ENZCon 2012

conference proceedings

10th – 12th December

UNIVERSITY of OTAGO

DUNEDIN, NEW ZEALAND

Published by:
Electronics Group
Department of Physics
University of Otago
www.physics.otago.ac.nz/reports/electronics/ENZCon2012.pdf

ISBN: 978-0-473-23275-7 (Print)
ISBN: 978-0-473-23276-4 (CD)
ISBN: 978-0-473-23277-1 (Online)
Copyright remains with the author(s) of each paper contained in these Proceedings. Please contact the authors if you wish to reproduce their work. Responsibility for the contents of these articles rests upon the authors and not the editors nor ENZCon Inc.

Organising Committee:
Tim Molteno
Colin Fox
Bev Reynolds
Hilary Lawrence

Reviewers:
Andrew Bainbridge-Smith
Gary Bold
Will Browne
Dale Carnegie
Colin Fox
Michael Hayes
Ann Harlow
Hilary Lawrence
Ken Mercer
Tim Molteno
Pieter Nieuwoudt
Amal Punchichewa
Gerard Rowe
Jonathan Scott
Craig Watterson
Steve Weddell

Credits:
Cover image: 'Circuit Board Sneakers' gabrielshaw.com
Logo photo (back page): Colin Fox
LaTeX editor: Hilary Lawrence, Hyuck Chung and Colin Fox
using LaTeX's 'confproc' class by V. Verfaille

Welcome from Tim Molteno, Conference Chair

Welcome to the nineteenth Electronics New Zealand Conference. This year ENZCon is being held at the University of Otago, and is being hosted by the Electronics Group at the Department of Physics. At Otago, we pride ourselves in providing a good mixture of work and play, sometimes at the same time. The Poster session - with beer and pizza - is an example of this, and we hope it will be well attended.

We have four oral sessions covering a wide range of topics, including our first ever Special Session on Electronics and Engineering Education. Thank you to Dale Carnegie, who did much of the work to make this session possible, and will also Chair the session. Other sessions continue to show that Electronics is moving with the times in New Zealand, and include presentations from first time student presenters, as well as industry and academic presenters.

Thanks to Bev Reynolds for her great organising and administration skills. Hearing "Oh, I've already done that" in response to several last minute requests has made the rest of the committee's job far easier. Special thanks to our sponsors RFTest Solutions, Hoare Software Research, Nichecom and Electrotest. The input of the Trade Sponsors helps keep ENZCon affordable and accessible.

We hope that everyone has a fun, informative conference.

Dr. Tim Molteno,
Conference Chair,
Hilary Lawrence
Organising Committee

Conference Program

Day 1 : Mon 10th Dec.

Registration at St David Seminar Room 6, 12:00 - 13:00

Opening Address by Gary Bold, 13:00

Oral Session 1 (13:15): Electronics and Engineering Education, Chair: Prof. Dale Carnegie

1 Threshold Concept Knowledge in Analogue Electronics: Support and Assessment
A. Harlow, M. Peter and J. Scott

7 NCEA as a Predictor of Success in Tertiary Engineering Study
D.A. Carnegie, C.A. Watterson, P. Andreae and W.N. Browne

13 Prototype of a New Engineering Masters Project Model: Working with Marketing and Software Faculties to Commercially Kickstart University Research
J. Scott, B. Chang, C. Jameson, L. Payne, D. McLaren, A. Barr-Smith, C. Came and C. Hechei

19 Encouraging Success – measure, mentor and engage
G.B. Rowe, M. Shepherd and C.R. Smaill

25 Teaching introductory embedded systems with the UCFK4
M.P. Hayes

31 Increasing Student Engagement in First Year Engineering
C.A. Watterson, W.N. Browne and D.A. Carnegie

Day 2 : Tue 11th Dec.

Oral Session 2 (09:00): Sensors & Wireless, Chair: Dr Micheal Hayes

37 Pose Tracking using Ultrasonic Sensors in a Small Scale Enviroment
R.SH. Liu, C.K. Wong and P.P.K. Lim

43 AWSAM – Miniaturised Wireless Nodes for Structural Health and Environmental Monitoring
A. Ivoghlian, K.I-K. Wang, Z. Salcic

49 How to protect against RSA side-channel attacks using redundant number hardware
I.L. Manuel and A. Bainbridge-Smith

55 Comparison of Time and Frequency Domain Methods for Analysis of Bearing Damage
H.J. Lawrence and T.C.A. Molteno

61 Estimation of Indoor Local Average Power of Wireless Sensor Network
K. Yu, S.M. Berber and M-J. Ho

Oral Session 3 (14:00): Imaging & Inference, Chair: Assc Prof. Colin Fox

67 Compressively Sensed Hyperspectral Image Recovery using Total Variation Minimisation by Approximation
D. Eason, W.Y.L. Lee and M. Andrews

73 Removing Floor Reflections to Improve Robot Navigation with a 3D-TOF Camera
S.G.J. Kelly and A.A. Dorrington

79 A GPS Signal Synthesiser for Systems Integration Testing of the Transient Array Radio Telescope
C.F. Shaw and T.C.A. Molteno

85 Tree Twist: a spiral grain scanner for increment cores
M.P. Hayes and N. Pink

Poster Session Rm G13-G14 Science III Building (18:00)

91 Compressive Hyperspectral Image Sensing Restoration via Joint TV-L1 Regularisation
W.Y.L. Lee, D. Eason and M. Andrews

97 Cascomp BJT Amplifier vs. Traditional Configurations
H. Jull, T. Balsom and J. Scott

103 Mobile SPOT: A Wireless Sensor and Actuator Node for Applications based on Internet of Things
U.D. Atmojo, Z. Salcic and K.I-K. Wang

109 Vision-based Positioning of Robots
A.A. Evina and C. Tsui

115 Efficient Feature Based Visual Servoing for Onboard Computer Vision
S. Duncan, M.P. Hayes and A. Bainbridge-Smith

121 Bayesian Inference of Position from Short Samples of GPS Signals
T. Holmes and T.C.A. Molteno

129 Low-power Telemetry System for Wildlife Tracking
B. Panckhurst, P. Brown and T.C.A. Molteno

135 Bachelor of Engineering Re-design with a Project Spine at Massey University
A. Punchihewa

141 Experiments with local moves for MCMC applied to electrical impedance tomography
E. Ma and C. Fox

Day 3 : Wed 12th Dec.

Oral Session 4 (09:00): General Electronics, Chair: Dr Tim Molteno

145 MechBass: A Systems Overview of a New Four-Stringed Robotic Bass Guitar
J. McVay, D.A. Carnegie, J.W. Murphy and A. Kapur

151 Adaptive Noise Cancellation for the Electroencephalogram
A. Kandasamy, R.D. Jones and S.J. Weddell

157 Use of Kinect Sensor in a Robotics Development Project
P.F. Nieuwoudt, R. Mackill and R. Kumar

163 Real Time Monitoring of Cat Feeding Behaviour
S. Al Souti, D. Bailey and D. Thomas

169 **Index of Authors**

Threshold Concept Knowledge in Analogue Electronics: Support and Assessment

Ann Harlow, Mira Peter*, and Jonathan Scott#

*Wilf Malcolm Institute of Educational Research
Faculty of Education
Univerisity of Waikato
Hamilton, New Zealand
e-mail: aharlow@waikato.ac.nz

School of Engineering
University of Waikato
Hamilton, New Zealand

Abstract: In electrical engineering, as in other academic disciplines, there exist special, threshold concepts, where students often get stuck but which once grasped reveal new ways of thinking about a subject. Two surveys, student interviews and focus group discussions, and students' assessment were directed at learning of threshold concepts and their pre-cursors. Results suggest that one of the precursor concepts, current flow, may be a threshold concept in itself. A model of precursor and threshold concepts assessment and additional student-support for learning threshold concepts is suggested.

