PEDAGOGICAL REFORMS OF DIGITAL SIGNAL PROCESSING EDUCATION

By

MICHAEL CHRISTENSEN

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2009

To my parents, for their unceasing dedication, exceptional example, and tireless love; and to those educators past and present for which I have had the fortune to be a spectator, for kindling an ardor of passion within me and always serving as a boundless source of inspiration

ACKNOWLEDGMENTS

With great appreciation for his passion for improving the state of engineering education, I thank my advisor, Dr. Fred Taylor, without whose generosity, humility, and patience this pursuit would have never proven possible. I am indebted moreover to Mr. Jim Kurtz for his invaluable insight and professionalism during my time as a research engineer, whose perspective on the engineering process and whose appreciation for the delicacy of management have served as a model to which I aspire. I finally recognize all of those whose unmitigated love and support I have received over the years; I only hope I have proven able to reciprocate it in as pristine a fashion as I received it.

TABLE OF CONTENTS

LIST OF FIGURES

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

PEDAGOGICAL REFORMS OF DIGITAL SIGNAL PROCESSING EDUCATION

By

Michael Christensen

May 2009

Chair: Fred Taylor
Major: Electrical Engineering

The future of the engineering discipline is arguably predicated heavily upon appealing to the future generation, in all its sensibilities. The greatest burden in doing so, one might rightly believe, lies on the shoulders of the educators. In examining the causal means by which the profession arrived at such a state, one finds that the technical revolution, precipitated by global war, had, as its catalyst, institutions as expansive as the government itself to satisfy the demand for engineers, who, as a result of such an existential crisis, were taught predominantly theoretical underpinnings to address a finite purpose. By contrast, the modern engineer, having expanded upon this vision and adapted to an evolving society, is increasingly placed in the proverbial role of the worker who must don many hats: not solely a scientist, yet often an artist; not a businessperson alone, but neither financially naive; not always a representative, though frequently a collaborator. Inasmuch as change then serves as the only constancy in a global climate, therefore, the educational system - if it is to mimic the demands of the industry - is left with an inherent need for perpetual revitalization to remain relevant.

This work aims to serve that end. Motivated by existing research in engineering education, an epistemological challenge is molded into the framework of the electrical engineer with emphasis on digital signal processing. In particular, it is investigated whether students are better

served by a learning paradigm that tolerates and, when feasible, encourages error via a medium free of traditional adjudication. Through the creation of learning modules using the Adobe Captivate environment, a wide range of fundamental knowledge in signal processing is challenged within the confines of existing undergraduate courses. It is found that such an approach not only conforms to the research agenda outlined for the engineering educator, but also reflects an often neglected reality: that the student who is free to be creative, free to err, and free to self-correct is emblematic of the profession – past, present, and future – to which he or she unwittingly aspires.

CHAPTER 1
THE STATE OF ENGINEERING

Obstacles of the Past

Though the role of the engineer has expanded as a function of socioeconomic progress and technological innovation over time, many of its overarching principles have been preserved. Beyond the commonly perceived sense of civil service and technical sophistication, the discipline of engineering has differed notably from pure science and mathematics in its end objective. The process of design, though rooted in the scientific method and often requiring a mathematical basis, has proven to be characterized ultimately as an art in its own right. Whether through the nuance of approximation or the motivation for abstraction, the engineer has always been implicitly tasked with framing such intellectual pursuits relative to some element of the common good – be it local, national, or global – in an effort to balance desire with capability.

For the engineer of generations past, this reality has been well captured through the prism of global conflict. During the Cold War, for example, the perception of existential threat was raised to a sufficient degree to warrant significant governmental influence. As described in [1], feats such as the launch of Sputnik I placed American technological prowess in question, leading to federal intervention through legislation such as the National Defense Education Act and through later incarnations like the Perkins Federal Loan Program, resulting in nearly quadruple the number of terminal engineering degrees during the height of tensions between the United States and the Soviet Union. Indeed, by this metric, provocative uncertainties such as the nuclear arms race and mutually assured destruction were in fact a boon to the engineering profession. Similar contextual descriptions arise in examination of electronic advancements during World War II, including modern radar systems [2], electronic countermeasures, communication theory

[3], and the fundamentals of signal processing that serve as the canvas for pedagogical development.

Though war serves as an excellent example of the advancement of engineering, the educational philosophy surrounding the profession has, until recently and irrespective of the state of international crisis, remained monolithic and homogeneous. In its inception, engineering was treated as an apprenticeship in most countries, with the United States being one of few sovereign powers to delegate its resources to engineering through the traditional university system by way of legislation like the Morrill Act [4]. More poignant yet, this description brought with it a broad and impacting implication: the engineering curriculum was to serve, in either model, as inherently and intentionally incomplete. All the while, left unfinished, pedagogical models were not a documented focus; the aforementioned disruptive technologies and the need for education of the highest caliber meant the student was tasked with a focus on the theoretical truth of a now practical reality. Questions of how such individuals learned or what best facilitated their learning were trumped by the discrete sensibility of what, according to the climate of the time, was needed to be known. As such, the now traditional pedagogical methods introduced and reinforced during this era, namely instruction by lecture and reference to text, still linger in the consciousness of student and educator alike, and only now is the profession as a whole gaining insight into the consequences of those decisions.

Challenges in the Present

In the present time, the engineering profession finds itself tasked with a multitude of responsibilities, not the least of which, one might argue, is the recapture of the identity to which it once laid claim. Though a persistent engagement in global conflict remains a source for innovation, the urgency to respond to it has yet to be received by the profession at large in the same fashion as the past; no archetypal technical adversary has arisen within the national consciousness, nor has any prototypical figure carried the proverbial banner of the engineering philosophy, which, in so doing, might galvanize a resolve within the populace for action. Indeed, the urgency of the modern generation is not fueled by an existential threat, but rather, as noted in [5], by an urgency for reform, stemming from a global marketplace and from which the populace of engineering is largely unaware.

In that respect, the American educational system faces an uphill battle both in recruitment and retention. The graduation of engineers has maintained a progressive decline since the end of the Cold War, with recent studies [6] indicating that more than half of all engineering students fail to complete their undergraduate curriculum and that the enrollment of women has declined systematically in each of the last ten years [7]. Governmentally funded reports [8] indicated the graduation of 70,000 engineers within the United States in 2004. By contrast, though common knowledge within the field has suggested an increase in the number of individuals educated in foreign institutions, research [9] shows that the gap is in fact smaller than once thought between the United States and emerging technological superpowers like China, India, and South Korea. Graduates in the latter nations, according to such studies [8], reach the hundreds of thousands, and statistics such as these do gain public editorial attention [10]. The contention regularly made within such literature is that the prohibitive American sociopolitical climate and improved economic conditions abroad are largely the cause for such convergence.

It is no doubt fair to ask, therefore, what the profession has done to curb this trend. Though there is an agreement that marketing to the younger generation is necessary [11] and that doing so as early as middle school [12] or within the early collegiate developmental cycle [13] is ultimately profitable, the scope of the problem reaches beyond any singular institution to tackle. Retrospective remarks from leading educators [5] unequivocally assert that governmentally funded efforts have been “studied to death” by agencies like the National Research Council, the National Science Foundation, and the American Society for Engineering Education. According to such far-reaching groups, the complication comes not with an ignorance of what to do, but rather with a failure to do it. Instead of a model of leadership [14], prominent figures in the process have observed a blasé reaction to the situation. The synopsis from [5] summarizes the challenge:

> ... when talking with individual faculty members, I sense a pervasive attitude that the ‘system ain't broke’, that it does not need to be fixed/changed. This attitude is not a resistance to change, but rather a sense of what we're doing is good, therefore it doesn't need to be changed.

Framed another way, no external agent has been present of sufficient influence as the government in previous decades [1] to induce an incentive to change, nor has a movement within the collegiate profession arisen to mandate it.

In fact, despite periods of significant growth [15], major funding institutions such as the National Institutes of Health have not even had their financial repositories keep with the rate of inflation [1]. To address these symptoms, arguments have been presented [4] both for and against the need for faculty to have an experiential, industrial basis for their academic philosophy. Doing so, it is argued, contributes to and partially validates their educational identity and entices the industrial sector to maintain a symbiotic relationship with the academic system. Since teaching responsibilities traditionally only compose one facet of a professor’s obligations, one is left to wonder how these alarming situations will be addressed, and moreover by whom. It is with recognition of addressing this challenge that the research field of engineering education has emerged.

CHAPTER 2
THE POSITION OF THE ENGINEERING EDUCATOR

Research in Engineering Education

Motivated by decades of research in educational psychology, the research field of engineering education currently exists in its formal infancy, spearheaded largely by the American Society for Engineering Education (ASEE) and its attendant publications. As discussed in [16], the agenda for this field has been compartmentalized into five major areas of investigation. The first field of *engineering epistemology* seeks to quantify the body of knowledge needed by the engineer, recognizing its time-varying structure. Closely related, studies in *engineering learning mechanisms* focus on the technical repertoire that is developed to conduct oneself as a life-long engineer. A third area of *engineering learning systems* examines the means, environments, and distribution methods by which engineering is gained. In so doing, a fourth area of *diversity and inclusiveness* tackles the challenges of fusing these concepts within a global framework, both for the purposes of equity and utility. Finally, *engineering assessment* addresses the testing criteria and measurements by which engineering knowledge is gauged. Though all five of these topics are coupled to one another, metrics in each area have been developed, and methods for conducting rigorous research [17] applicable to all such areas have aided in galvanizing the discipline.

Since most educational research techniques to date have employed some novel element within existing coursework, whether by the tacit or explicit deviation from traditional pedagogy, further seminal research [18-19] has revealed useful approaches in the development of curricula. Examination of educational coalitions funded by the National Science Foundation illustrated a systematic process by which success and productivity could be measured. Beyond initial development, the willingness to lead programs and implement them in a modular manner proved greatly beneficial in convincing fellow faculty to follow suit. Development alone, however, fostered stagnation, resulting in the need for intrinsic mechanics to promote and retain utility. Though this description is arguably general, the spirit of its structure is a reflection of the challenge currently within the pedagogical environment and has served as a motivational source for this research.

