A research-oriented project that motivates undergraduate students in digital signal processing

Stephen So
Griffith School of Engineering, Gold Coast campus, Griffith University, QLD, Australia.
Corresponding Author Email: s.so@griffith.edu.au

Structured abstract

BACKGROUND

Digital Signal Processing (or DSP) is an important area whose applications pervade several areas of modern electrical and electronic engineering, such as information and communication systems, digital control systems, power engineering, mechatronics, biomedical engineering, etc. Therefore it is a course that comprises two aspects: a strong mathematical and theoretical part as well as a practical aspect. More often than not, students get bogged down with the theory without seeing the applications, and this has been observed to adversely affect their motivation for studying the course. Therefore, the over-arching goal in digital signal processing education is to link these two aspects together in a coherent and comprehensive way, such that it not only assists in improving student understanding of DSP theory, but also allows them to fully appreciate its applicability and effectiveness in practice. This paper describes the final project (Biometric speaker verification) in 4307ENG Advanced Digital Signal Processing that is offered in the ECE program within the Griffith School of Engineering on the Gold Coast campus. This project consisted of a research-oriented 'extensions' process that encouraged and motivated students to further investigate competing or improved algorithms and then report on their findings.

PURPOSE

We wish to report on the design of a final project in the course Advanced Digital Signal Processing that served to improve the understanding of DSP theory as well as implementation skills of undergraduate students, through a research-oriented 'extensions' process, which motivated them to investigate further, encouraged them to take ownership of the project, and allowed them to better appreciate the discipline.

DESIGN/METHOD

Feedback and opinions were obtained from two cohorts of students who studied 4307ENG Advanced Digital Signal Processing on the Gold Coast campus.

RESULTS

The feedback obtained showed the positive benefits of the use of 'extensions' as the main driver for further investigations into the system. Friendly competition developed among the student groups to design and implement the best performing system. Students felt that the project had improved their understanding of the DSP concepts as well as aided in seeing their interactions in a complete DSP system.

CONCLUSIONS

The design of the final year project was successful in combining the DSP theory taught and their practical implementation, as well as allowing students to better appreciate the applications. Moreover, through the research-oriented 'extensions' process, it was effective at motivating students to experiment and investigate competing algorithms as well as report on their findings.

KEYWORDS

Digital signal processing, research-oriented projects

Introduction

Advanced Digital Signal Processing (4307ENG) was a final year elective course that was taught in the four-year Bachelor of Engineering in Electronic and Computer Engineering at Griffith University on the Gold Coast campus. The course built upon the third year core course, 'Digital Signal Processing' (DSP), by introducing final-year students to the concepts of random and stochastic digital signal processing, non-parametric and parametric spectral estimation, as well as Wiener and adaptive filters. Traditionally these topics have been taught in postgraduate courses, rather than undergraduate. However, the applications of statistical digital signal processing are numerous in the electronic engineering industry, such as computer technologies, mechatronics, communications, biomedical engineering, digital control systems, etc. Therefore, McClellan et al. (1998) described the existence of a "gap between what is taught at a typical university and what graduating engineers need to bring into the workforce". The 4307ENG Advanced Digital Signal Processing course aimed to fill this gap by exposing final-year undergraduates to this important area. Many of the industrial applications mentioned above require *real-time signal processing*. Therefore, there exists a strong need to train undergraduate engineering students in the application and development of DSP algorithms running on powerful special-purpose *digital signal processors* (DSPs) (Chassaing, et al. 1993; Zacharias and Conrad, 2007).