Keywords: analogue electronics, threshold concepts, learning support, assessment

1 INTRODUCTION

The threshold-concept theory developed by Meyer and Land (2003) has motivated a growing interest in contemporary research on teaching and learning of threshold concepts. According to the theory, in each academic discipline there exist special concepts that once grasped reveal new and previously inaccessible ways of thinking about that subject. Threshold concepts have been linked to ontological shifts (Meyer, Land, & Baillie, 2010), changes in identity, and shifts in subjectivity that come with the reconfiguration of a learner's prior conceptual framework. These changes are central to what it means to become an artist, economist or engineer. Alas, it is threshold concepts that students often find troublesome and where they frequently 'get stuck' (Davies, 2006; Harlow, Peter, Scott, & Cowie, 2011; Meyer & Land, 2005, 2006). Arguments have been made (Meyer and Land, 2003; 2005) that after first encountering threshold concepts, students may spend time in a 'liminal' state prior to crossing the threshold for deep understanding. Until students cross the threshold they are only able to mimic deep understanding and repeatedly fail to solve conceptually identical problems when these appear in new contexts. Thus, it is crucial to uncover why some students find it troublesome to understand and to express knowledge of threshold concepts; likewise, it is important to determine how students undergo a transformational, or even a creative, experience in the liminal space of learning.

Since the formulation of the threshold concept theory, researchers have been focusing on identifying threshold concepts in their field using the attributes identified by Meyer and Land (2003): threshold concepts are *transformative* (they change the learner's whole way of thinking), *irreversible* (they are hard to unlearn, as riding a bicycle), *integrative* (they connect into many diverse niches of a discipline), *bounded* (they mark the edge of a discipline), and they are potentially '*troublesome*' (difficult to grasp; counter-intuitive).

The research reported in the present paper builds on our previous and recent work on the identification of threshold concepts in the first-year electronics engineering and on our investigation of the impact of a threshold concept-informed curriculum and pedagogy on students' learning and retention in electrical engineering (Scott, 2010; Scott, Harlow, Peter, & Cowie, 2010a, 2010b; Scott, 2012). In our previous research five 'key' threshold concepts were identified and two were selected for a renewed curriculum for the first-year paper which was piloted in Semester A of 2010 (Scott, Harlow, Peter & Cowie, 2010a). The outcomes of our mixed-method evaluation study over three years (2010 - 2012) have been used to inform further reviews of the course, including further consideration of the threshold concepts. We have examined the relationships between students' achievement and their perceptions of the challenge of learning and understanding the identified threshold concepts. Additionally, we have explored the affordances of various threshold-concept-teaching opportunities and the students' perceived barriers to learning (Harlow, Peter, Scott & Cowie, 2011; Harlow, Scott & Peter, 2012).

Collectively, the outcome of these investigations has resulted in a revised curriculum to focus on stepping through these targeted concepts using a scaffolded, problem-based learning approach. The findings suggest that the changes in pedagogy, based on threshold concept theory, are worth pursuing as they have the potential to address current concerns around the declining number of tertiary graduates in electrical engineering world-wide (Entwistle, Nisbet, & Bromage, 2004).

Our goal in the present research was to investigate the impact of using the *Immediate Feedback Assessment Technique* (IF-AT) (Epstein et al. 2002) on students' learning of threshold concepts and to tap into students' understanding of threshold concepts and their precursors.

By the term 'precursor concept' we refer to an idea that must first be mastered by the student before a particular question addressing a threshold concept can itself be understood. The precursor concept might itself be a threshold concept, but it might equally be simple knowledge. As a trivial example, we might seek to test a student's understanding of friction in the context of thread and knots by means of a question about the action of an overlocking stitch. Such a question would fail if the student did not know the particular action of an overlocking machine, and what constitutes an overlocked stitch. The knowledge of overlocking is necessary to understand the main question. In this example the dependence is obvious, but particularly when the precursor is a threshold concept that dependence can be less conspicuous. We might postulate that this arises either because the questioner is less likely to register that the student does not understand the precursor idea, or because the deep interconnections of threshold concepts make the dependence distant and subtle. In our work, we have noticed that the danger can lie in assuming that because the student passed a previous course involving the precursor threshold concept that they have passed the threshold, and understand it.

2 METHOD

In the first-year analogue electronics course there are typically 100 to 140 students, some of whom are international students. Students come from a variety of disciplinary backgrounds across the spectrum of science subjects taught at secondary school. There are also students with work experience in technical fields. In the present paper we did not determine if results differed by subgroup as that was not the focus, however, this analysis is definitely in our future plans.

Theories and problems are introduced in weekly lectures which provide motivation and a view of the 'big picture'. Most of the class time is spent on 'tangibles', 'ponderables', and 'visibles', which are essentially hands-on activities, interesting questions and problems, or simulations, conducted in labs where students have hands-on experiences with equipment to solve problems, and in problem-based tutorials.

For this research we designed two surveys—one at the start of the course and one in the last week of the course—tapping into students' perceptions of their understanding of taught threshold concepts, and activities that helped them learn. Of the 119 students enrolled in the first-year analogue electronics course, 105 took part in survey 1 and 69 took part in survey 2. Results from these surveys were combined with students' mid-term and final examinations achievement scores on threshold concept and precursor to threshold concept questions to obtain a better understanding of student knowledge of threshold concepts.

Using various concept inventories that have been created in similar fields (Evans et al., 2003; Hestenes, Wells & Swackhamer,1992), the lecturer built a 'concept inventory' to cover the three-year course (Scott, Peter & Harlow, 2012). In order to maximise first-year students' understanding of threshold concepts the lecturer focused on only two threshold concepts, Thévenin's theorem and dynamic resistance.

To assess students' understanding of threshold and precursor concepts the lecturer used IF-AT testing system which provides students with immediate feedback about the accuracy of their answers and allows students to continue answering a question until they discover the correct answer (see Figure 1). Using IF-AT the lecturer created an interactive learning opportunity for students and a more informative assessment opportunity for the lecturer.



Figure 1: An IF-AT question and 'scratchie' card

3 RESULTS AND DISCUSSION

The aim of this research was to address several issues.

3.1 What counts as 'troublesome knowledge' in analogue electronics?

Survey data from 69 students provided an answer to this question. Concepts that students found the most troublesome were: dynamic resistance (67%), Thévenin (55%), transistor action (52%), and op amps (51%).

> *Dynamic resistance of diodes is difficult. It is the resistance of diode is changing as you change the current and voltage, which is not normal. I think it is probably because when I look at the question, they say 'go and use the source and the voltage…' To my mind I see several things that could be the source voltage, so for example I could be using a dc source and an ac, and I think 'is it supposed to be the dc or the ac one I am supposed to be using?'(year-one student explaining why he found dynamic resistance troublesome)*

Support from staff, particularly in the lab and tutorials, was judged to be helpful by all students. Students, in particular those who did well in the course, liked the lecturing style (66%), and both those who did well and those who did not do well in the course liked the IF-AT tests (57%): perhaps for different reasons. Working in pairs was appreciated the least by students

who had done well in the course (29% of A-level students) in contrast to other students, especially those who had not done so well (i.e., 54% of B-level students, 56% of C-level students and 71% of D-level students). More fully-worked examples, particularly in lectures, were appreciated and requested by all students. Some students (19%) noted that there did not seem to be a reliable textbook that they could mine for examples of the conceptual problems they found to be troublesome.