That having been stated, numerous transcendent challenges remain for the general engineering educator. While the tools of how people learn in pre-collegiate levels have been well developed and continue to serve as a cornerstone to ongoing research [20-21], the vision for the engineer of upcoming decades [11] is influenced by the proliferation of technology within the classroom, whether as an intended and independent injection of pedagogy or as a response to its embrace within the general culture of which students are a part. The institution of student interactivity with technology has become increasingly commonplace and, at least within electrical engineering, an embedded staple of the curriculum [22-23]. For the signal processing student in particular, programming environments such as Matlab continue to serve as an invaluable resource for the visualization and computation of tasks once deemed infeasible in previous generations of instruction. The capacity to support fluency in such topics, however, is often at odds with the limited time available to satisfy curricular expectations for a given course. As of the time of this research, the university system as a whole, whether influenced by accreditation or industrial committee, has not uniformly accepted the burden to introduce support for such types of programmatic proficiency by way of courses serving solely that end. Nonetheless, the demand for technology in the classroom remains paramount, along with assurances that its inclusion is not subject to the traditional pitfalls of unconstrained media.

To that end, the aforementioned desire for the modular deployment of learning systems is juxtaposed with the sobering reality that students are both capable (oftentimes more so than their instructors) and willing to circumvent the spirit and intent of a learning system in favor of an improved result. Whether through exploitation within programming objectives [24] or in efforts to circumvent the opportunity [25], the stark truth of academic honesty cannot be overlooked. Indeed, an alarming number of students within engineering have admitted to cheating during the course of their education, and strong statistical correlation has been found between academic dishonesty and other deleterious behaviors [26]. As a result, though the infusion of technology can yield extensive benefit, a subtle but necessary condition of security also applies. Educational systems deterred by or motivated to avoid these challenges oftentimes further consign the most representative elements of such pedagogical evolution to participation in so-called capstone sequences [27]. In so doing, many of the ultimate goals of the process have been shown to be lost in the proverbial ether, pointing to a larger reality.

In that regard, though the technical hurdles arising within such an educational overhaul can be significant, fundamental misconceptions about the worth of given approaches are arguably more challenging to dispel. Technical expertise has proven time and again to be an insufficient measure of educational aptitude [20,28], and yet the proverbial "null hypothesis" to simply make coursework less challenging for students neither guarantees nor implies perceived benefit by the student. In fact, the contrapositive is also a myth warranting dismissal: demanding and extensive coursework, which readily translates in the parlance of a student to a "difficult" professor, does not in fact systematically translate to poor evaluation of the course or its administrator [29]. Moreover, recognizing disparities in learning styles between the educator and the student [30], whether as a result of generational gaps or personality differences, can have an immediate impact on the manner in which assessment is conducted [31] and ultimately can serve as a direct indicator of how likely a student is to be retained in the process [32].

Educational Research and Digital Signal Processing

Educators have shown themselves aware of these realities to varying degrees, and learning paradigms have been constructed accordingly to address weaknesses and enhance perceived benefits. In the context of engineering education, a nomenclature unto itself has been developed to describe such compositions [33-35], though all invoke some measure of problem solving. Indeed, educational media for digital signal processing (DSP) with this objective in mind have come in an array of forms. For example, the development of summative assessment methods for enhancing proficiency with Matlab has been a focal point for prominent educators in signal processing [36], and the emergent prominence of Texas Instruments devices has brought with it a marketable demand for educational efforts to aid in developing foundational and application-based expertise in DSP principles [37-41]. However, be it due to the relative infancy of the engineering education process or simply the ignorance to it, few within electrical engineering or DSP in particular have attempted to frame such efforts in a manner from which a more general understanding can be derived. By contrast, this research seeks to expand upon the author's previous trials in signal processing education [42] and to posit a number of critical epistemological considerations.

CHAPTER 3
DECISION-ORIENTED LEARNING

An Argument Favoring Error

Engineering has been often characterized as iterative, developmental, and responsive to unplanned shortcomings. In this way, though engineering designs become increasingly driven by detail, they are equally subject to an increasing tolerance of and compensation for more catastrophic circumstances. When framed in this light, one might rightly argue that the decision process for the engineer is multi-faceted; having an understanding of successful operation is conditioned as much upon knowledge of why alternative techniques are unsuccessful. Indeed, arguments have been made both within scientific journals [43] and in the mainstream consciousness [44] that the greatest innovation stemming from scientific endeavor (of which engineering is a part) is contingent upon the necessity of creativity and the willingness to make mistakes. When pedagogy is framed in a pejorative “academic” tense, however, such opportunities are inextricably lost and the traditional summative assessments methods used no longer measure or reflect one’s preparedness for a career, but rather merely the quality of recitation. Said more simply, the student who is without the expectation to make decisions and experience their consequences – be they beneficial or detrimental – is void of a critical element of the engineering process, without which their pedagogy is fundamentally incomplete.

This belief has served principally as the motivation for this research, producing several operational objectives. For the DSP engineer in particular, this mandates the framing of real-world problems in a fashion that replicates a typical design condition. First, the deployment of these learning mechanisms, henceforth referred to as modules, must be sufficiently accessible to the student at his or her current level of topical knowledge while still challenging and expanding upon it. In so doing, they must moreover house sufficient latitude to tolerate faults in the learning process and support remediation prior to the completion of the effort. Finally, to ensure consistency and induce increased comfort with the activity, student exposition to such modules must occur with sufficient regularity.

With these ideals in mind, investigation into available outlets for this notion of decision-oriented learning exposed an additional constraint. Previous efforts to engage students more directly [45-46] were met with hesitancy to participate with a traditional adjudicative figure. Believing this was at least in part a result of social inhibition and simultaneously recognizing the aforementioned need for the infusion of technology, focus turned to the use of software development as an equivalent means of communication. Previous efforts have been developed and assessed with respect to software [22, 36, 47-48] and a similar learning paradigm has been explored in other areas of engineering [49], but such pursuits have largely relied on custom-designed systems to serve that explicit end. In examining alternatives with a wider appeal, the use of the Adobe Captivate development environment for learning module creation demonstrated considerable potential profit. Used by industries for traditional training exercises in the workplace and currently in its third version, the functional framework of the system lends itself to concoct a unique experience as a function of constrained options through the principle of branching logic. This is achieved on the application level by publishing the Captivate design into one of several possible file formats, including HTML, Shockwave Flash, and a standalone executable. In addition, the final product can be viewed at high resolution while containing features like embedded audio and full-motion video to support a particular illustrative task. Though developmental time for a given module is a direct function of the desired sophistication – all modules for this research required no more than forty hours each to support narrative audio and video – the learning curve for the use of Captivate is extremely limited, aiding in the effort for future distribution to other faculty.

To test the efficacy of the learning philosophy as well as the modules themselves, three major topical objectives were pursued and summarily introduced to students enrolled in signal processing courses in the Department of Electrical and Computer Engineering at the University of Florida [50]. In each module, the student is placed in a scenario that requires them to participate in a fashion similar to that of a working engineer. The modules challenge the student to make immediate use of the knowledge he or she has recently acquired in traditional lecture, examining topics in filter design and DSP hardware and software, motivated by discussions in [36, 38, 51].

Filter Design Learning Modules

The devised scenario places the student in the position of the singular DSP engineer in a multidisciplinary team, working to design and implement a rudimentary transponder technology for a collection of unmanned aerial drones in a potentially hostile environment. As shown in Figure 3-1, each drone is assumed to transmit a sinusoidal tone to a base controller at a given frequency and for a specified length of time. Data from multiple drones is collected and filtered at the controller. A Fast Fourier Transform (FFT) is performed on the filtered data, followed by tone detection and frequency shifting on each incoming signal for verification and security, which is viewed by the learner as a black box. Nonetheless, once complete, a valid input tone then undergoes an Inverse FFT (IFFT) and amplification before being transmitted back to the drone. This process is repeated multiple times to decrease the chance of false positives.

To increase system complexity, the fictive team has been tasked with adequately suppressing a jammer in the filtering phase to decrease the probability of a false detection. This entire procedure is further encapsulated in a real-time constraint to eliminate the possibility for input overflow. With this information in hand, individual learning modules are provided to students that address specific design challenges, first with a finite impulse response (FIR) approach.

Finite Impulse Response Module

The first module begins by addressing sample rate restrictions as a function of the bandwidth limit indicated. Four options are provided: an aliased sample rate, critical sampling, extreme oversampling, and satisfactory oversampling. In the event students respond with any of the first three conditions, the module branches off to a remediation phase in which their choice is questioned. The ‘critically sampled' case is given as an inadequate reason due to the need for a non-causal filter, and the ‘extreme’ oversampling case is deemed improbable due to the real-time constraint, which the resultant filter order would likely violate. In the event these reasons are not recognized, a final layer of remediation is provided, at which point the appropriate answer is given as a basis for continuation.

Students are then tasked with managing a multitude of drones simultaneously. This is equivalently cast as selecting an FFT length in order to determine the satisfactory frequency resolution, much as in Figure 3-2. In contrast to the first decision layer, however, two ostensibly correct choices are provided along with two errant choices. Should an errant choice be selected, remediation is provided that indicates either of the two remaining choices is satisfactory, with the student left to select one. This decision influences the branch down which the remainder of the module will follow.

To address potentially disruptive effects, the student is provided with the received jammer and signal powers. In so doing, they are tasked with designing a symmetric, linear phase FIR that will yield a specified signal-to-jammer ratio for an out-of-band tone jammer, which is indicated to be sufficient for the tone detection process and is illustrated in Figure 3-3. An incorrect solution digresses into an explanation of frequency response and with it a graphical framework of the problem.

With most critical design parameters in tow, the module then aims to simulate incremental collective knowledge gains (i.e. the time-elapsed design process) by injecting information about the system speed requirements via other members of the project team as well as presumed personal research. Specifically, quantitative timing estimates are given concerning the tone detection and amplification processes as well as the documented benchmarks for FFT and IFFT processing, as shown in Figure 3-4. The student is then tasked with determining the amount of time remaining in the processing interval allotted. Four options are once again provided, three being discrete values and the fourth suggesting that insufficient time is available.