The objectives of the course content can divided into two aims: (i) *providing* the *theoretical foundations* in statistical digital signal processing that rely heavily on probability, statistics, optimisation, and differential calculus; and (ii) *developing sophisticated algorithms* for representing, analysing and processing stochastic signals in software and on embedded DSPs. There are many challenges in teaching a Digital Signal Processing-based course that satisfies these aims. According to Melton et al. (1999), acquiring scientific knowledge requires three pieces of information: (i) "a level of abstraction and modelling"; (ii) "an *application* that motivates concept comprehension"; and (iii) "understanding the links between the two". Primarily, there is often a gap in DSP education between understanding the mathematical foundations and seeing their relevance in a practical context (Albu and Malakuti, 2009). Therefore, it was important to emphasise the applications of digital signal processing (Zoltowski et al., 1996) in order to link and reinforce them with the theory taught (Chassaing et al., 1993).

Motivation is a key issue in DSP education. The abstract and conceptual nature of the course content (for example, with probability density functions and Gaussian random processes) often fails to motivate engineering students, and it has been noted before that students testify theoretical courses as "only endured, not enjoyed" (McClellan et al., 1998). Undergraduate students in electrical engineering find signal processing courses difficult and perceive it as abstract mathematics (Marozas & Dumbrava, 2010). The question therefore is how can one assist students to understand and reinforce the theoretical content with applications; gain the necessary practical skills to prepare them for industry; and motivate them in a research mode to investigate further? According to McClellan et al. (1998), DSP has some natural advantages "to motivate further study and questioning". Furthermore, projects should be offered that are based on "familiar realistic systems". The idea is that by allowing students to relate common real-life systems to the underlying fundamental DSP theory in a hands-on project setting, they will not only be able to link the two aspects (theory and application), but will also be motivated to further explore and experiment with new techniques. Therefore, it is possible to weave a research-oriented component into the hands-on project that guides students in the process of undertaking DSP research.

In 4307ENG, an emphasis on speech signal processing was chosen due to the expertise of the original course convenor, Professor Kuldip K. Paliwal. Speech processing encompasses the applications of speech compression, speech recognition, speech synthesis, and speech enhancement (Rabiner and Schafer, 2011). Speech and image processing have traditionally been areas that were covered only at postgraduate levels (Zoltowski et al., 1996). The processing of one-dimensional speech signals is somewhat easier than two-dimensional images; hence speech processing can be placed at a suitable level for undergraduate project work. Furthermore, the use of speech-based applications allows students to relate the theory to the real-life activities they perform every day, such as speaking and listening to others, either in person or over mobile telephony. Speech-enabling technologies are also very prevalent these days. For example, Apple Inc.'s Siri personal assistant is available on iOS-based mobile devices, where spoken user commands are interpreted by Siri (speech recognition) and the results of the query are spoken back (speech synthesis).

In this paper, a research-oriented final project for 4307ENG will be described. It is a project that seeks to reinforce the theory via applications for students, gives them practice in development and implementation, and fosters motivation for further investigation. Specifically, the use of a research-oriented 'extensions' process will be discussed and its effectiveness will be evaluated using quantitative and qualitative feedback from students over two semesters.

Curriculum and assessment structure

The curriculum and primary assessment structure for 4307ENG Advanced Digital Signal Processing was originally designed by Professor Kuldip K. Paliwal from the Griffith School of Engineering at the Nathan campus. It is presented here to show how the course content linked in with the laboratory experiments and ultimately the final project.

Table 1 presents the curriculum for the course, which is taught over a 13 week semester. The topics are covered in the lectures using the chalk-and-talk style (So, 2012) and involve a fair degree of mathematical rigour. These topics are complemented by hands-on laboratory experiments, which were designed to provide students with practice in algorithm development in both MATLAB and the C language running on an embedded DSP. The embedded DSP used in this course is the Texas Instruments TMS320C6713 Digital Starter Kit (DSK). It is important to provide hands-on DSP experiments as they have been reported to have a "major positive impact on their understanding of basic DSP concepts" (Chassaing et al., 1993).

Various DSP applications and the latest research were discussed in each topic to reinforce the concepts that were covered. For example, after students were introduced to the concept of Gaussian probability density functions (PDFs) in Topic 1, they were shown how Gaussian PDFs, in the form of Gaussian mixture models (GMMs), can be used to improve the efficiency in wideband speech coding systems (So and Paliwal, 2005). This provided the link between fundamental concepts that the students were learning and the current research.