3.2 The assessment tool

The idea that assessment should be able to identify variation in progress (Meyer and Land, 2008) inspired our search for a way to assess threshold concept understanding for the purposes of re-teaching or extra tuition for those who need it.

Data from the mid-semester assessment was used in this analysis. The assessment that was designed in 2012 using the IF-AT method was able to isolate sources of conceptual difficulty that related to students' prior knowledge. Without specific prior knowledge (e.g., holistic current flow, graph interpretation and 3D representation) students could not be expected to understand the more complex modelling required by Thévenin's theorem (Scott & Harlow, 2012). In addition to poor basic understandings, (e.g., holistic current flow, graph interpretation and 3D representation) students who had difficulty with or guessed the answers to the threshold concept questions had little idea of the inter-relationships between concepts (e.g., some students could not relate the pictures of equipment to a drawn circuit diagram). By assessing understanding of the pre-cursor concepts and of the threshold concept in the same test it was easy to identify students who might have guessed the answer to a threshold concept question.



When the switch S1 closes in the circuit above, the current in R1 will: a) Increase, b) Stay the same, c) Decrease, d) Fall to zero, e) There is not enough information in the question to tell.

Figure 2: Example of pre-cursor exam question that assesses student understanding of current flow

Seventeen students who failed to correctly answer precursor questions (i.e., did not get full marks) but answered the threshold concept question correctly, (i.e., got full marks) were interviewed to gain insight into their understanding of threshold concepts. In all cases, the students who had failed to answer the pre-cursor questions correctly had guessed the more difficult threshold concept questions or used a process of elimination. The results suggest that these students neither understood the threshold concepts nor the pre-cursor concepts.

Without the necessary prior knowledge, these students would have likely failed to complete the course successfully—even if they passed the course there would be a high probability that they would continue their studies without understanding the threshold concept. As Meyer and Land (2008) put it "the necessary pre-liminal ontological shift for the programme was not deemed to have taken place." (p.75).

In 2008, Meyer and Land suggested that traditional assessment should be abandoned since students may often get the right answer yet retain fundamental misconceptions. Instead, assessment should reflect the way students make an ontological shift when they grasp a threshold concept (i.e., when students internalise knowledge and think from the point of view of an expert). Still, in designing such assessments the question remains: How do we know that our students have 'got it' or when they 'became unstuck'?

3.3 Entry-level physics

For this analysis we have combined data from 105 students who completed survey 1 with their achievement on the final exam questions and their final grade for the course. Results revealed that students with NCEA Level 3 physics on average achieved at a similar level in the final exam as those those who did not have NCEA Level 3 physics, $t(97) = 0.85$, $p > 0.05$.

Similarly, high school physics did not ensure the correct answer to a precursor question—only 71% of students *with* NCEA Level 3 physics answered the precursor question correctly, and a lack of high school physics did not prevent students answering correctly the threshold concept questions—75% of students *without* NCEA Level 3 physics answered a threshold concept question correctly. Note that students who achieved NCEA Level 3 physics may not have necessarily studied electronics as part of their high school physics. Regardless, results indicate a need for a lecturer to dedicate initial lecture(s) to precursor ideas, such as current flow, and to assess if students have grasped these basic ideas some time prior to mid-semester examination.

3.4 Pre-cursor concepts

For this analysis we have combined data from 67 students who completed survey 2 with their achievement on the final exam questions.

At the end of the course, 62 students rated their understanding of current flow, a precursor concept for Thévenin's theorem, as good or excellent while only five students felt that they did not understand this concept. However, the results of their achievement data revealed that many students did not understand the idea of holistic current flow, even at the end of the course. Namely, on 10 exam questions, out of 15 that required an understanding of current flow, a greater proportion of students provided an incorrect answer than a correct answer regardless of how students rated their understanding of current flow (e..g., have trouble understanding current flow vs. no trouble understanding current flow). For example, of 62 students who said that they had *no trouble* understanding current flow a significantly greater proportion gave an incorrect answer

to six questions requiring an understanding of this concept.

The findings show that it is essential that assessment (particularly if the assessment consists of a multi-choice test) consists of both the precursors and the threshold concept questions as this combination would provide the most comprehensive view of the students' grasp of threshold concepts.

3.5 Perceived understandings

For this analysis, data from 64 students who completed survey 2 were combined with their achievement on the final exam threshold concept questions requiring an understanding of Thévenin's theorem and their total year grade. Thirteen students rated their understanding of Thévenin's theorem as excellent or very good, 27 students rated their understanding as good, and 24 rated it as limited or none.

Analysis of students' achievement on the threshold concept questions revealed that on 3 out of 5 questions a significantly greater proportion of students who rated their understanding of Thévenin's theorem as excellent or very good more frequently answered those questions correctly than the other two groups of students (100% vs. 52% and 87%, $\chi^2(2) = 14.53$, $p < 0.05$; 71% vs 33% and 42%, $\chi^2(2) = 5.55$, $p = 0.06$; and 93% vs 78% and 42%, $\chi^2(2) = 12.74$, $p < 0.01$). However, regarding the other two Thévenin questions in the examination there were no significant differences in the proportion of correct answers for the three groups of students. For one of the questions, all three groups of students had a high proportion of correct answers (93%, 93%, and 100%) and for the other, all three groups had a very low proportion of correct answers (29%, 15%, 17%).

Results (see Table 1) suggest that although a perceived understanding of the threshold concept was not a good indicator of a student's overall grade for the paper, it was a good indicator of how well students answered questions involving the threshold concept (Thévenin's theorem).

	Perceived understanding of Thévenin		
	A group	B group	C group
	Excellent, very good	Good	Limted, none
Average grade	**73%**	60%	64%
Thévenin exam questions correctly answered	**88%**	54%	57%

Table 1: A group more frequently answered Thévenin questions correctly than B and C group

3.6 Interim assessment

For this analysis, mid-semester achievement and final examination achievement data from 119 students were combined. The results revealed that there was a a small but significant correlation between students' achievement on the mid-semester assessment and their achievement on the final examination $r(119) = 0.4$, $p < 0.05$. and between their mid-semester assessment and their overall grade $r(119) = 0.5$, $p < 0.001$. Thus, mid-semester assessment (interim assessment) was only an indication of potential examination marks. Lab-book marks were more indicative of high/low total grades than interim assessments.

These findings may have implications for the type of assessment used during a course. Formative assessment using an IF-AT test was trialled with year two students as a learning exercise. Students worked in groups of three or four debating their understandings and producing a group answer. This exercise was followed by the lecturer going over the most difficult questions, making sure that the threshold concept questions were explained in detail. Students found this way of assessment for learning to be both engaging and useful.

4 GENERAL DISCUSSION

The findings show that it is essential for students to have a good understanding of the precursor concepts and that these need to be assessed early in the course. The understanding of basic concepts cannot be underestimated for learning more complex concepts. In fact, given the results reported here, it is possible that holistic current flow may actually be a threshold concept.

With different learners taking different routes and time to fully grasp the threshold concept, whereby they are transformed from thinking like a novice to thinking like an expert, the questions of how we might assess this process arise. Will passing a test, or the final examination at the end of the course mean that those who passed have grasped the threshold concept? Given that it is well established that learners may spend a considerable length of time in an unstable, transformational conceptual space before they fully grasp troublesome knowledge, should we even be considering assessment of a threshold concept in a student's first year?

From our findings over two years we have come to an understanding that students who are introduced early to a threshold concept will take on board whatever aspects of that threshold concept that they can, based on their prior knowledge, their propensity for follow-up study, and their ability to articulate their understandings. They find that encountering the threshold concept again and again, across different and appropriate contexts allows them the necessary time and space to follow through the logic yet again, to recall aspects of what they did the previous time, to recognise that a problem needs an application of the concept, and eventually to 'come to know' or to 'own' the knowledge so that they wonder why they found it so troublesome at the start (Harlow, Peter, Scott & Cowie, 2011). This process is not necessarily linear, and it may take well into the second year before 'the penny drops' or 'the light-bulb glows'.