In this way, the prior decision concerning FFT length is amplified. Namely, a student who chose to use a longer length FFT, under the arguably vacuous assumption of the benefit to higher resolution, comes to find that the process cannot be completed coherently when accounting for the resultant increase to delay, whereas the student selecting a satisfactory but shorter length can in fact succeed. Should a student not acquire such an observation or alternatively select an inappropriate estimate of headroom, remediation is once again provided to demonstrate the contributing factors to the timeline.

The module then concludes by reinforcing the major themes of the exercise and extending them to the working environment. Indeed, though the student was in some sense led blind down a portion of the proverbial path, this is indicative of an engineering design which either is evolving or ill-defined, both of which occur in practice and are representative of the incremental and iterative process. Hints are then given concerning the next module, which aims to tackle the problem in a new light.

Infinite Impulse Response Module

Once students have received preliminary instruction on IIR systems through traditional pedagogical means, the second module in the course is deployed. In it, several notable changes are made to the design plan to reflect discussions within the team and the contract supervisor. Specifically, the concept of the jamming system is revisited, wherein the argument is made (by proxy) that a broadband jammer would be more probabilistically effective to combat drone tones at presumably unknown frequencies, particularly those potentially within the filter pass band, as illustrated in Figure 3-5. Cognizant that a power-limited jamming system will be inhibited by increased operational bandwidth, the proposition is put forth that each drone could utilize a larger frequency range while simultaneously transmitting side band tones as a primitive countermeasure. In so doing, should the jammer, though described as broad band, in fact act more like a pulse across the chosen range, then such side bands could produce an artificial (i.e. desired) hit without disrupting communication. From this idea the module expands.

The student is reminded that the choices made and information provided in the previous module impact the current condition. Specifically, timing estimates for the ‘black boxes’ were conditioned on specified sampling rates and lengths, with a change to rate likely having a larger impact than a change to length. Therefore, having to assume an invariant sample rate, the student is tasked with selecting an operational bandwidth and sideband spacing – the latter equivalent to a message frequency – that will give the jammer the greatest challenge. As this is achieved whenever adjacent sidebands overlap, two seemingly suitable choices are provided, one utilizing the entire unaliased bandwidth and one that does not, along with two incorrect choices. Should a student make an erroneous choice, the decision conditions are outlined in greater detail and only legitimate options are provided.

With a selection in place, the module then addresses how to combat the cumulative effects of the drone side bands. The argument is made that such side band tones would ideally wish to be suppressed in the filter block so that the tone detection, IFFT, and amplification processes do not receive unnecessary contributions. Moreover, lower memory requirements would promote the compact design that the sponsor is seeking. To that end, and without mandating an explicit design technique, the student is left to determine the minimum amount of memory that would be required to store the filter coefficients to achieve such suppression.

Irrespective of the student's decision, the module then breaks from the 'scenario' mode of Captivate and switches to a 'training' mode utilizing full-motion video recording. At such time, the narrator rehashes the problem at hand and makes use of Matlab to illustrate the design challenge. The concept of an IIR comb (notching) filter is introduced using the Filter Visualization Tool as a viable solution to the problem, as shown in Figure 3-6. Both plausible alternatives are addressed, for which it is demonstrated that failing to use the whole unaliased bandwidth would result in an unsatisfactory design. In contrast, making such use not only would prove viable, but also would simplify an IIR methodology to an FIR result. Moreover, the student is led to conclude that significant memory savings can be had by this approach due to the need to store only a small fraction of the redundant filter coefficients. Once complete, the student is returned to the more familiar scenario environment and left to continue the exercise.

Having established a fixed resolution and sample rate, a final challenge is presented in the form of timing estimation. The student is faced with assessing the latency resulting from successive drone transmissions. This process becomes feasible to calculate due to a constant group delay resulting from the filter simplification. As in previous instances, a student failing to select an appropriate estimate is reminded about the relationship between frequency resolution and sample length and how it, along with the filter delay, impact the timing specification.

The module concludes with a synopsis of what was learned. The concept of simplification through generalization is reinforced, and practical applications of the study are introduced, eliciting the similarities between the project description and technologies such as frequency modulation and spread spectrum techniques.

Digital Signal Processing Module

The third learning module aims to realize the ideals of digital signal processing by placing the student within the confines of a multidisciplinary team once again, but now with a more far-reaching aim. The functional concepts of real-time signal processing and radar signal processing in particular are introduced under the guise of remote surveillance of endangered species. Students are introduced to the problem at large by eliciting the need from their manager for an end-to-end signal processing solution that can operate within arctic climates in an effort to capture migratory patterns of creatures otherwise elusive to humans. In view of the significant posed interest associated with the effort, the student is immediately tasked with developing both a suitable hardware platform as well as a robust detection methodology to support a panoramic sensor array. In the former case, considerations about price and mass production are put into the framework of the scenario without further immediate elaboration, while in the latter case the sensor architecture is illustrated identical to Figure 3-7. As shown, four sensors form the array, each sensor sweeping a quadrant (or swath) of space by way of four narrow, steered beams, which the student is told responds to the presence of irradiated energy. This process produces sixteen refined position estimates within the circular viewing position.

Hardware Decisions

Since the software is irrevocably conditioned upon the hardware, the student begins by a survey of the Texas Instruments High Performance series of DSPs in a manner consistent with what a working engineer would find by way of product guides, informative websites, and pricing charts as of December 2008. First, the C6474 is introduced as the top-of-the-line DSP and critical specifications such as its parallel processor architecture and peripherals are advertised to potentially entice the student. At the same time, its limitations are exposed, namely in its lack of modularity and customization and, due to its novelty, its particularly high cost. Motivated by these shortcomings, a second processor, the C645x series, is introduced as being a compromise of many ideals. Though not as fast as the 6474 device, its numerous options are articulated, including a wide range of available memory spaces, operating regions, and processor speeds. Because of such latitudes, however, the price remains competitive with the 6474, and only in mass production would significant fiscal benefits be possibly observed. To once again play on these pitfalls, the C641x is presented as the lowest end alternative that might satisfy the design conditions. Citing its versatility to be competitive with the 645x in processor speed and operating environments, the only listed limitation is in its fixed memory space.

From these descriptions, the student is tasked, albeit naively, with making a preliminary selection of hardware. Should he or she select the high performance 6474, the module branches off into a separate discussion based on how, due to a lack of multiple operating ranges, the 6474 would fail in an arctic climate. Should the student attempt to reason principally on the basis of cost and therefore select the C641x series, another branch is formed that performs a more detailed analysis of the claims to versatility. Indeed, for arctic climates, a maximum operating frequency of only 600 MHz is given as the only supported clock rate, which, when compared to the alternatives, proves less than alluring. In either case, or if the student did initially select it, all discussions then return (albeit on independent paths) to the 645x series chip. Because of the need for low temperature outfitting, the estimated cost per unit of $300 is in fact more than even the high performance 6474, but in view of no other options, the student is led to believe that it is the best decision.

With a general make of processor selected, the module then branches to inform the student on the differences in product models, both in price and performance. Though both the 6454 and 6455 models use an internal 1 GHz clock, the 6455 doubles the amount of available internal memory and supports inter-processor communication by way of Serial RapidIO. In view of the limited conceptual framework of the software design at this juncture in the design, no decision is required of the student at this stage and instead business issues are placed into focus.

Clock Division

As a narrative interlude, the fictive manager in the scenario then asks the student to place his findings into a fiscal context, namely with projected memory and timing budgets. Indeed, a detection algorithm that produces numerous points of interest, though seemingly robust, could, in later stages, prove to be detrimental in cost due to the need for additional memory. With this reality placed in context and the sensor modality emphasized once more, the concept of array signal processing is introduced as a spatial-temporal problem, for which the two-dimensional FFT is selected to serve as a worthwhile mechanism for potential species identification. The student is then steered to consider possible radix selections for the production of such FFTs by way of TI documentation, for which only powers of two and four are shown as supported.

With a firm memory restriction of eight kilobytes in place for each full scan of the field of view put in place to guarantee that the system can run for a significant duration without the need to purge its contents, the student is then left to estimate the data collection timeline. Specifically, he or she must select a an integer divisor of the 1 GHz clock rate as a sampling rate that will allow the maximum number of 16-bit samples to be collected within each region and within a two microsecond time frame. Should the student select an unsatisfactory divisor, separate remediation branches are introduced that explain how such a memory budget per scan implies a 0.5 kB allocation per region, which, at two bytes per sample, implies a maximum of 256 possible samples. The remaining algebra is illustrated along with block diagrams from the TI documentation to reinforce the decision and the conceptual operation of the system moreover, after which the processing timeline associated with the 2-D FFT is investigated in more detail.

Processing Timeline and Device Selection

Though the FFT mechanism ostensibly serves as a graphical means for recognition, the student is reminded that it alone cannot serve as an automated detection methodology. As a first blush response to this, the suggestion is presented that all elements within the 2-D FFT result could be compared to a static threshold to declare a point of interest. Being provided the instruction cycle length for comparisons and further restricting the processing sequences to be serial, the student is then asked to estimate a timeline for both the collection and detection phases of the signal processing for the entire panoramic field of view.   Should a selection be made that challenges the expected response, the scenario branches into a discussion on the technique for calculating the 2-D FFT, ultimately converging upon a final estimate.

A proverbial wrench is then thrown into the mix by asserting a subtle reality to the problem: not all detections will ultimately be of interest. All the while, a mandate has been put forth to operate the system interminably with transmission of relevant information each second. The student therefore is left to realize that some method of pruning is necessary, be it by an increase to the sophistication of the detection process itself or by a hard limit on the number of accepted detections. However, in so doing, the narration argues, all of the necessary criteria are in place to make an informed decision concerning a processor model, as shown in Figure 3-8.

To that end, the student is asked to select a DSP that will satisfy the aforementioned expectations while restricting detections to five per swath (e.g. quadrant). Such inquiries are solicited by way of estimating the amount of remaining available memory, with the notional specification of "negative memory" implying the need for a more robust processor (6455) instead of its inferior counterpart (6454). Errors made within this period are afforded immediate remediation in view of the closure of the hardware topic.