Table 1: Curriculum for 4307ENG Advanced Digital Signal Processing

Topics	Laboratory Experiment	DSP applications highlighted
Topic 1: Random variables and processes	Lab 1: Random variables (MATLAB)	Gaussian mixture models for speech coding
Topic 2: Non-parametric spectral estimation	Lab 2: Non-parametric spectral estimation (*) (MATLAB)	Spectral subtraction in speech enhancement
Topic 3: Parametric spectral estimation	Lab 3: Autoregressive spectral estimation (*) (MATLAB)	Speech coding, speech synthesis, and speaker verification
Topic 4: Wiener filtering	Lab 4: Noise removal (*) (MATLAB and DSK)	Wiener filters in speech enhancement
Topic 5: Adaptive filtering	Lab 5: Adaptive differential pulse code modulation (MATLAB and DSK)	Noise cancellation

The laboratory experiments labelled with a (*) in Table 1 led up to the final project. In other words, the MATLAB and C-based algorithms developed by students in these laboratory experiments formed critical sections of the system in the final project. These algorithms include:

- Implementing overlapping frames (Lab 2)
- Implementing the Hamming window (Lab 2)
- Implementing the Levinson-Durbin algorithm (Lab 3)
- Real-time signal processing on the DSK (Lab 4)

Table 2 shows the assessment structure in 2010 (the structure in 2012 had minor weighting adjustments to the final examination). Students undertook prescribed laboratory experiments in Weeks 2 to 9 and then spent the remaining time working on the final project. More details on the final project will be given in the next section.

Table 2: Assessment structure of 4307ENG Advanced Digital Signal Processing

Assessment Item	Weighting (%)
Laboratory reports × 5 (*Weeks 2 – 9*)	25
Final project (*Weeks 10 – 13*) • Technical defence (3%) • Project report (5%) • Basic MATLAB implementation (3%) • Basic DSK implementation (6%) • Extensions (8%)	25
Final examination (3 hrs)	50

A research-oriented final project

Description of the *basic system*

The original task assigned to the final project was to develop a biometric speaker verification system in both MATLAB and the C language running in real-time on the C6713 DSK. A speaker verification system validates the identity of the user by using their speech, similar to the use of a password to verify a username. A block diagram of this system is shown in Figure 1. The *basic system* (i.e. with no extensions) was briefly described to the students in a generic way. Working in pairs, the students were then expected to work out the specifics of the system implementation by using their knowledge from the lectures and previously written code from the laboratory experiments to produce a functioning system.

In the MATLAB implementation, eight utterances of the word ‘money’ from five different speakers were provided as wave files to students. The student then chose a subset of these files as training utterances, which were then used to train the system and produce reference patterns for each speaker. The remaining subset of utterances was then used for the verification phase.

For the embedded DSP implementation, students were expected to develop a system that could be trained using the voice of one member of the group in real-time and then verified the identity of an unknown user. The result of the verification process was indicated by the system via the printing of the words “ACCEPT” or “REJECT”, activating an on-board LED, or generating an appropriate tone. The recognition accuracy of the system can be defined as the ratio of the number of correct verifications with the total number of trials. The basic system had a mediocre recognition accuracy of around 60% and hence there was plenty of room for improvements.

Limited assistance was provided to students in the programming of the system, since it was expected that the prescribed laboratory experiments prior to the project had provided them with sufficient experience in their debugging and implementation skills.

