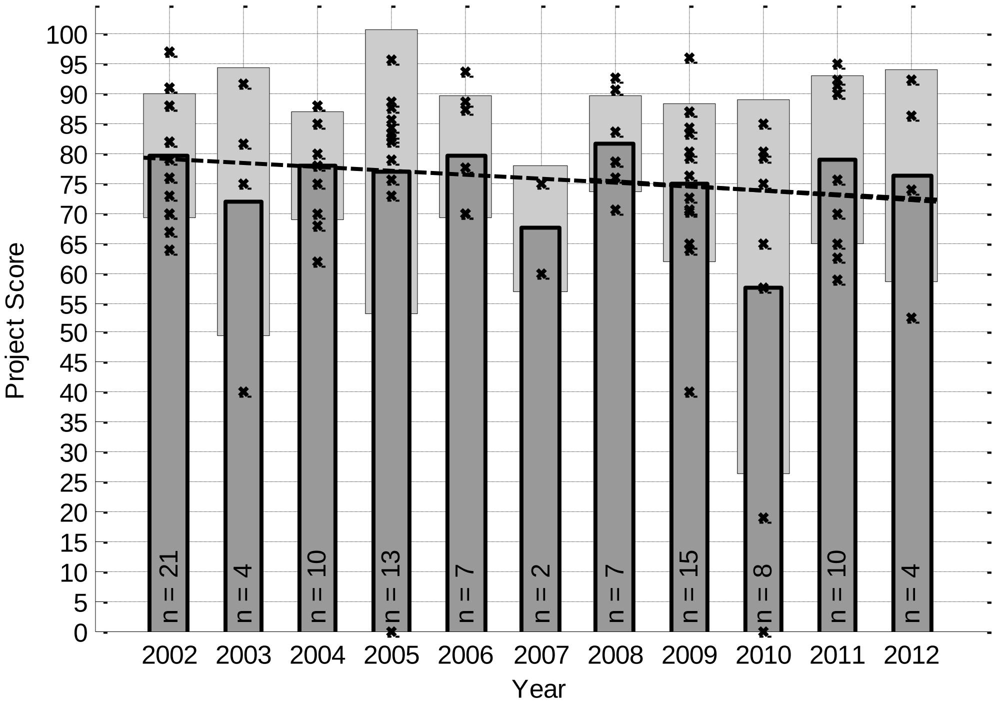

Figure 5: Course Project Performance

To conclude our discussion, we provide select performance data taken from the Fundamentals of Engineering (FE) examination administered by the National Council of Examiners for Engineering and Surveying (NCEES). Due to its relevance to course content, we limit our attention to performance on the signal processing component of the FE exam. Figure 6 presents the performance of 197 NDSU ECE students over the (available) period from 2007 to 2012 relative to the national average (the data of Figure 6 is used with permission from the NCEES, which generally considers FE exam data as confidential and proprietary). Before making any statements about the data in Figure 6, several cautions need to be made. First, the students represented in Figure 6 are taken from the entire NDSU ECE department; some of the students will have taken the DSP course while many more will not have taken the DSP course. Nearly all students, however, take the DSP instructor's required junior-level signals and systems course. Next, it is not possible to extract individual student FE exam data, which limits the type of comparisons and inferences that can be drawn from the data. While most FE exam data is from first-time test takers, it is possible some data comes from students retaking the FE exam. Lastly, although data is presented for 2007 to 2012, the data for these years does not directly correspond to earlier course-level data for the same 2007 to 2012 period. In other words, students taking the DSP course during any given year are not necessarily taking the FE exam during the same year.

Keeping these cautions in mind, we still obtain useful information from the data in Figure 6. First we see that NDSU ECE student performance in the signal processing area is relatively flat or, perhaps, slightly increasing, which is a trend that is quite similar to the pretest data provided in Figure 1. We see that NDSU ECE student performance is almost always above the national average, sometimes substantially so (on average, 7.34% above the national average for the 2007 to 2012 interval). Together, this suggests that consistent and effective signal processing education at NDSU. More importantly, however, is the fact that the FE exam is external and independent of the DSP course and its instructor. Thus, Figure 6 serves as a form of baseline data that, noting its consistency with Figure 1, strengthens the validity of our observations and conclusions. The data of Figure 6 does not suggest, for example, that the improvements in final exam performance seen in Figure 2 are artificially caused by the instructor "teaching to the test" or some other bias.