Software Decisions

To address the implementation of the detection algorithm, the focus is then turned to the production of software algorithms in Matlab, with the aim that their architecture would prove suitable for porting directly to a DSP. To test student knowledge of programming principles and their relation to the target development environment, three unique implementations are compared and contrasted. In the first example, the implementation is flawed in that it fails to collect the maximum number for each region, but rather will select as many samples per region until the maximum is reached. In so doing, the possibility for failed detections increases based upon the concentration of information relative to the static threshold provided; data from the first swath processed could, by this faulty implementation, eliminate the inclusion of any subsequent detections. Should this reality fail to be observed, the associated code is explained in detail to demonstrate its limitations.

A second implementation brings into question the worth of a static threshold in the presence of sufficiently high strength, periodic data generated by random integer permutation, shown in Figure 3-9. The underlying consequence of such a circumstance, one might rightly argue, is the unnecessary expenditure of time associated with the discarding results that could be physically attributed to leakage or other deleterious effects. Even in the presence of ideal circumstances, phenomena such as bipedal and quadrupedal gait are able to induce significant

Doppler shifts when receiving “irradiated energy” as previously described. With these things in mind, the student is tasked at this layer of the scenario with estimating the number of discarded points of interest under the assumption of a now correctly working parsing algorithm. Unlike the first implementation making use of three-dimensional vectors, the second format utilizes the structure data type in Matlab along with an indexing system more akin to that of a target development environment such as C. As an inhibiting factor, however, internally supported algorithms for sorting are used to provide contrast. In this way, student competence is tested both programmatically and conceptually. With that, though remediation is provided for errant decisions, separate branches are nonetheless instituted to reinforce the subtleties described, principally to serve as a foundation for the radar signal processing module to follow.

The final form of the detection algorithm mimics the target environment most closely by separating the 2-D FFT functionality into individual steps. Furthermore, rather than making extensive use of dynamic vector nullification, iterative processes are utilized to illustrate the need for a more modular transfer of code. Though functionally identical to the preceding form in both input and processing, the student is tasked with determining the effect of changes to the static threshold with regard to how many quadrants are reviewed and what type of values are retained.

Module Conclusion

Due to the increased problem complexity relative to preceding modules, a summative discussion is instituted to conclude the module, chiefly to elaborate more upon the idiosyncratic aspects of each Matlab implementation. Particular emphasis is placed upon demonstrating what types and kinds of functions are typically supported within both the TI DSP and Matlab development environments, including variable instantiation, two-dimensional vector concatenation, and element nullification. A salient point, however, is asserted to the student: the contrived scenario arguably does not act as a real-time system. This is motivated chiefly by the observation that the amount of time to process data was well in excess of the time required to collect it. The implication of this reality is that observations would be temporally (and possibly spatially) discontinuous, which, when needing to observe instantaneous motion such as with endangered species, could prove catastrophic. With such a description, the student is reminded of the necessity for clear problem definition above all else.

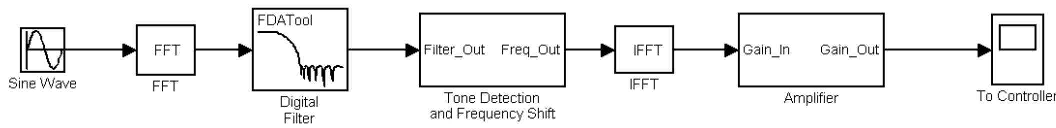

Figure 3-1. Filter design module transponder block diagram.

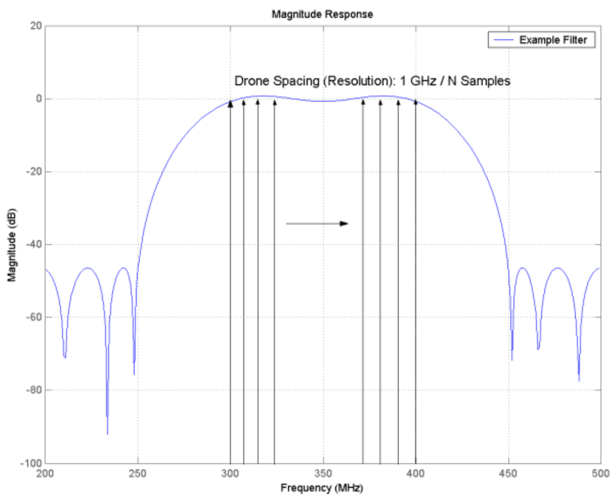

Figure 3-2. Filter design module resolution clarification.

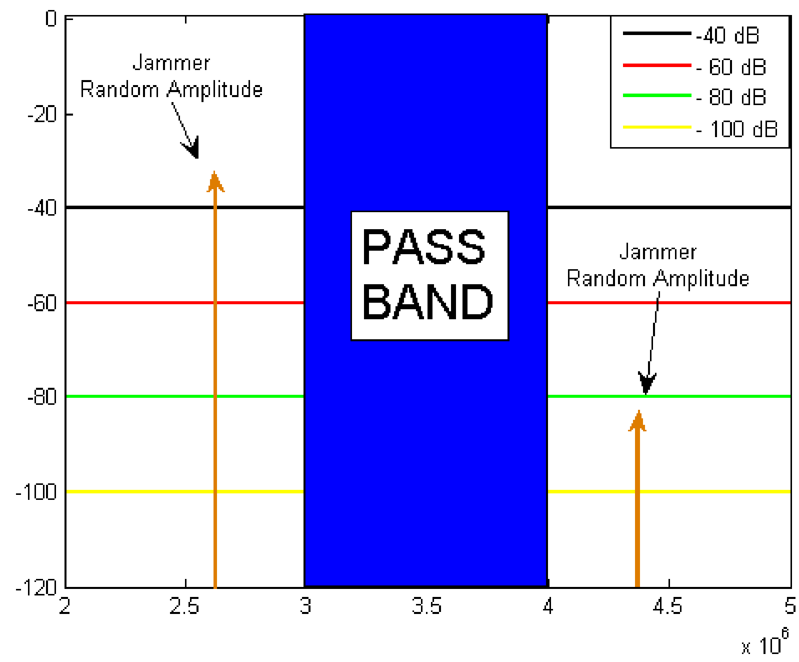

Figure 3-3. Illustration of out-of-band tone jammer.

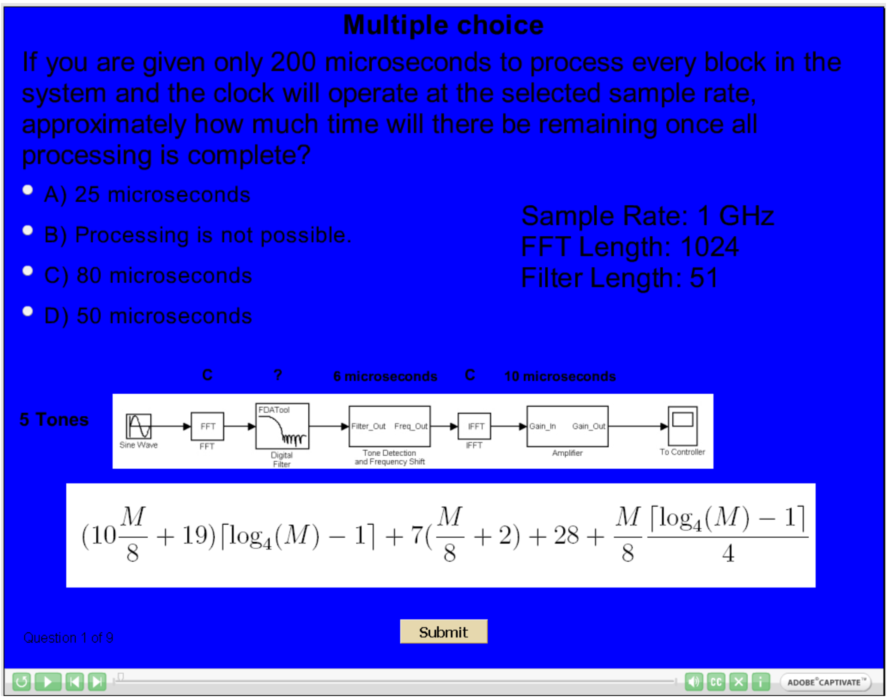

Figure 3-4. FIR module timing calculations.

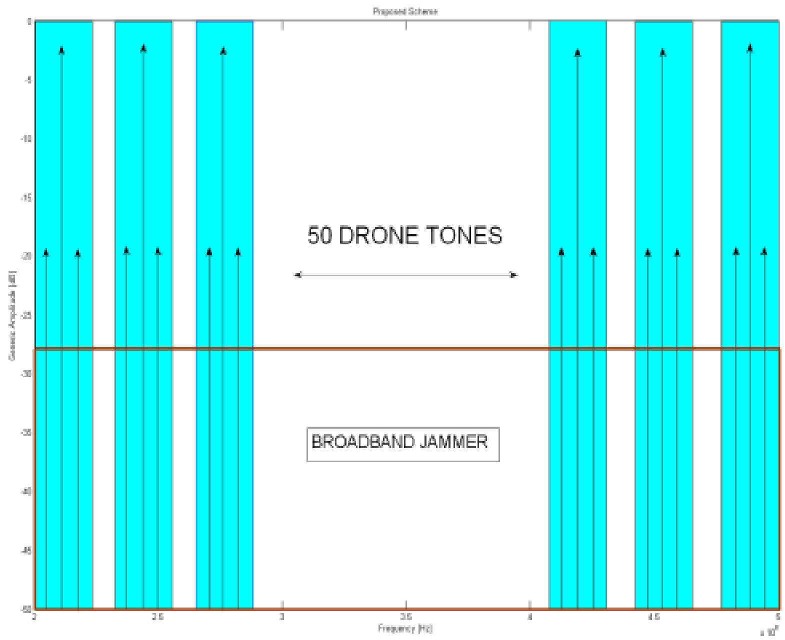

Figure 3-5. Broadband jammer illustration.