> *There are a few things that are difficult like Thévenin, that is highlighted as being tricky and he spends quite a few lectures just so that we can pick it up. When we first did it I had no idea, and thought, "Sure, I'll take your word for it." But now I can do most of it on my own. I would recognise it now – it's hard, but I can do it. It is just remembering how to simplify the circuit down, because everything has its own set of rules – like whether it's in parallel or series, there is a lot to remember. Once you pick it up it's OK, but when you are learning it is a struggle. (year-one student explaining why it took time to grasp Thévenin)*

This is particularly relevant to the chosen threshold concept, in this case, Thévenin's theorem, as students were introduced to it when first learning circuit theory early in year one, and it was not until year two when transistors were introduced that many students said 'Thévenin clicked'.

5 FURTHER WORK

The design of a comprehensive test is now ending its first year of trialling and to date the lecturer has found it extremely difficult to ensure a reliable and valid test of student understanding. The lecturer is continuing to refine the questions in his assessments and to devise ways of correlating the threshold questions with the precursor questions. He hopes to show that assessment of the more troublesome areas in analogue electronics is all that is necessary to know if a student has had a transformative experience and will re-enrol as a successful student in year two. The key to unlocking what the students are actually doing when they select an answer has been found by interviewing students, which may not be practical without the researcher's help. So the next step is to refine the questions more so that each response gives an indication of where the student has gaps in his knowledge. This may be done in conjunction with ensuring that each threshold concept question is associated with a precursor question, and only when both are answered correctly, does the student receive full marks and is deemed to have understood the troublesome knowledge.

While the lecturer concentrates on refining the assessment tool for year one students, our research focus to support a summer scholarship student to on develop online tutorials in the areas that students are known to have difficulty understanding. An electronics engineering graduate student, who has shown an interest in student learning by demonstrating in labs and tutorials since the start of this project, will be working to design and then monitor online tutorials for students who, regardless of their achievement grades all say they want extra worked examples to help them learn. Some of the findings of the way students prefer to work will be taken into consideration in the design of the online tutorials. Student access data will allow us to track student learning behaviours which may inform an improved teaching and assessment approach.

6 CONCLUSION

Our work has been conscious of making the threshold concepts the 'jewels in the curriculum' (Land & Meyer, 2008). They have helped the lecturer to bring into focus the destination points for students, but as threshold concepts are known to take time to be understood, students continue to have difficulty with them in the first-year course. We are coming to the conclusion that there is a need for more support for precursor knowledge in electronics to be taught at the year one level.

7 ACKNOWLEDGEMENT

The authors wish to thank the Teaching, Learning and Research Initiative (TLRI) for ongoing funding to support this work, now part of a two-year research project to investigate the teaching and learning of threshold concepts across disciplines.

REFERENCES

Davies, P. (2006). Threshold Concepts: How can we recognise them? In R, Land & J.H.F Meyer (Eds.), *Threshold concepts and troublesome knowledge* (70–84). London, England: Routledge.

Entwistle, N., Nisbet, J., & Bromage, A. (2004). *Teaching-learning environments and student learning in electronic engineering.* Paper presented at Third Workshop of the European Network on Powerful Learning Environments, Brugge.

Epstein, M., Lazarus, A., Calvano, T., Matthews, K., Hendel, R., Epstein, B., & Brosvic, G. (2002). Immediate feedback assessment technique promotes learning and corrects inaccurate first responses. *The Psychological Record, 52*, 187-201.

Evans, D., Gray, G., Krause, S., Martin, J., Midkiff, C., Notaros, B. Pavelich, M., Rancour, D., Reed-Rhoads, T., Steif, P., Steveler, R., & Wage, K. (2003). Progress on concept inventory assessment tools. Session T4G, 33rd ASEE/IEEE *Frontiers in Education Conference* (FIE), November 5-8, 2003.

Harlow, A., Peter, M., Scott, J., & Cowie, B. (2011). 'Getting stuck' in analogue electronics: Threshold concepts as an explanatory model. *European Journal of Engineering Education*. DOI:10.1080/03043797.2011.606500

Harlow, A., Scott, J., & Peter, M. (2012). Students' perceptions of travel through the liminal space: Lessons for teaching. Paper presented at the *4th Biennial Symposium on Threshold Concepts,* Dublin July 27-29, 2012.

Hestenes, D., Wells, M., & Swackhamer, G. (1992). Force concept inventory. The Physics teacher. 30(3), 141-158.

Meyer, J. H. F., & Land, R. (2003). Threshold concepts and troublesome knowledge: Linkages to ways of thinking and practising within the disciplines. In C. Rust (Ed.), *Improving student learning. improving student learning theory and practice – 10 years on*

(pp. 412–424). Oxford, England: OCSLD. Retrieved from http://www.tla.ed.ac.uk/etl/docs/ETLreport4.pdf

Meyer, J., & Land, R. (2005). Threshold concepts and troublesome knowledge (2): Epistomological considerations and a conceptual framework for teaching and learning, *Higher Education, 49*(3), 373–388.

Meyer, J., & Land, R. (2006). *Overcoming barriers to student understanding: Threshold concepts and troublesome knowledge*. Oxon, England: Routledge.

Meyer, J.H.F. and Land, R. (2008) Threshold concepts and troublesome knowledge (5): dynamics of assessment", paper presented at the 2nd International Conference on Threshold Concepts: Threshold Concepts from Theory to Practice, Kingston, Ontario, Canada, 25–27June.

Meyer, J., Land, R., & Baillie, C. (Eds.). (2010). *Threshold concepts and transformational learning.* Rotterdam, Netherlands: Sense Publishers.

Scott, J. (2010). Over the threshold: learning and student retention in electronics courses. *TDU Talk, Issue 5*, (18-19). (Special issue of University publication on teaching in science).

Scott, J., & Harlow, A. (2012). Identification of threshold concepts involved in early electronics: Some new methods and results. *Australasian Journal of Engineering education, 18*(1), 1-8.

Scott, J., Peter, M., & Harlow, A. (2012). An electronics threshold concept inventory: Assessment in the face of the dependency of concepts. Paper presented at *The IEEE International Conference on Teaching, Assessment and Learning for Engineering (TALE2012)*, Hong Kong, 20-23 August, 2012.

Scott, J., Harlow, A., Peter, M., & Cowie, B. (2010a), Exploring threshold concepts in electronics engineering. Presentation at the *Third Biennial Threshold Concepts Symposium: Exploring transformative dimensions of threshold concepts,* The University of New South Wales and the University of Sydney, Australia. Retrieved from http://www.thresholdconcepts2010.unsw.edu.au/Abstracts/Scott_Harlow_and_Cowie.pdf

Scott, J., Harlow, A., Peter, M., & Cowie, B. (2010b). Threshold concepts and introductory electronics. In Proceedings of the *2010 Australasian Association for Engineering Education Conference*, The University of Sydney, NSW, Australia. Retrieved from http://researchcommons.waikato.ac.nz/bitstream/10289/4917/1/ThresholdConcepts.pdf

NCEA as a Predictor of Success in Tertiary Engineering Study

Dale A. Carnegie, Craig Watterson, Peter Andreae, Will N. Browne
School of engineering and Computer Science
Victoria University of Wellington
Wellington, New Zealand
Dale.Carnegie@vuw.ac.nz

Abstract: The New Zealand Government is moving towards restricting access to tertiary education and implementing a managed entry scheme. It is therefore important to be able to predict whether a student has a reasonable likelihood of succeeding in tertiary engineering study. New Zealand secondary schools mostly operate on the National Certificate of Educational Achievement (NCEA) model whereby subjects are assessed on the basis of discrete individual modules. The paper compares the NCEA results to first year grades in tertiary engineering subjects obtained from many New Zealand providers of the BE degree to determine whether these NCEA grades can be used as a predictor of success or failure in tertiary engineering programmes. This is the first nation-wide survey of its kind and has yielded surprising results. For example, predicting success based on whether a student has *Not Achieved* is more insightful than basing on achievements, which is counter to the basis of many school league tables worldwide.