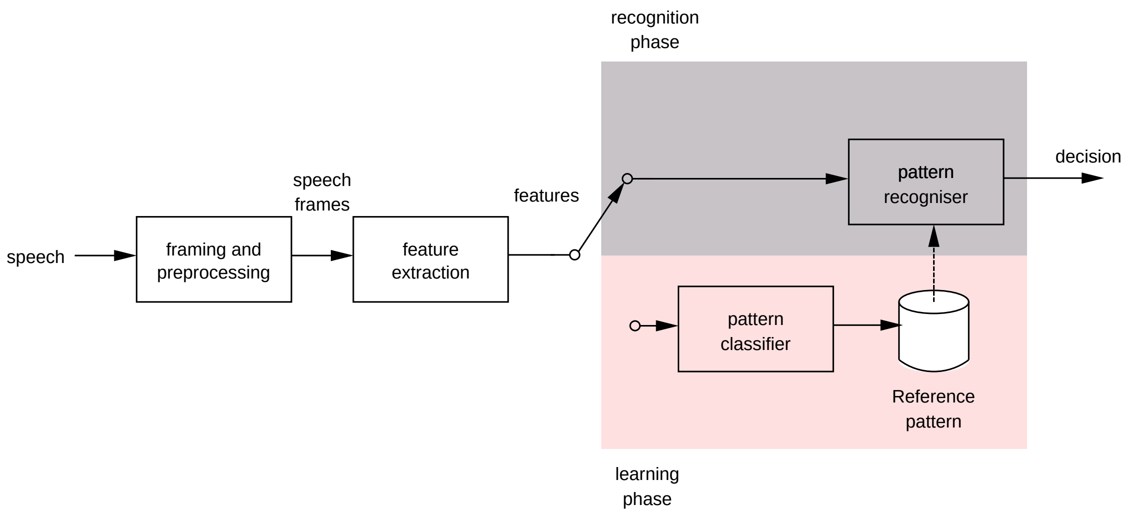

Figure 1: Block diagram of the biometric speaker verification system to be developed in the final project

Project 'extensions': motivating students for further study

According to the marking criteria, students who were able to correctly implement the *basic system* described in the previous section, were awarded up to a maximum mark that is less than the total of 25 marks. In order to achieve the total of 25 marks, students needed to modify their basic system to include 'extensions'. A suggested list of 'extensions', as shown in Table 3, and the approximate marks that they attract, was provided to the students. These 'extensions' are well-established methods that were investigated and reported in the speech processing research literature. They were purposely not covered in the course content. It then became the task for students to undertake research into each of these techniques by reviewing the literature. The teaching staff members were available to provide general guidance on these 'extensions' but ultimately, it was the students' responsibility to understand, implement, and experiment with each 'extension'. The marks for each 'extension' were chosen based on their relative level of difficulty and complexity.

Table 3: List of 'extensions' and corresponding marks

Extensions	Marks
Effective voice activity detection or VAD	1%
Pre-emphasis filter and Hamming window	1%
Overlapped frames	1%
Linear prediction cepstral coefficients	2%
Cepstral liftering	1%
Adaptive threshold for EER	2%
Delta and acceleration (delta-delta) features	1%
K-means classifier	3%
Gaussian mixture model-based classifier	3%
Mel frequency-warped cepstral coefficients	3%

Building upon a basic system that had room for improvement was crucial to the success of the project in motivating students. For example, a real-time DSP project from the California State University, Chico, involved students working independently on a self-tuning real-time adaptive filter, which needed to be optimised based on several performance metrics such as signal-to-noise ratio (SNR), adaptation speed, convergence stability, etc. The study by Chassaing et al., (1993) concluded that the DSP students had been "very highly motivated by the laboratory component, and no longer view[ed] DSP as abstract and esoteric".

In the case of the biometric speaker verification system, the performance metrics include recognition accuracy (RA), equal error rate (EER), false acceptance rates (FAR), false rejection rates (FRR), robustness to noise, etc. Each extension that was correctly implemented would mostly improve some or all of these measures incrementally and these were to be measured by the students. It was observed that a form of friendly competition started to develop in the class, where student groups were attempting to tune their system to achieve the best metrics, such as FAR or EER. Therefore, there was a notable sense of motivation in the class as well as a drive to optimise the system to the best they can. By doing their own research and implementation of the 'extensions', students had developed a sense of ownership of their implemented system.