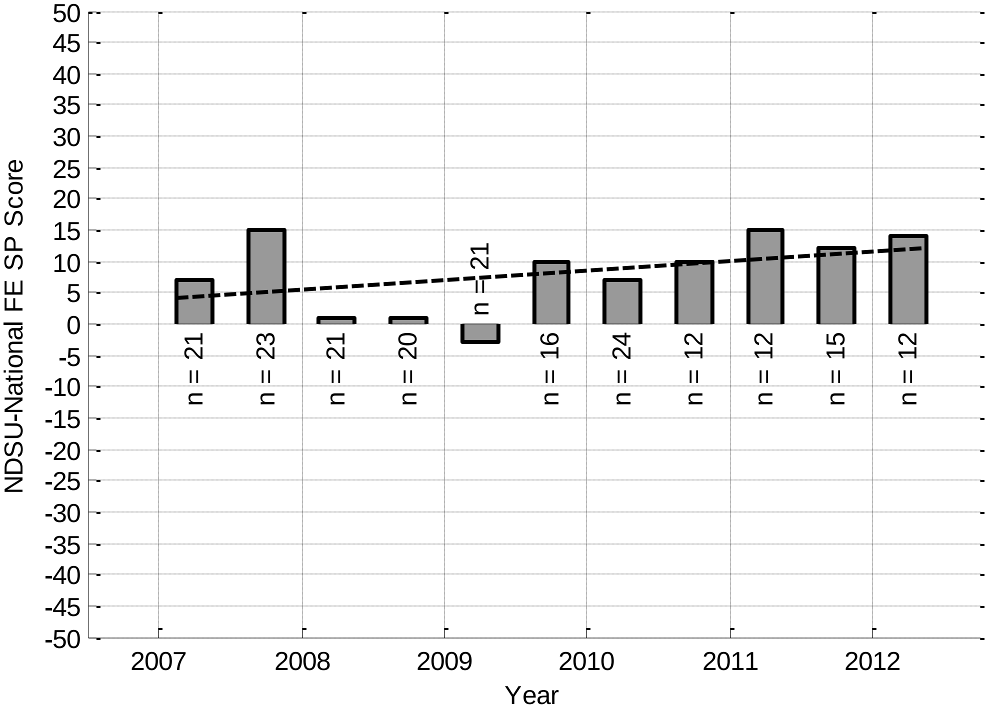

Figure 6: Relative FE Signal Processing Performance

Conclusions:

In this paper, we describe multiple components of a longitudinal study of student performance in an elective applied digital signal processing course, including DT SSCI exam performance, standardized final exam performance, student-identified book error data, project performance, and FE SP component exam performance. Taken across all measures, performance trends support that students receive consistent and effective signal processing education at NDSU. Data also supports a consistent level of student preparedness for the applied DSP course and modest performance improvements accompanying times when the course utilized a supplemental DSP manuscript. The data supports that through conscientious instructor efforts to better engage students, performance will improve. In this case, by emphasizing critical reading and evaluation of the DSP manuscript, student course performance, as measured by the final exam, generally improved.

Longitudinal performance studies depend on careful planning and execution. The current study would have been improved with more complete data recording (student-level DT SSCI exam performance during initial years, unintentional missing DT SSCI data from year 2006) and consistent use of detailed and fixed evaluation rubrics for all performance measures (somewhat lacking in this case for project evaluation). Still, as the current study shows, it is worthwhile to investigate performance trends even in the case of imperfect data. Longitudinal analysis can help validate teaching approach and methodology, can identify areas of weakness, and can provide unexpected results that may surprise even seasoned instructors, such as stable student course preparedness and consistent standards and performance throughout a course's history.

References:

[1] Chassaing, R., DSP Applications Using C and the TMS320C6x DSK, John Wiley & Sons, Inc., New York, 2002.

[2] Chassaing, R., Digital Signal Processing and Applications with the C6713 and C6416 DSK, John Wiley & Sons, New Jersey, 2005.

[3] Welch, T. B., Wright, C. H. G., and Morrow, M. G., Real-Time Digital Signal Processing from MATLAB to C with the TMS320C6x DSK, CRC Press, Taylor & Francis Group, LLC, Florida, 2006.

[4] Proakis, J. G. and Manolakis, D. G., Digital Signal Processing Principles, Algorithms, and Applications, 4th edition, Pearson Prentice Hall, New Jersey, 2007.

[5] Porat, B., A Course In Digital Signal Processing, John Wiley & Sons, Inc., New York, 1997.

[6] Van Valkenburg, M. E., Analog Filter Design, Holt, Rinehart and Winston, Inc., Fort Worth, 1982.

[7] Lathi, B. P, and Green, R. G., Essentials of Digital Signal Processing, Cambridge University Press, anticipated publication in late 2013.

[8] Wage, K. E., Buck, J. R., Welch, T. B., and Wright, C. H. G., "The Signals and Systems Concept Inventory," *Proceedings of the 2002 ASEE Annual Conference*, Session 1532, pp. 1-29, June 2002.