Keywords: prediction of engineering success, engineering preparation, managed entry

1 INTRODUCTION

Victoria University of Wellington (VUW) began offering a professional Bachelor of engineering degree in 2007 in these specialisations of: engineering and Computer Systems (ECEN), Networking (NWEN), and Software engineering (SWEN).

In the first year of engineering study all students engage in computer programming and algorithms courses. ECEN students heavily engage in mathematics (including calculus) and physics. NWEN and SWEN students engage in discrete mathematics, statistics and an introductory physics course.

Victoria University of Wellington (VUW) has a mainly open-entry policy to its Bachelor of engineering (BE) degree – anyone meeting the general university entrance requirements with satisfactory secondary school performance in mathematics (and for one specialization, also satisfactory performance in physics) can enter the first year of the BE programme. In order to maintain the quality of the graduating students, students must attain a B grade average (65% or above) over their first year engineering subjects in order to be able to progress.

The consequence of this B average grade requirement has been that only about 40% of students who initially enroll in the BE degree progress, and of these about one third will take more than one year to achieve this B average over the required courses. This attrition rate was historically acceptable to the university due to the "bums-on-seats" Government funding model that encouraged universities to enroll as many students as possible, and even to actively compete against other universities in order to obtain an increased share.

Our Government has now dramatically altered its funding structure and universities now have a cap placed on their enrolment figures and are not only *not* funded for any additional students they enroll, but can be actively penalized by way of a substantial financial fine. Additional penalties are imposed for any courses that do not pass 55% (moving to 60%) of the Government funded enrolments. As a result, VUW adopted a managed entry policy, where access is restricted to those who have performed "adequately" in their final year of secondary school study. The determination of what qualifies as performing "adequately" is presented in section III, "Secondary School Grading Scheme". A positive outcome of this is that there is now an increased effort to determine what factors influence student success in engineering study.

A complicating factor is that the primary entry qualification from New Zealand secondary schools is the National Certificate of Educational Achievement (NCEA) qualification. Since this system is unique to New Zealand, there are no international benchmarks that can be used even as a precursor to forming a useful predictor. Consequently this paper outlines a prediction approach based on a national survey of first year engineering stu-

dents. The prediction quality of NCEA has been found to be poor, and hence it is recommended to develop bespoke diagnostic tests to increase accuracy of prediction.

At the date of this paper preparation, an updated correlation and data-mining study is underway. It is expected that these results will be available in time for the for the Conference presentation. An extended version of this paper was presented at EDUCON2012, Marakech, Morocco [1].

2 NCEA

This paper assumes that the reader is familiar with the NCEA curricula and grading offered in the bulk of NZ secondary schools. If not, information is easily obtainable on-line [http://www.nzqa.govt.nz/qualifications-standards/qualifications/ncea/], otherwise the reader is referred to [1].

There are no dedicated engineering subjects in the NZ secondary school curriculum. Prior to 2011 there were many "technology" subjects. At first glance there is an impressive range of these, including subjects such as biotechnology, electronics & control technology and information & communication technology. However an examination of the assessment criteria yields little difference between these subjects and other technology offerings such as materials technology or food technology. Such assessment includes phrases such as "Develop a means for multi-unit production of a technological outcome" (AS90352), "examine how technological practice is influenced by responsibilities to the wider community" (AS90773), "develop a proposal for a production process for a client" (AS90792). The vague nature of these criteria translate into a very varied offering of these subjects amongst secondary schools and the quality of provision is often heavily dependent upon the skills and enthusiasm of individual teachers.

Importantly, until very recently, secondary schools have not been producing students with experience in computer programming. "Computer studies" is offered at most schools, but this deals more with computer applications and the acquiring of a working knowledge of computer development tools rather than instructions on how to program. These Technology orientated courses are discussed further in the following section.

"Core" mathematical and science subjects are assessed against *achievement standards* that use the grades of "Not Achieved", "Achieved", "Merit" and "Excellence", whilst most of the modules in the technology area (which includes computer studies) are assessed against *unit standards*, which only use the pass/fail grades of "Not Achieved" or "Achieved".

In practice, a significant problem arises from this: high performing students typically avoid unit standard subjects and aim for subjects for which Merit and Excellence grades are available. This immediately means that our best students tend to avoid the subjects that under the current curricula format might have more of an engineering philosophy.

The grading into "Not Achieved", "Achieved", "Merit" and "Excellence" is also "non-linear" – Excellence is generally hard to obtain, and is not equivalent to the A- to A+ range of university grades. Furthermore, there have been significant criticisms of the grading scheme, because of inconsistencies in the standards applied from year to year, a significant degree of subjectivity in the awarding of an Excellence over a Merit, rigid marking standards that have been applied that favor regurgitation of material rather than creative thinking and that heavily penalize students with excellent understanding but minor inaccuracies or carelessness.

VUW has trialed a managed entry model. Whereas previously students merely had to meet the "University Entrance" requirement (42 level 3 credits over three areas, plus literacy and numeracy requirements), for guaranteed entry, students now must also score at least 150 points based on their NCEA level 3 scores. This score is calculated on the basis of student's BEST 80 level 3 NCEA credits in approved subjects. The credits are then weighted by their grade – 4 points for Excellence, 3 points for Merit, and 2 points for Achieved. Students with less than 150 points are put on a waiting list, and only admitted if there are still places after those with guaranteed entry have been admitted.

While this scheme is a little more selective than University Entrance, it is not specific to any subject area, and does not have significant predictive power of success in engineering. What is needed is a more precise predictor of success in engineering study based on secondary school NCEA results that could provide effective guidance to both prospective engineering students and to the university staff who must advise and admit them.

3 FORMING A PREDICTOR

We chose to limit the forming of a predictor to the "digital" areas of engineering, specifically Electrical/Electronic, Computer Systems, Software, Networking and Mechatronics. The "digital" label is not completely accurate, but serves to differentiate these areas of engineering from areas such as Civil, Mechanical or Chemical engineering. The motivation for this restriction stems primarily from the engineering offerings at the authors' host institution, but importantly also because this collection of digital engineering specializations shares common relevant subjects.

Beginning this study, we expected that the strongest predictive subjects would be the core science papers, particularly those that directly lead into engineering study, specifically physics and mathematics (with calculus).

To inform this study, data were obtained from several New Zealand universities that offer a BE degree in these "digital" engineering subjects. This also ensured our data would not exhibit any regional or university bias. These data consist of the NCEA level 3 results and the first year course grades for every student who enrolled in digital tertiary engineering study at VUW in 2009, 2010 and 2011, and at the University of Canterbury, Massey University and the University of Waikato in 2009.

To keep the data manageable, we removed NCEA results for subjects that were most distant from engineering, and analyzed only science, technology, numeracy and literacy relevant NCEA subjects, specifically; Mathematics with Calculus, Mathematics with Statistics, Physics,

Chemistry, Biology, English and Computer Studies. Of these subjects, Computer Studies is assessed using the "unit standard" pass/fail model; all other subjects have the Achieved, Merit, and Excellence grades of achievement standards.