Figure 2: Texas Instruments TMS320C6713 DSK used in the final project

Assessment methods

In the final week, each student group was required to briefly present and demonstrate their working system. This formed part of the technical defence, where the teaching staff asked each student questions about the technical specifics of their system in order to gauge their understanding of the project. Each student was expected to discuss about the 'extensions' they used, how they worked, and explain their effects (positive and negative) on the final recognition performance. The use of technical defence or *viva voce* also allowed students to practice effective oral communication skills that will be of benefit to them for their final year dissertation. One student commented that this course and others had "used technical defences very well".

Another assessment item was the project report, which was written in the style of a four-page, two-column IEEE transactions paper. This item added another aspect of research to the project, namely the production of a (mock) journal publication. The quality of the literature review as well as the clear presentation and discussion of results were assessed in the project report.

Discussion of assessment design

Biggs and Collis (1982) introduced the SOLO (Structure of the Observed Learning Outcome) taxonomy that describes and classifies the quality of student learning outcomes into five levels: prestructural, unistructural, multistructural, relational, and extended abstract. The prescribed laboratory experiments were designed to provide *multistructural* understanding, where different and unrelated signal processing techniques and algorithms were introduced to students. At this stage, students only understand each technique independently but do not see the connections. The final project provides the necessary *relational* understanding, where the loosely connected concepts covered in the laboratory experiments are integrated into a fully working system. The 'extensions' process of the project prepares students with the research and critical thinking skills, which puts them on the path to generalising the techniques to other fields and applications (*extended abstract* understanding).

Applicability to other engineering disciplines

It is worth mentioning that the research-oriented project outlined here is applicable to advanced elective courses in other disciplines of engineering. The main idea is to design a project that encompasses and integrates the same techniques covered in the course content and has a basic implementation that permits further room for improvement. By providing students with the necessary resources and literature, a suitable (and achievable) 'extensions' list and weighted marking criteria can be used effectively in a small-to-medium-sized class.

Evidence of success

Interdependence between final project and final examination marks

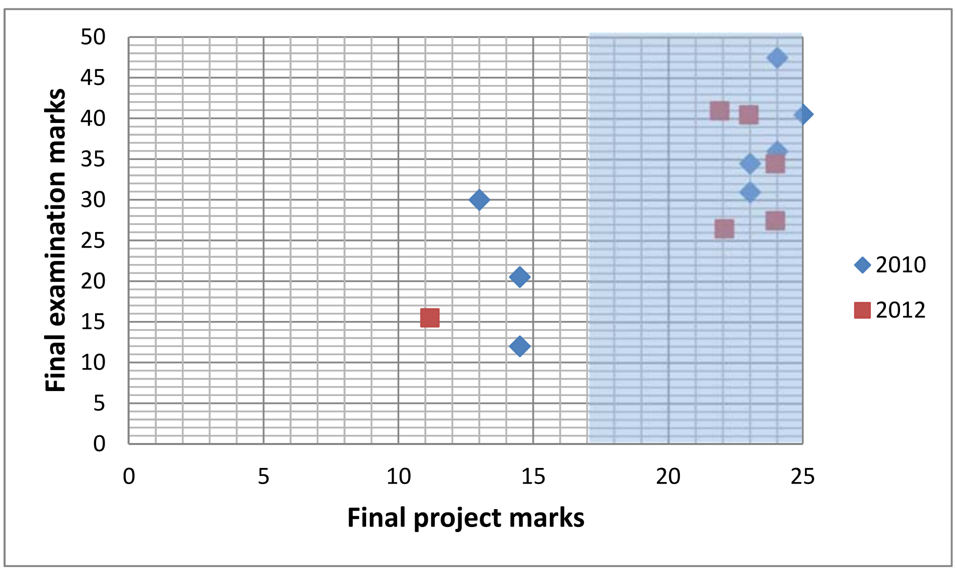

Figure 3: Graph of student project marks versus final examination marks in 4307ENG in Semester 2, 2010 and 2012. Shaded region shows final project marks obtained through implementing the 'extensions'