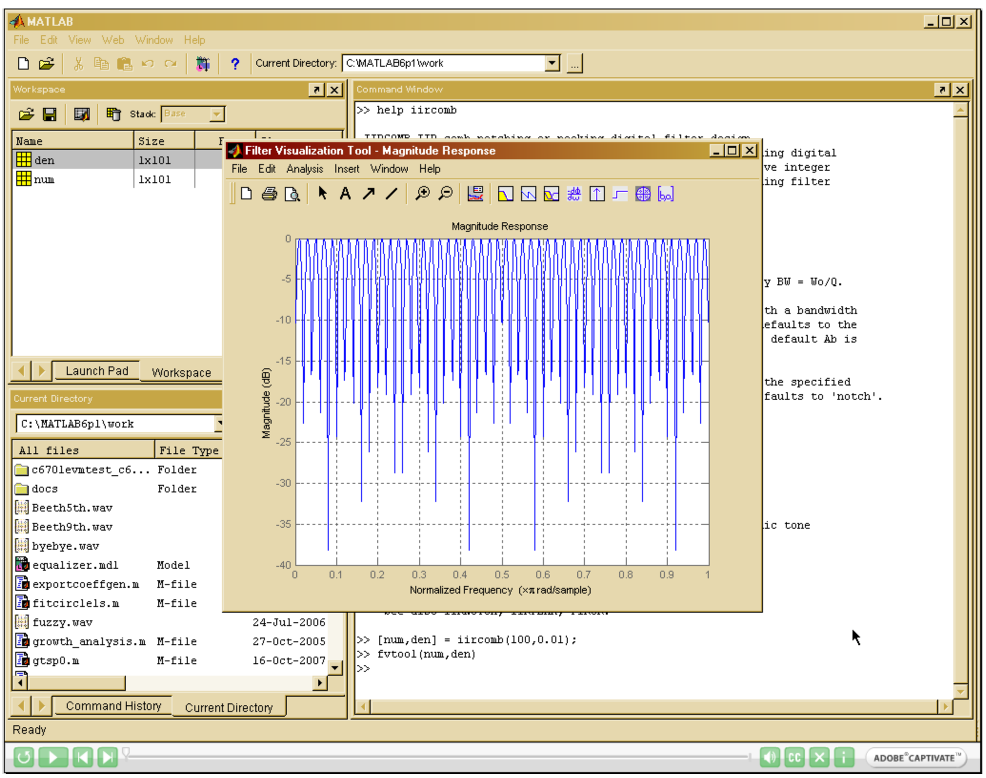

Figure 3-6. Matlab full-motion video capture.

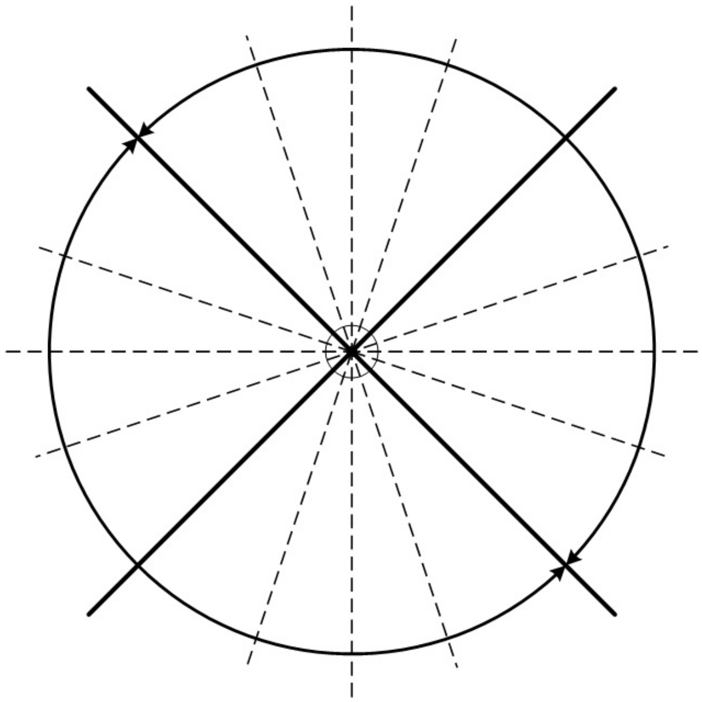

Figure 3-7. Panoramic sensor model.

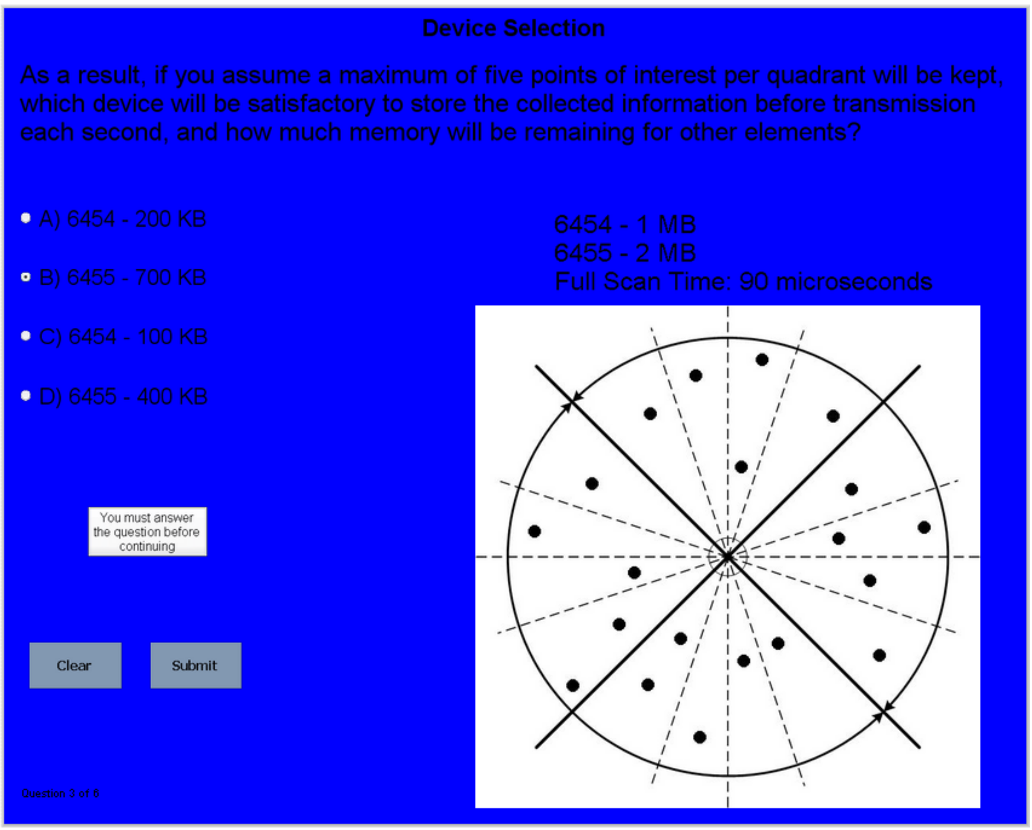

Figure 3-8. DSP processor selection.

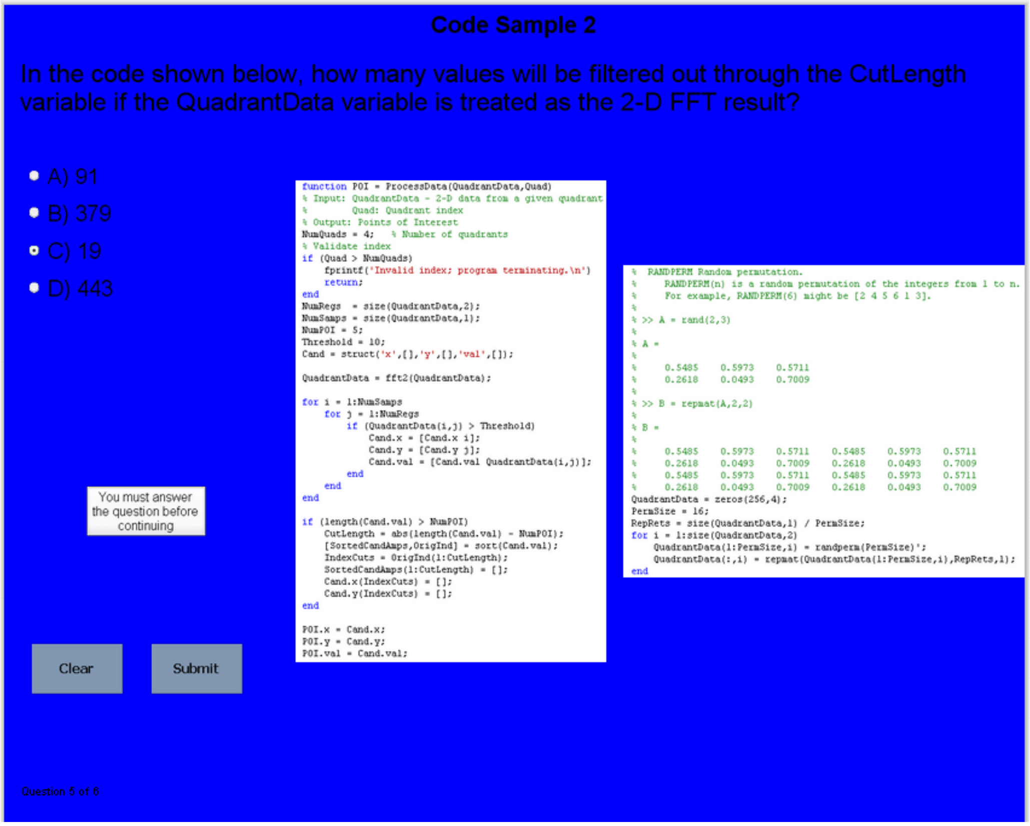

Figure 3-9. Implementation with high strength, periodic data.