We only considered the first year university grades because the variabilities of the different tertiary institution's offerings after the first year made comparisons extremely difficult.

Attempting to predict tertiary success on the basis of secondary results is not new. Several international and national studies have examined if academic achievement at secondary school was an indicator of success at tertiary level. Two Australian studies, Everett and Robins [2] and McClelland & Kruger [3], found that school results predicted some performance in relevant tertiary (especially science) subjects. Similarly, in a report for the New Zealand Ministry of Education, Engler (2010) [4] noted that good secondary school performance in Mathematics with Calculus, Physics and Chemistry predicted some level of improved performance in tertiary engineering study. One problem with both studies is that the predictive quality of the associations is not strong.

Our research differs from these previous studies in that we consider more than just a success/failure in engineering study and because we have full access to NCEA data, and hence are able to look for correlations within individual modules as well as within a holistic subject.

4 CORRELATION GRADE ANALYSES

The NCEA data was prepared to run a series of correlation and data mining analyses. For each of the subjects listed in the previous section, the total number of Achievement, Merit, and Excellence credits for each module were calculated. Due to the perceived subjective differentiation between Merit and Excellence grades, these two figures were combined.

Realizing that a predictor could be formed from a combination of subjects rather than on any one specific subject, separate value fields were created with broader subject groupings. Specifically these combined

(1) Physics and calculus
(2) Calculus and statistics
(3) Physics, calculus and statistics
(4) All science subjects

Each of these groupings was characterized by both the number of Achievement credits and the sum of the Merits and Excellence credits. These NCEA data were compared to the GPA (grade point average) of their overall first year university GPA grade.

To form an initial indication of the relevance of each NCEA subject to tertiary engineering success, a simple correlation calculation was performed. It was initially run over the entire data set of all students. The set was then divided into two, one for students in ECEN type specializations (that is those areas of tertiary "digital" engineering that require calculus and physics) and the other for students in NWEN/SWEN type specializations (that do not require calculus or advanced physics).

For the entire data set, the highest correlation score was 0.612 which related the sum of all the science Merit and Excellence scores (this excludes Achieved level scores) to GPA. It is not surprising that in general a student who performs well overall in secondary school science should tend to perform well in tertiary engineering. However, breaking the data into the two subdivisions gave a different result. For the ECEN students, this total of science Merit and Excellence scores was still the top ranked correlation at 0.722, but for the NWEN/SWEN students, the science Merit and Excellence scores dropped to a correlation of only 0.448. This was the seventh highest correlation, though the highest correlation for these students was only 0.513, which was with the combined English Merit and Excellence scores.

For the ECEN students, the next highest correlation score was the sum of Merit and Excellence scores for physics and calculus at 0.691, followed by the sum of Merit and Excellence scores for physics, calculus and statistics at 0.668. The implication is that for hardware electronic/mechatronic engineering study at a tertiary level, a strong performance in the mathematics and physics subjects at secondary school is the best indicator for success, where a strong performance means a reasonable number of Merit and Excellence grades. However, there must be some caution associated with this claim. A correlation score around the 0.7 level is not very strong, and is definitely not strong enough to form a reliable predictor for making entry decisions.

Of particular note is that all of the measures that included the "Achieved" level in NCEA credits at secondary school resulted in an inferior correlation score. For example, the correlation score for all credits for science and mathematics subjects (including credits that were only Achieved), drops to 0.525 for all students, and drops to 0.620 for ECEN students only. This decrease is not surprising since we are now including students who were not up to the Merit or Excellence level, but indicates that just achieving physics and mathematics credits at secondary school may not be sufficient basis for tertiary study. However, the wide range of the "Achieved" grade makes generalizations problematic.

For the SWEN/NWEN students, there were no significant correlations between their NCEA results and university GPA. The correlation scores such as 0.513 for Merits and Excellences in English, 0.463 for all grades in science, 0.453 for all grades in physics and calculus, and 0.448 for Merits and Excellences in science demonstrate that none of these subjects appear to provide good preparation for, or good predictors of success in the computing focused engineering subjects.

Poor performance at university is not always due to academic reasons - a number of students who enroll at university "fail to engage", often for cultural, social, or financial reasons. It is possible to identify many of these students by the predominance of "E" grades over most of their first year courses. The goal of our analyses was to find good predictors of the potential for academic success at university. We were concerned that these low GPA's for non-academic reasons might have been confounding the correlation scores and masking the effectiveness of some of the predictors.

To explore this concern, we re-ran the correlation analysis on two further data sets, first removing all the students whose GPA was less than 1 and then removing all students with a GPA of less than 2. However the correlation scores dropped for both of these reduced data sets. Hence the students who we identified as failing to engage in tertiary study do not appear to be confounding the correlation scores, and eliminating them from analysis does not give more informative results

As a final correlation study, we wondered whether the NCEA grades could be used to predict "significant success", that is, could we identify those students who would have a B+ or greater average (GPA>6.0)? The answer was no, with the top correlation scores again being lower than for the full data sets.

A correlation score that was particularly surprising was the near-zero and at times even *negative* correlation between first year tertiary engineering GPA and secondary school Computer Studies for all streams of students. Given that computing is core to the SWEN/NWEN streams and required for all the digital engineering students, it might be expected that there would be at least as much correlation with secondary school Computer Studies as with all science subjects, if not more. The fact that good performance in secondary Computer Studies might even be associated with poor performance in SWEN/NWEN is rather worrying.

One part of the explanation is that the topics included in Computer Studies courses are highly variable between secondary schools, but most commonly, the emphasis is not on programming, but rather on computing applications and tools. Another factor, which may account for the negative correlation, can be found in the consequences of the "unit standard" assessment of the Computer Studies modules. Such subjects tend to be avoided by students aiming for university partly because of a widespread prejudice about the academic status of unit standards, and partly because high achieving students generally want the option of getting Merit or Excellence grades. Furthermore, the Computer Studies credits do not count for the first two subject areas for University Entrance (UE), and therefore constitute a greater risk for students for whom getting University Entrance is critical.

Although not the focus of this paper, we note that a brand new secondary school curriculum in digital technologies is currently being rolled out which is likely to change this situation significantly. Firstly, it includes modules on computer programming and computer science, in addition to the applications centered modules. Secondly, it will all be assessed using Achievement Standards, and will count towards UE in the same way as other subjects. We expect that over time, this will increase the academic status of computing in the secondary schools and also provide a much better foundation for university study in engineering and computer science.

In summary, for ECEN students there seems to be a clear indication that we should strongly consider Merit and Excellence scores in the science and mathematics subjects in the forming of a predictor. Further, physics and calculus generally score higher in the correlation calculations than chemistry, biology or statistics. For NWEN/SWEN students, the trend is not so clear. The correlation scores are significantly lower than the ECEN, and English rates highly in some tests, as does the total number of credits attained in science (rather than just the Merits and Excellences).

5 DATA-MINING ANALYSES

The correlation analyses discussed above were designed to try to find a predictor based on simple combination of NCEA subjects, with a division between ECEN and NWEN/SWEN students. The highest correlation score obtained of 0.722 for ECEN students and 0.513 for all NWEN/SWEN students is not sufficient to form a reliable predictor. A significant limitation in such a correlation analysis is that it relies on the human analyst to propose the appropriate groupings of subjects of grades (for example the combination of physics and calculus Merit and Excellence grades) to analyze. We also expect that different modules within a single subject may not be equally effective for predicting later performance. For example, within the mathematics subject, calculus related modules may be more relevant than geometry based modules.