Figure 3 shows a graph of the student marks in the final project plotted against their final examination marks for the semesters in 2010 and 2012. Note that the course was not offered in 2011. Since 4307ENG was a specialised elective course in the ECE program, student numbers tended to be rather low. It can be seen that there is a general trend where students who scored well on the final project also tended to do well in the final examination. There were a few students who did not appear to fit this trend. For example, the 2010 student who achieved 13 in their final project managed to score 30 out of 50 for the final examination, which as can be seen was a mark that was slightly higher to those in 2012 (who got less than 30 in the final examination) but did fairly well in the project. It can be inferred that the former student was stronger on the theoretical side than the practical side, and vice versa for the latter 2012 students.

The blue shaded region in Figure 3 represented the marks obtained through the implementation of ‘extensions’. It can be seen that 10 out of the 14 students had successfully implemented the project ‘extensions’ and from this group, seven of them achieved higher marks on the final examination than those who had not.

Student feedback (quantitative)

Quantitative feedback from the student questionnaire is presented in Table 4 and is a combination of responses from the students in 2010 and 2012. It can be seen that the majority of students were positive about the effectiveness of the final project in linking theory to applications.

Student feedback (qualitative)

Table 5 lists the typical student responses obtained from both the student experiences and questionnaire. There were more responses from the questionnaire because it was more focused on their views of the final project. The student feedback here was very positive on the final project. Some students were pleased to have worked on a project with real world outcomes and relevance, which provided the necessary context. Another student appreciated how various concepts covered in the course were brought together into the project.

Table 4: Five-point Likert scale result from student feedback (Legend: SA – strongly agree; A – agree; N – neither agree nor disagree; D – disagree; SD – strongly disagree)

Question	SA	A	N	D	SD
Q1: *The final project of this course has increased my level of interest in digital signal processing.*	63.6%	36.4%	0%	0%	0%
Q2: *The final project of this course did not demonstrate digital signal processing concepts being put into practice.*	0%	0%	0%	27.3%	72.7%
Q3: *The final project of this course has linked together theoretical concepts covered in the course.*	63.6%	36.4%	0%	0%	0%
Q4: *The final project of this course did not provide me with the experience of implementing a complete digital signal processing application*	0%	0%	0%	27.3%	72.7%
Q5: *The use of both the MATLAB and the C6713 DSK development platforms in the final project of this course was beneficial to my engineering studies.*	54.6%	36.3%	0.1%	0%	0%
Q6: *The prescribed laboratory experiments earlier in the semester have not assisted me in undertaking the final project for this course.*	0%	0%	0%	54.6%	36.4%
Q7: *The final project in this course has improved my understanding of the relevant theoretical concepts covered in the lectures of this course.*	63.6%	36.4%	0%	0%	0%

Table 5: Student responses from student experiences and questionnaire

Responses from the student experience (What did you find particularly good about this course?)
“The labs and practical components were very applicable”

"The practical aspect was incredibly helpful. The project was very engaging and it seems like we've finally worked on a project with immediate real world outcomes. The entire subject was very interesting and kept me engaged the entire semester - it was very enjoyable."
Responses from the student questionnaire on the final project
"Excellent final project. Was challenging, diverse and provided me with a physical outcome. Positively influenced my professional development in signal processing" *"Advanced digital signal processing made me understand the dynamics in signal processing to a significantly greater level, beyond a point I didn't even think of. It was a challenging course to get my head around but largely beneficial. One thing could be to elaborate on the use of such processing techniques used in new technology and apps in phones these days. To me the durbin levinson algorithm was quite an interesting topic.. ;)"* *"Brilliant example of how to bring concepts together in a project that was both engaging and fun, and also had real world relevance which gave context to what we were doing."* *"It was a very fun and memorable project. Practical application of concepts, particularly advanced DSP concepts, is invaluable."* *"This was an incredibly exciting project, and helped integrate a number of digital signal processing concepts learnt throughout the course. Having first hand experience in designing systems like this is a very valuable asset."* *"The advice given by the lecturer inside and outside of course hours on the theoretical topics was invaluable to getting the most out of this project. The motivation of the lecturer made this project so successful."* *"Stephen So helped greatly in my understanding of this project and sparked for me an interest in DSP that I didn't have previously."*