CHAPTER 4
ASSESSMENT

Existing Assessment Methods

Traditional psychometric analysis is conditioned upon a variety of factors, all of which ostensibly serve to enhance objective sensibilities. For example, to assess the efficacy of a particular technique, students are frequently partitioned into two groups: one having access to the material in question and one using only traditional pedagogical techniques, the latter being described as a control group. In so doing, no assumption is made on the worth or harm of the technique, and the effort serves to stratify populations while still being unified under identical methods of testing knowledge. Certain assumptions concerning the distribution and stratification of these groups can indeed influence the types of conclusions that can be drawn, and natural statistical traits further impact this reality, particularly in the representation of given populations. Indeed, adequately large populations are often necessary to satisfy well-known statistical convergence requirements, implying limitations in the sample space either require sophisticated techniques to interpolate or extrapolate behavior, or otherwise inconclusive results follow. Functionally similar procedures are maintained in medical experiments with the use of placebos for demographically diverse populations, and these methods are so pervasive as to be a commonly accepted practice. Nonetheless, though additional tests can be made concerning correlation, causation, and belief in dichotomous hypotheses, the existing framework and curricular offerings have shown to limit available approaches.

Employed Testing Methodologies

To that end, though the modules themselves were developed in the sequence previously described, the structure and content of existing undergraduate courses in signal processing influenced the order and manner in which the modules were distributed. Moreover, contrary to the author's intended method of use, students were ultimately not required by their instructor to participate in these exercises and were informed during their lectures and associated announcements that involvement would be entirely optional. It therefore should be stated unequivocally that the results stemming from this effort, particularly attempts to measure the effectiveness of the modules themselves, are not necessarily reflections of the product itself, but rather of the mechanisms used to deliver it. To that end, though unintended, this investigation has revealed that the approach students take with respect not only to the notion of optional material but also to their participation in electronic pedagogy can be reliably and quantitatively shown to largely be an act of expedited guesswork. Whether this therefore serves as a reflection of the arguments made in preceding chapters that unperceived stagnation and fettered creativity are at the root of the problem has not undergone the necessary level of scientific scrutiny, but adequate evidence has been collected to suggest and highlight the belief that regurgitation impacts a student's capacity for abstraction, and reinforces the commonly held belief that students largely oblige themselves only to that which is mandated to them. These effects were first made manifest with freshman and sophomore students and its effects were exacerbated as students persisted in the curriculum.

Testing with Introductory Undergraduates

The Fall 2008 semester was used for the purpose of initial screening of the modules themselves as well as their corresponding assessment techniques. A section of EEL3135 (Signal and Systems) under the administration of Dr. Fred Taylor was provided as a sample population. Within such a course, the typical student has had no previous engineering coursework and is only assumed to have gained the requisite mathematical repertoire to analyze the critical aspects of linear systems. (Several enrolled students were found to be graduate students, presumably auditing the course; their performance was omitted in subsequent analysis.) Because of this, only the FIR module was utilized for the purpose of assessment.

After students had been introduced to the critical aspects of FFTs and FIR filters, a testing technique was devised that would be of sufficient generality to be solved by a student using only material obtained through traditional pedagogical means, yet of sufficient specificity as to benefit from the inclusion of the learning module. Through the comparison of performance on the testing material both prior to and after the distribution of the module, the goal was to assess the difficulty of the problem itself as well as the learning module for students with the experience cited.

Testing question and solution

To that end, on 13 October, a single pre-test question was provided to students to be completed via the UF Learning Support Services assessment software prior to the next lecture of 15 October 2008. The question was as follows:

> Consider a DSP which receives data from 5 unique analog inputs simultaneously. Each input has an Analog-to-Digital Converter (ADC) with a common user-selected sampling rate. If each input stream consists of 2048 (buffered) samples that must all be filtered through a single 51st order symmetric, linear phase FIR within 10 milliseconds at the selected rate, what is the lowest nonzero frequency that can be represented?

Students were provided five options for a response, four of which assessed understanding of filter mechanics and a fifth which allowed the student to select none of the answers provided. A maximum of one hour was afforded to solve the problem.

With such a question, notional concepts of constant group delay and resolution are emphasized within the framework of a real-time requirement, much like the module itself. To answer the question, the student must first observe that since the lowest nonzero frequency is equivalent to the frequency resolution, the lowest sampling rate that will achieve the timeline will serve as the basis for the unaliased bandwidth. Division by the number of samples per input yields the desired resolution.

The options provided for answers intend to address typical conceptual errors when attempting to solve problems of this type. For example, the inclusion of one option would be correct in the case of parallel filter banks instead of the serial one so described. Similarly, the inclusion of another option results from failing to take into account the architecture of the filter, which reduces the group delay by a factor of two. The remaining response options were modifications or fusions of these themes to force the student to discriminate between nearly equal choices.

In similar form, a post-test was administered after the lecture of 15 October to be completed before the lecture of 17 October, with changes to the pre-test question in the number of buffered samples as well as the timeline restriction. Options were scaled accordingly and students were provided an identical period of time.

Preliminary testing results

A total of 42 students completed the pre-test, with 15 answering correctly, an approximate 35.7% success rate. By contrast, 39 students completed the post-test with 12 answering correctly, an approximate 30.7% achievement. Though the -5% differential in and of itself proves concerning – one could argue its statistically insignificance relative to random guessing – a deeper challenge arose. Specifically, a more thorough analysis of the pre-test and post-test results revealed an alarming oversight: many students (beyond the differential of three) did not participate in both examinations, requiring their performance to be discarded.

Of those students who participated in both tests, the results were mixed. Since the LSS system monitors the amount of time a given assessment takes prior to submission, this time differential was used as a filtering agent. Specifically, a student who spent an insignificant amount of time (15 seconds or less) in either instance was taken to have guessed; the guessing contingent was one of three well-defined ones. The remaining two groups, all found to have spent a plausible amount of time solving the problem, either successfully answered the problem correctly on both occasions or did not. Of those who did not, however, there was no clear correlation found between time investment and accuracy of result with the exception of one student, who did not answer the pre-testing question correctly but did so on the post-test while spending considerably less time doing the latter.

Given the inherent inconclusive nature of these results, an alternative viewpoint was instituted to aid in later testing. Specifically, since the principal goal was to determine the difficulty of the question and the use of the modules, all other assessments for the semester were examined. As shown in Figure 4-1, the average student spent approximately 10.5 minutes on the pre-test compared to approximately 6 minutes on the post-test, but with marginally different results. At the same time, most instances of high performance were clustered toward short solution times. The interpretation of this observation is partially obvious and partially insightful. To the former end, it stands to reason that the student who knows how to solve a problem is more likely to do so more quickly than the problem he or she does not. At the same time, and to address the latter aspect, questions whose complexity rests principally in calculation and not conceptual formulation can inherently require more time to solve, of which several of the assessments shown are a part, as are problems in similar introductory courses such as solid state electronics. As a result, though the question posed for students of this introductory course was posed once again to students of a senior undergraduate section for the purpose of comparison, subsequent questions were framed around the intent to have solutions readily accessible within the time frame described.

Testing with Senior Undergraduates

The Spring 2009 semester was utilized for the remainder of testing within EEL4750C (Introduction to Digital Signal Processing), a technical elective taken by undergraduate students after having completed introductory courses in both continuous and discrete signal processing. The timeline for the distribution of the learning modules was influenced by the nature of the coverage within the course so as to ensure that all topics contained within a particular module were covered in advance through established pedagogical techniques. To that end, the third learning module covering DSP choices was distributed first, with the remaining modules to follow. A total of 43 students were enrolled for the entirety of the exercises to be described.

Digital signal processing module test questions and solution

Immediately following a series of lectures on the FFT and its variants, a pre-test was distributed as a supplemental document during the lecture of 18 February 2009 to be completed as a paper submission preceding the lecture of 20 February. Two questions were put forth as follows:

> (1) Suppose you wish to process two-dimensional data: one part spatial, one part temporal. Much to your surprise, the DSP available to you for processing supports 1-D FFTs of radix size R={2,3,4,5,8,16}. As a result, if an N-point 1-D FFT of radix R takes N log_R N clock cycles to compute, what is the fastest that a 2-D FFT can be computed if you must collect *at least* 3100 temporal samples and 3800 spatial samples?
>
> (2) You're trying to save power in completing calculations, and are led to believe the power consumption will increase linearly with an increase to clock speed. The DSP requires 10 nW of power to calculate the 2-D FFT of (1) at its nominal 2.2 GHz clock rate. If such a clock can be divided down by any integer of your choosing, what is the fastest the 2-D FFT can be computed while requiring no more than 6 nW of power?

The first question, beyond reinforcing the notional concepts of two-dimensional transforms, attempts to elicit knowledge about the use of zero padding to improve computational complexity while still recognizing the worth of atypical radixes. The inclusion of a radix table proved necessary to ensure that students were not left with solving a nonlinear optimization problem as a function of constrained sample length. In particular, though a radix-5 FFT is supported and would result in the least amount of zero padding of the temporal samples (e.g. 3125), the logarithmic complexity in fact decreases if the zero padding is extended to 4096 samples and a radix-16 FFT is used. Using the same reasoning, the 3800 spatial samples are similarly zero padded to 4096 and a radix-16 FFT is again used. Recognizing additionally that the 2-D FFT is composed of successive invocations of 1-D FFTs, the complexity thus scales at double the square of the zero-padded length.

As with the initial distribution of testing, five options were given for answers. The four errant choices were constructed to siphon potential conceptual flaws in the temporal and spatial lengths and radixes, respectively. For example, one choice made use of radix-5 and radix-2 FFTs for lengths of the nearest power of the radix (e.g. 3215 and 4096). Another used the most generic form (radix-2) for both dimensions, while another option used radix-5 and radix-16. A final option was generated as a bogus compilation of numbers. In providing these choices, the degree of conceptual error, outside of guesswork, is more readily manifest.