To explore a richer range of relationships between NCEA grades, we performed a second set of analyses using tools from the Weka Data Mining Software from the University of Waikato. Specifically, we used the J48 Decision Tree Classifier with standard settings. An important property of this tool is that one can provide it with a large number of potentially relevant features and it will identify a small subset of those features that are most effective for discriminating the required classification. When using the tool, we were able to provide it with all the NCEA features, including grades for individual modules of courses as well as summary groupings.

We used two different criteria for building the decision tree classifiers. The first was a binary classification into students who's GPA would be at least a B average – i.e. a GPA score of 4.5 or above – against those whose GPA was less than 4.5. For VUW students, this corresponded to whether they would be allowed into the second year of engineering, and corresponded to a straightforward predictor of success.

However, for advising students, this binary classification is not necessarily the most useful result unless the predictor has very high reliability. A more useful classification is a ternary classification into those who would clearly fail to achieve our minimum B average (GPA < 3.0), those who were in the middle ground (who perhaps could succeed if appropriate intervention was in place) (3.0<= GPA<6.0) and those who would clearly succeed (GPA >= 6.0). If we can obtain reasonably reliable predictors for these categories, then the first group can be advised against attempting engineering, the last group can be confidently accepted, and the middle group can be advised carefully and provided with additional support.

An initial request to the software was to construct a tree for the binary criterion, using just the Guaranteed Entry Score (GES) feature (calculated using the system outlined earlier). As can be seen from Figure 1, the software identified a GES score of 153 as the most discriminating threshold (which is not dissimilar to VUW's determined score of 150). The numbers in the shaded rectangles indicate the number of students that fall into that classification, and the number of students incorrectly

classified. A "0" in the rectangle indicates the classifier predicts failing to meet the criteria of GPA ≥ 4.5, and a "1" that it predicts success (ie GPA ≥ 4.5). From Figure 1, we can see that of the 54 students that scored a GES under 153, 76% obtained less than a 4.5 GPA in their first year of engineering at VUW; of the 96 students who scored a GES over 153, 62% obtained at least a 4.5 GPA.



Figure 1. Simple classifier based only on GES score

Although this classification somewhat separates the students, it is not a sufficiently reliable predictor of success. We then constructed a tree using the ternary criteria, hoping to find clearer predictors. Figure 2 presents the results of this. In this, and subsequent figures, "0" in the shaded rectangle refers to a clear fail classification, "1" middle ground, and "2" clearly succeed. Students scoring above 227 on the GES ranking are predicted to clearly pass (6 incorrectly classified from 20 students), those between 170 and 227 would be "middle ground" (31 incorrectly classified from 61 students), and those scoring below 171 are identified as failing (27 students incorrectly classified).



Figure 2. Ternary Classifier based on GES

It is clear that these predictors are also insufficiently reliable to be used for advice or admissions. Not surprisingly, the GES alone is not good enough for predicting success in engineering.

We then provided all the NCEA results to the software for it to build a decision tree based on finer descriptions of the data. Figure 3 shows resulting tree. Interestingly, the first decision node in the tree uses the GES score, indicating that the system found the GES score to be the single most informative predictor of performance. This is perhaps encouraging, since this is the primary factor that the university currently uses for admission decisions. The decision tree then considers the total number of physics credits, the total number of Merits and Excellences in mathematics, the total number of mathematics NCEA credits, and the statistics results. Note that in these figures, mathematics with statistics is coded as "Stats", mathematics with calculus is coded as "Maths". "A" refers to the number of Achievement standards (excluding Merits and Excellences), "ME" to the sum of the Merit and Excellence standards.

In spite of having access to all the individual units, it still finds that the groupings by subject are the most informative, which provides some support for the correlation analysis in the previous section. It also supports the result from that section that physics and calculus performance at secondary school is important, and confirms the usefulness of grouping the Merit and Excellence grades (as used in the "Maths ME" decision node). However, as can be seen by the fractions of incorrect classifications, this still performs poorly as a predictor.



Figure 3. Classifier based on NCEA scores

To look for better predictors, we took advantage of the availability of "Not Achieved" results, (at least for VUW students) that were released to us midway through this year. Unfortunately, these were not available from the other universities due to New Zealand's privacy laws.

The best classifier based on all possible NCEA scores was complex and had a poor reliability. We then simplified the data, by including only the "Not Achieved" grades and the summed Merit and Excellence grades, and used the simple binary classification criterion. In the following figure "N/V/Y/Z" represents the different ways that a student can gain a "Not Achieved" grade (this includes failing the subject and being absent for the assessment).

The classifier over all students was still not reliable, but it did suggest that the "Not Achieved" grades may be a useful diagnostic for predicting students who have a low chance of success.

We therefore run the data mining system separately on the two groups of students (ECEN and NWEN/SWEN). Figure 4 shows the decision tree for ECEN students only under the ternary classification criterion. The tree is comparatively simple and reasonably accurate.



Figure 4. ECEN Classifier based on NCEA scores

The most informative feature is the number of physics credits that were failed. If none were failed, the student is predicted to easily pass. The next feature is a specific NCEA module, 90636 (the integration module of level 3 mathematics with calculus). This classifier is the first of our predictors to take an individual NCEA module rather than consider the subject as a whole. The predictor can

be understood as saying that passing all of the physics modules is good indicator of success in ECEN, but students who failed some physics but have passed the integral calculus may succeed, perhaps with additional support.

For the NWEN/SWEN students, the generated decision tree is much more complex, and in the authors' opinion has so many dependencies that in fact it is the noise in the data that is being modeled rather than an underlying pattern. It is presented in Figure 5 for the reader's interest, but the authors do not consider that any useful conclusions can be drawn from it.

The conclusion from the data mining analysis essentially confirmed the correlation analysis, with the addition of a particular mathematics module identified as a useful predictor for performance in ECEN.



Figure 5. NWEN/SWEN Classifier based on NCEA scores

6 CONCLUSION

This paper has presented the results of a series of tests to determine whether NCEA grade data can be used as a reliable predictor of performance in tertiary digital engineering study. To inform this study, data were obtained from many of the providers of a Bachelor of engineering in these defined digital engineering areas.

To first establish whether particular secondary school subjects could be used as a direct predictor, a series of correlation analyses were performed. The correlation scores by themselves were too low to form a reliable predictor, however, they did indicate that a differentiation is required between the hardware ECEN type students (electronics, computer systems, mechatronics) and the NWEN/SWEN type students (networking/software). ECEN students had a heavy reliance on secondary school ability in calculus and physics. This was much less evident for the NWEN/SWEN students, with certain analyses indicating that performance in English was more important.

An interesting finding, and a consequence of the unit standard versus achievement standard assessment of Computer Studies (and technology subjects in general), was a negative correlation with university first year GPA. We obviously would be reluctant in the extreme however, to reject students into a Software engineering course because they had completed a Computer Studies course at secondary school!

The strongest correlation for ECEN type students was the combined Merit and Excellence scores in calculus and physics, a trend that was often exhibited in the overall number of Merits and Excellences in science subjects. However, the strongest correlation at 0.72, is not sufficient to form a reliable predictor.

A more in-depth analysis of the data was undertaken using data mining software that also considered individual NCEA modules rather than the subject (or groupings of subjects) as a whole. As a first pass, the new VUW Guaranteed Entry Score of 150 was established as being a reasonable level, although a significant number of students were incorrectly classified indicating that an individual assessment of students is still required.

This data mining was able to utilize the "Not Achieved" NCEA data for VUW students, whereas the national correlation analyses did not have access to this. Again for ECEN students, the importance of calculus and physics at the Merit and Excellence level was established. Students who only manage an "Achieve" level in these subjects are very likely to struggle at the university level. At the least they must be encouraged to take the lower level mathematics and physics subjects and to make full use of the additional assistance provided.