Conclusion and further work

The design of the final project was successful in combining the DSP theory taught and its practical implementation, as well as allowing students to better appreciate the applications. Moreover, through a research-oriented 'extensions' process, it was found to be effective at motivating students to experiment and investigate various algorithms as well as report on their findings. The goal of the project was simple: develop a biometric speaker verification system that performed better than the basic system. While provided with general guidance on what appropriate 'extensions' were available, students performed their own research from the literature and took ownership of their project. A friendly competition developed among student groups in the class, where they were optimising their systems to deliver the highest performance, while at the same time, appreciating the benefits brought about by using alternative algorithms. The feedback from the students on the final project was positive and encouraging.

Given the small class size and nature of the course (final year, specialised elective), the influence of the 'Hawthorne effect' in the observations cannot be discounted. Therefore, further work in the form of applying the proposed final project structure to other cognate engineering courses and observing the trends is needed to reinforce the findings of this study.

References

Albu, A.B. & Malakuti, K. (2009). *Work in progress – problem-based learning in digital signal processing*. Paper presented at 39th ASEE/IEEE Frontiers in Education Conference, San Antonio, TX.

Biggs, J. & Collis, K. (1982). Evaluating the Quality of Learning: The SOLO Taxonomy. New York: Academic Press.

Chassaing, R., Anakwa, W. & Richardson, A. (1993). *Real-time digital signal processing in education*. Paper presented in IEEE International Conference on Acoustics, Speech, and Signal Processing, Minneapolis, Minnesota.

Marozas, V. & Dumbrava, V. (2010). Motivating the students to study the basics of digital signal processing by using virtual learning environment. *Electronics and Electrical Engineering.–Kaunas: Technologija*, 102(6), 87-90.

McClellan, J.H., Schafer, R.W. (1998). A changing role for DSP education. *IEEE Signal Processing Magazine*, 15(3), 16-18.

Melton, D.E., Finelli, C.J. & Rust, L.M. (1999). A digital signal processing laboratory with style. Paper presented at 29th ASEE/IEEE Frontiers in Education Conference, San Juan, Puerto Rico.

Rabiner, L.R. & Schafer, R.W. (2011). Theory and applications of digital signal processing. Upper Saddle River: Prentice-Hall.

So, S. (2012). *Refined 'Chalk-and-Talk' of Lecture Content: Teaching Signals and Systems at the Griffith School of Engineering*. Paper presented at the Australasian Association for Engineering Education Annual Conference, Melbourne, Vic.

So, S. & Paliwal, K.K. (2005). *Multi-frame GMM-based block quantisation of line spectral frequencies for wideband speech coding*. Paper presented at the IEEE International Conference on Acoustics, Speech, and Signal Processing, Philadelphia.

Zacharias, J.J. & Conrad, J.M. (2007). *A survey of digital signal processing education*. Paper presented at SoutheastCon, Richmond, Virginia.

Zoltowski, M.D., Allebach, J.P. & Bouman, C.A. (1996). : Digital signal processing with applications: A new and successful approach to undergraduate DSP education. *IEEE Transactions on Education*, 39(2), 120-126.

Acknowledgements

The author wishes to acknowledge Professor Kuldip K. Paliwal from the Griffith School of Engineering for the original development of the course and assessment structure of 4307ENG Advanced Digital Signal Processing. This course provided the basis for studying the effectiveness of a research-oriented process on the Gold Coast campus. The author also wishes to thank the anonymous reviewers for invaluable feedback and suggestions.

Copyright statement