To address the second question, the observations are arguably less profound. Due to a linear scaling of power consumption with frequency, the 6 nW requirement relative to a 10 nW baseline implies that no greater than 60% of the clock frequency can be used. Since this value equates to 1.32 GHz, the clock operating at 2.2 GHz must be divided by two in order to produce a 1.1 GHz clock. The resultant speed is merely the number of cycles calculated in the preceding question divided by this rate. To stratify responses once again, five options were provided. The first errant choice used a radix-5 and radix-16 FFT with a 1.1 GHz clock, while another mistakenly used the same radixes but for an undefined 1.32 GHz clock. A third option used the correct radixes but with an improper clock frequency, and a final flawed choice made correct use of the clock frequency but using the aforementioned bogus calculation.

Digital signal processing module testing results

Students were presented with the two questions on 20 February 2009 to be due by 23 February in pre-test form, after which the learning modules were made available. From there, minor modifications were made to the relevant quantities as post-test questions, available on 2 March and due by 4 March. Nine students abstained from the process entirely.

Using personal solution rates as a litmus test for probable guessing (given *a priori* knowledge of the correct technique and answer), the tests were examined both for accuracy as well as the speed with which such selections were made. Though it is fair to argue that students could have had exposure to the questions prior to their participation, such an argument implicitly convicts the student of some measure of academic dishonesty, which the experiment would have been ill equipped to tackle. As such, those responses produced at a faster rate were taken to suggest guesswork, and were lumped into categories of correct guessing and incorrect guessing. To account for the presumed effect of a faster solution rate as a function of prior exposure, a 20% attenuation was injected into the guessing measure, making the solution speed that much faster; though this value is effectively heuristic, it proved to adequately delineate between cases in which students were probable to have randomly selected an option and those who simply could not solve the problem presented.

The first question of the pre-test yielded an additional sixteen students not participating. Of those who did, eight produced the correct answer for a mean response of 26.67%, but only two did so within a plausible time frame. To that end, the average time investment across all participating students was approximately 7.76 minutes, with seventeen students investing under the threshold of a minute of time. Similarly, for the second pre-test question, nineteen students did not participate, and of those remaining who did, ten students answered correctly, but only one did so without being probabilistically conditioned on guessing, be it due to time investment or errors in the previous question.

For the post-test, a reduction to the seven students answered the first question correctly, one previously having done so in the pre-test. Of those demonstrating improvement, four of the six had invested a plausible amount of time in their responses, with the remainder guessing. Students not falling into any of these descriptions produced inconclusive results as a function of their time investment or did not complete both activities. To that latter end, four of the six students who correctly answered the first pre-test question correctly in any form ultimately answered the post-test incorrectly, and five of the six did the same for the second question. By contrast, four of fourteen went from incorrect to correct for the first question while four of eleven did so for the second question.

Filter design modules questions and answers

Distributed on 25 March 2009 to be due within three days, students were presented with several questions covering the topics of FIR and IIR filter design and an external reminder was produced concerning the optional nature of the events. The three questions addressed the concept of both tone and broadband jamming without making direct use of such terminology. To that end, the pre-test questions were as follows:

> (1) Consider a DSP that receives data from five unique analog inputs simultaneously. Each has an Analog-to-Digital Converter (ADC) with a common user-selected sampling rate. If each input stream consists of 2048 (buffered) samples, all of which must be filtered through a single 51st order symmetric, linear phase FIR while experiencing no more than ten millisecond latency at the selected rate, what is the lowest nonzero frequency that can be represented?
>
> (2) Suppose now a sinusoid operating out of the pass band of the filter interfered with one of the five analog inputs and it was known to have no more than 10 kilowatts of peak received power. For a maximally flat pass band, what minimum stop band attenuation

would be required of the filter in (1) to ensure a 1 volt signal in the pass band was at least 20 dB above the out-of-band tone?

(3) Now, instead of a single sinusoid interfering with one of the inputs, suppose there is a collection of sinusoids at each input having equal strength, with each sinusoid transmitted at a constant frequency spacing from one another across the whole unaliased band width of the system. If the filter order is kept fixed at 51 but can be redesigned to produce a "best effort" of attenuating each of these tones, what is the minimum amount of memory required to store all five output data streams and the filter coefficients if all elements are stored as 16-bit values?

The first question was identical to that provided to students in the previous semester in an effort to potentially glean measured improvement as a function of experience within the engineering curriculum. The second question addressed the fundamental concepts of frequency response and units of measurement, wherein the effects of filter attenuation behave additively when measured in decibels. Students were given options that either failed to adequately attenuate the disruptive tone or did not do so minimally relative to other choices. The third question was posed as the greatest challenge to students, requiring an ability to realize that the abstract description provided is representative of a comb filter. As such, memory savings can be had by intelligent programming, which is reflected in the learning modules themselves. With that, students were provided a combination of choices that included and excluded both the number of inputs, the contribution of filter coefficients to the memory budget, and the ability to interpret the results in the listed units of bytes.

Filter design modules testing results

Using personal solution rates and the Fall 2008 results as a metric for predicting guesswork, a total of 6.5 minutes were allocated to the pre-test questions and a similar solution rate increase of 20% was afforded to the post-test. Moreover, the three questions provided signified three distinct levels of difficulty, implying that a student was most probable to have solved the second question most easily, with the first question being of moderate difficulty, and the final question being the most challenging.

Participation in these activities (as well as other non-experimental procedures) dwindled considerably by the time of distribution. For the pre-test, for example, 29 of the 43 students did not participate. For the first question, two students answered correctly, with one student answering within an improbable time frame. The second question found four students answering correctly, with one doing so at a conceivably high rate. Four students also answered the third question correctly, but one did so without any accuracy in preceding questions (while still investing plausible amounts of time) and two others did so simply by a likely speedy guess.

In the post-test, 23 students did not participate. Though six students answered the first question correctly, four did not do so within a plausible time frame. Six students also answered the second question correctly, with two likely to have guessed. A mere two students produced correct solutions to the final question, however, and neither did so within a credible time line. Both students answering the first question correctly initially failed to do so again; one in four also did for the second question and two of three for the third. Conversely, five of ten students who answered the first question incorrectly were correct in the post-test, along with two of eight for the second question and one of nine for the third question.

Conclusions

As already suggested, the assessment of these modules comes with significant caveats. However, in recognition of limited participation and inconsistency in results, it is worthwhile to note that these habits are in fact consistent with other pedagogical examinations that were periodically given throughout the semester in the same medium and with the same restrictions. In this regard, differentiation of the worth to the learning modules relative to existing content is not possible by sheer virtue of such limitations, but does speak to a larger issue: students demonstrate either limited genuine dynamic problem solving capability or interest. In fact, the student body was shown to only consistently produce accurate solutions to problems when such problems were sufficiently similar to existing available material, irrespective of the pervasive nature of concepts such as frequency response. In this regard, the argument of rapid production of answers as a result of pre-existing understanding is reinforced. The converse of this hypothesis, however, is challenged by the results shown in this effort: expedient responses alone neither signify nor imply knowledge. Though this fundamental observation is in no way surprising, the frequency with which it occurs suggests that policies mandating participation reduce the haphazard attitude with which circumstances are approached. To that end, modifications to the manner in which this assessment was conducted could conceivably have significant impact on the conclusions drawn from the effort itself.

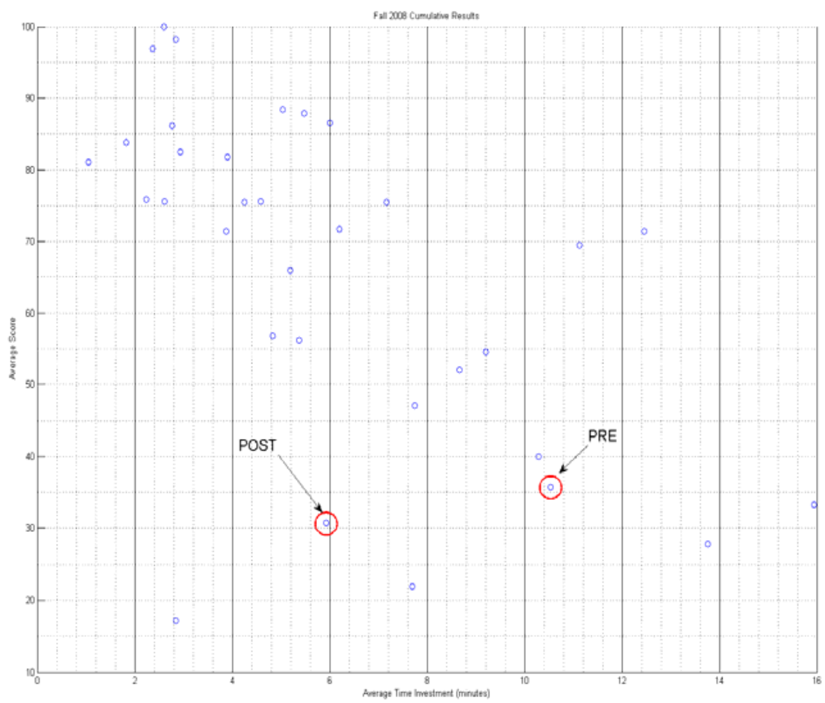

Figure 4-1. Fall 2008 cumulative testing results.

REFERENCES

[1] Selingo, J., “Game of Chance,” *ASEE Prism,* Vol. 17, No. 5, 2008.

[2] Buderi, R., *The Invention That Changed the World: The Story of Radar from War to Peace*, London: Abacus, Little, Brown, & Co., 1999.

[3] Shannon, C., and W. Weaver, *The Mathematical Theory of Communication,* Urbana and Chicago: University of Illinois Press, 1963.

[4] Gupta, M., “Is Industrial Experience Necessary for Teaching?” *IEEE Transactions on Education,* Vol. 31, No. 1, 1988, pp. 9-20.

[5] Wulf, W.A., “The Urgency of Engineering Education Reform,” *Journal of SMET Education,* Vol. 3, No. 3-4, 2002, pp. 3-9.

[6] Fortenberry, N.L., J.F. Sullivan, P.N. Jordan, and D.W. Knight, “Retention: Engineering Education Research Aids Instruction,” *Science,* Vol. 317, No. 5842, 2007, pp. 1175-1176.

[7] Loftus, M., “Why Won’t She Listen?” *ASEE Prism,* Vol. 17, No. 3, 2007.

[8] Committee on Science, Engineering, and Public Policy, *Rising Above the Gathering Storm: Energizing and Employing America for a Brighter Economic Future,* Washington, D.C.: National Academies Press, 2007.

[9] Gereffi, G., V. Wadhwa, B. Rissing, and R. Ong, “Getting the Numbers Right: International Engineering Education in United States, China, and India,” *Journal of Engineering Education,* Vol. 97, No. 1, 2008, pp. 13-25.

[10] Lemonick, M., “Are We Losing Our Edge?” *Time,* February 5th, 2006.

[11] National Academy of Engineering, *Educating the Engineer of 2020: Adapting Engineering Education to the New Century,* Washington, DC: National Academies Press, 2005.

[12] Cantrell, P., G. Peckan, A. Itani, and N. Velasquez-Bryant, “The Effects of Engineering Modules on Student Learning in Middle School Science Classrooms,” *Journal of Engineering Education,* Vol. 95, No. 4, 2006, pp. 301-309.

[13] Qualters, D., T.C. Sheahan, E.J. Mason, D.S. Navick, and M. Dixon, “Improving Learning in First-Year Engineering Courses through Interdisciplinary Collaborative Assessment,” *Journal of Engineering Education,* Vol. 97, No. 1, 2008, pp. 37-45.

[14] Hairston, G., “Progressive Educators Must Lead, Not Follow,” *Journal of SMET Education,* Vol. 1, No. 3, 2000, pp. 1-4.

[15] Nature Editorial Board, “A Crisis of Confidence,” *Nature,* Vol. 457, p. 635.

[16] Journal of Engineering Education Editorial Board, “Special Report: Research Agenda for the New Discipline of Engineering Education,” *Journal of Engineering Education,* Vol. 95, No. 4, 2006, pp. 259-261.

[17] Streveler, R., and K. Smith, “Conducting Rigorous Research in Engineering Education,” *Journal of Engineering Education,* Vol. 95, No. 2, 2006, pp. 103-105.

[18] Borrego, M., “Development of Engineering Education as a Rigorous Discipline: A Study of the Publication Patterns of Four Coalitions,” *Journal of Engineering Education*, Vol. 96, No. 1, 2007, pp. 5-18.

[19] Clark, M.C., J. Froyd, P. Merton, and J. Richardson, “The Evolution of Curricular Change Models Within the Foundation Coalition,” *Journal of Engineering Education,* Vol. 93, No. 1, 2004, pp. 37-47

[20] Bransford, J., A. Brown, and R. Cocking (Eds.), *How People Learn: Brain, Mind, Experience, and School*, Washington, D.C.: National Academies Press, 1999.

[21] Wankat, P.C., “Improving Engineering and Technology Education by Applying What is Known About How People Learn,” *Journal of SMET Education,* Vol. 3, No. 1-2, 2002, pp. 3-8.

[22] Canizares, C., and Z.T. Faur, “Advantages and Disadvantages of Using Various Computer Tools in Electrical Engineering Courses,” *IEEE Transactions on Education,* Vol. 40, No. 3, 1997, pp. 166-171.

[23] Walker, M.B., and J.F. Donaldson, “Continuing Engineering Education by Electronic Blackboard and Videotape: A Comparison of On-Campus and Off-Campus Student Performance,” *IEEE Transactions on Education,* Vol. 32, No. 4, 1989, pp. 443-447.

[24] Joy, M., and M. Luck, “Plagiarism in Programming Assignments,” *IEEE Transactions on Education*, Vol. 42, No. 2, 1999, pp. 129-133.

[25] Parker, A., and J.O. Hamblen, “Computer Algorithms for Plagiarism Detection,” *IEEE Transactions on Education,* Vol. 32, No. 2, 1989, pp. 94-99.

[26] Carpenter, D., T.S. Harding, C.J. Finelli, S.M. Montgomery, and H.J. Passow, “Engineering Students’ Perceptions of and Attitudes Towards Cheating,” *Journal of Engineering Education*, Vol. 95, No. 3, 2006, pp. 181-194.

[27] Hoole, S.R.H., “Engineering Education, Design, and Senior Projects,” *IEEE Transactions on Education*, Vol. 34, No. 2, 1991, pp. 193-198.

[28] Borrego, M., “Conceptual Difficulties Experienced by Trained Engineers Learning Educational Research Methods,” *Journal of Engineering Education,* Vol. 96, No. 2, 2007, pp. 91-102.

[29] Dee, K.C., “Student Perceptions of High Workloads are Not Associated with Poor Student Evaluations of Instructor Performance,” *Journal of Engineering Education*, Vol. 96, No. 1, 2007, pp. 69-78.

[30] Rosati, P., R.K. Dean, and S.M. Rodman, “A Study of the Relationship Between Students’ Learning Styles and Instructors’ Lecture Styles,” *IEEE Transactions on Education,* Vol. 31, No. 3, 1988, pp. 208-212.

[31] Vogt, C., “Faculty as a Critical Juncture in Student Retention and Performance in Engineering Programs,” *Journal of Engineering Education*, Vol. 97, No. 1, 2008, pp. 27-36.

[32] Roselli, R., and S.P. Brophy, “Experiences with Formative Assessment in Engineering Classrooms,” *Journal of Engineering Education*, Vol. 95, No. 4., 2006, pp. 325-333.

[33] Jiutso, S, and D. DiBasio, “Experiential Learning Environments: Do They Prepare Our Students to be Self-Directed, Life-Long Learners?” *Journal of Engineering Education*, Vol. 95, No. 3, 2006, pp. 195-204.

[34] Kumsaikaew, P., J. Jackman, and V.J. Dark, “Task Relevant Information in Engineering Problem Solving,” *Journal of Engineering Education*, Vol. 95, No. 3, 2006, pp. 227-239.

[35] Walls, M, “The ‘Undisciplined’ Interdisciplinary Problem: PBL and the Expanding Limits of SMET Education,” *Journal of SMET Education*, Vol. 1, No. 3, 2000, pp. 21-24.

[36] McClellan, J., C.S. Burrus, A.V. Oppenheim, T.W. Parks, R.W. Schafer, and H.W. Shuessler, *Computer-Based Exercises for Signal Processing,* Upper Saddle River, New Jersey: Prentice Hall, 1998.

[37] Chassaing, R., *Digital Signal Processing and Applications with the C6713 and C6416 DSK*, Hoboken, New Jersey: Wiley & Sons, 2005.

[38] Chassaing, R., *DSP Applications Using C and the TMS320C6x DSK*, New York: Wiley & Sons, 2002.

[39] Kehtarnavaz, N., and M. Keramat, *DSP System Design Using the TMS320C6000,* Upper Saddle River, New Jersey: Prentice Hall, 2001.

[40] Gan, W., and S.M. Kuo, “Teaching DSP Software Development: From Design to Fixed-Point Implementations,” *IEEE Transactions on Education,* Vol. 49, No. 1, 2006, pp. 122-131.

[41] Bose, T., “A Digital Signal Processing Laboratory for Undergraduates,” *IEEE Transactions on Education*, Vol. 37, No. 3, 1994, pp. 243-246.

[42] Christensen, M., R. Troche, and F. Taylor, "ENGAGE: A DSP Learning Paradigm for an Engineering Age," *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2008, pp. 2629-2632.

[43] Schwartz, M., "The Importance of Stupidity in Scientific Research," *Journal of Cell Science*, Vol. 121, 2008, p. 1771.

[44] Robinson, K., "Do Schools Kill Creativity?" 2006 TED Conference, http://www.ted.com/index.php/talks/ken_robinson_says_schools_kill_creativity.html, Accessed February 2009.

[45] Taylor, F., and J.D. Mellott, "SPECtra: A Signal Processing Engineering Curriculum," *IEEE Transactions on Education*, Vol. 39, No. 2, 1996, pp. 180-185.

[46] Taylor, F., R. Frankel, and N. Nezis, "The InvestiGATOR: A Studio-Based DSP Learning Paradigm," *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2004, pp.1013-1016.

[47] Dahm, K., "Interactive Simulation for Teaching Engineering Economics," *Journal of STEM Education*, Vol. 4, No. 3-4, 2003, pp. 1-4.

[48] Roskowski, A.M., R.M. Felder, and L.G. Bullard, "Student Use (and Non-Use) of Instructional Software," *Journal of SMET Education*, Vol. 2, No. 3-4, 2001, pp. 41-45.

[49] Jawaharal, M, A. Shih, and P. Schrader, "Use of Scenario-Based Learning in Teaching Statics," *Proceedings, ASEE Annual Conference and Exposition,* American Society for Engineering Education, 2004.

[50] Christensen, M., and F. Taylor, "A Scenario-Based DSP Learning Study," *Proceedings, 2009 IEEE Signal Processing Society 13th DSP Workshop & 5th SPE Workshop,* Marco Island, FL, January 4-7.

[51] Williams, A., and F. Taylor, *Electronic Filter Design Handbook,* New York: McGraw-Hill, 2006.

.

BIOGRAPHICAL SKETCH

Michael Christensen was born in 1982 in El Paso, Texas. The younger of two children, he moved with his family to Florida in 1988 and has since lived in the cities of Lakeland and Gainesville while garnering his education. He received his Bachelor of Science (Magna cum Laude) and Master of Science in Electrical Engineering degrees in 2004 and 2005, respectively, from the University of Florida. From 2003 to 2005, he served as an undergraduate laboratory instructor in communication systems while dual enrolled in graduate school. From 2005 to 2006, he undertook research in radar signal processing sponsored by the Army Research Laboratory to develop and enhance machine learning algorithms for airborne minefield detection. He then allocated 2006 to 2007 exclusively to engineering education research, with numerous guest lectureships at the University of Florida focused on target environment digital signal processing. Since then, he has simultaneously worked as a research engineer for the University of Florida in the development of end-to-end radar systems for numerous defense agencies and a contributor to multiple undergraduate electrical engineering courses focused in signal processing. He has additionally contributed regularly both to signal processing and educational literature with publications in books, magazines, journals, and conferences.

Since completing his degree, Michael has been in negotiation with a host of defense agencies to offer his services in radar signal processing as an engineer.