Given the low prediction capabilities of NCEA, VUW developed a diagnostic test for ECEN students at VUW that have shown very encouraging results. Space unfortunately does not permit a discussion of this diagnostic, however, when used in conjunction with the NCEA data, a reasonably reliable predictor of failure in ECEN students can be established. The use of this modified predictor has been used to provide early assistive intervention and even to suggest enrolment in an alternative degree.

Whilst the focus of this study was primarily on the academic predictors of engineering success, these must be viewed in the context of social, cultural and expectational influencers. It is well known in the literature that students who should perform well at university sometimes do not, and can actually fail abysmally. This cannot be predicted from school results since it can often be the newly found independence at university that alters a student's behavior. Our pastoral care agent has helped us to identify and resolve such issues before they permanently affect the student's chances of succeeding in engineering. This is presented in another paper.

7 REFERENCES

[1] Carnegie, D.A., Watterson, C.,A., Andreae, P., and Browne, W.N., "Prediction of Success in engineering Study", In Proceedings of the IEEE engineering Education Conference (Educon 2012), Marrakesh, Morocco, 2012, pp 57 – 65. ISBN: 978-1-4673-1455-8

[2] Everett, J. E., & Robins, J., "Tertiary Entrance Predictors of First-Year University Performance, Australian Journal of Education", Vol.35(1), 1991, pp. 24-40.

[3] McClelland, A.A., & Kruger, P.W., An Investigation of the Subsequent Performance in Tertiary Studies of Students Admitted Through the Queensland Tertiary Admissions Centre in 1989-90, Australian Government Publishing Service, Canberra, 1993.

[4] Engler, R., Are particular school subjects associated with better performance at university?, Ministry of Education, Wellington, 2010.

Prototype of a New Engineering Masters Project Model: Working with Marketing and Software Faculties to Commercially Kickstart University Research

Jonathan Scott, and Benson Chang*
Caine Jameson, Lesley Payne, Daniel McLaren, and Anthony Barr-Smith†
Chelsea Came, and Cordula Hechei.‡

*School of Engineering, University of Waikato, Hamilton, New Zealand. e-mail: jonathanscott@ieee.org

†Computing and Mathematical Sciences, University of Waikato, Hamilton, New Zealand.

‡Management Communications, Waikato Management School, University of Waikato, Hamilton, New Zealand.

Abstract: We describe a Master of Engineering (500-level) project modelled on the real-world arrangement where engineers work with marketing and software groups to prepare a product for commercialisation. A 4-member software team to develop and test embedded firmware and support applications on a mobile platform was provided through a final-year undergraduate software-engineering project course based outside the engineering school, in a separate faculty. A marketing team consisting of interns prepared logos, product names, and advertising materials, with input from a creative 200-level class. This team also considered possible exit strategies based on analysis of the market size and activity. This marketing effort was organised through the management communications group in the magagement school. The masters student acts as project manager and it is their remit to guide the product towards release on the crowd-sourced venture-capital site `kickstarter.com`. A small but original product idea is required to provide a viable vehicle for the project. Financial commitment to manufacture, even on a small scale, represents a novel outcome for a university project.

Keywords: Continuing education, Electronics engineering education, Venture capital, Interdisciplinary collaboration, Project management, User interfaces, Digital command and control (DCC), Microcomputer applications

1 Introduction

Engineers are expected to have team skills. In modern product development groups they typically work with programmers and marketing personnel to develop a manufacturable prototype in parallel with documentation, marketing materials, and marketing strategies. In the "Silicon Valley" model, engineers are also stockholders and may be involved with business decisions such as IP protection approaches and potential corporate exit strategies. On the other hand, undergraduate programs typically emphasize technical learning, although modern accreditation programs demand an amount of teamwork and management skills. Much of the practical detail of product development is left to be picked up "on the job", possibly through industrial placements such as those integrated into the degree structure of Antipodean universities.

Our thesis is that we ought to be able to provide the real-world learning through a Masters program. The program model reported in this manuscript is a compact version of what really happens in our industry. We address the questions "How might a small version of what happens in industry be arranged at a university?" and "How can such a program be assessed?"

2 The Model

For the purposes of the experiment, the masters student, the second author of this manuscript, is regarded as the

project manager (PM). In this particular instance he is also the "chief hardware engineer". The supervisor has the responsibility of establishing liason with colleagues in computer science or software engineering, and with marketing or management communications, and generally providing "the environment".

To fit viably in the time available, the project must either have been proven feasible beforehand, say through the existence of a prototype that demonstrates the technology, or be of a level of difficulty such that the hardware could be constructed in about half of the available time. A masters program to corresponds to something between 1200 to 1800 hours. By way of comparison, universities typically consider a full-time academic employee to contribute 1500–1800 hours per year once allowance is made for leave, overheads, and inefficiences.[1] The provision of a vehicle project is not as serious a demand as it might appear at first. There were several candidate projects available to the authors, including various toys, gadgets, and small appliances. Websites such as `kickstarter.com` are a fertile source of inspiration, and variations on exiting themes provide more opportunities.

For the experiment reported here, we chose a small electronic toy using a microcontroller with embedded firmware that demonstrated some novel ideas. The history and technology of this particular project, and its suitability for this experiment, is covered in section 3.

3 The Technology

The Electronic Engineering problem of transmitting both power and control data over the same pair of wires has arisen in a variety of contexts, including domstic power delivery, household automation, and remote control of small tethered tools and toys. The model railway industry provided an elegant solution for this several decades ago. It was developed by Lenz in Germany for the Märklin and Arnold companies in the 1980s. [1] The system is today known as "DCC", standing for Digital Command and Control. It was accepted as the global standard, published in February 1994 by the NMRA, and has been adopted by all major model railway manufacturers. [2] After the fashion of TCP/IP it is documented as a series of standards and recommended practices, or RPs, available from [2].

The DCC standards make no specification about the user interface (UI). A major problem impeding the acceptance of DCC in the marketplace has been the overly complex and technical format of controllers. Figure 1 epitomises the problem. The advertisement in the figure depicts a complex-looking piece of equipment that at first glance could be any instrument from an engineeing lab, rather than a consumer toy. The caption proudly displayed—"If you can use a TV remote, then you can use the Elite"—tells the potential buyer that he or she is in for a frustrating, technical experience. The adoption of DCC has been very slow, in spite of its technical elegance and universal adherence to the standards that guarantee compatibility of various suppliers' products. Browsing a high-circulation trade magazine such as [5] serves to confirm that the technical experience promised by the product in figure 1 is the norm. Interviews, conducted in 2009 in collaboration with the psychology school, discovered that in clubs where DCC is widely deployed, a small fraction of the members understood the technology, and typically provided de-facto technical support to the others.

Over a period of about two years, staff and students at our institution have solved the problems of assigning and remembering addresses, in much the same way that USB solves that problem in comparison with bus connection of IEEE488-equipped instruments, but without an upstream data channel or cooperation of the adressees, as this is not generally available on DCC systems. [2] Also, a new UI has been designed that takes advantage of the scanning technology to eliminate the need for a keypad. The new controller compares with designs of the type depicted in figure 1 as an iPhone compares to phones of early 1990s vintage: It is vastly easier to use, but in the end does the same thing, save elegantly, in accordance with accepted design principles. [4]

The hardware constructed to verify feasibility of our ideas represented a prototype for a potential product. The university declined to patent these ideas, considering the small market addressed by DCC. A processor turn and relayout of the PCB on the original design and the development of a manufacturable form factor and enclosure seemed like a suitable hardware task, perhaps a few hundred hours work. In the end, we decided to make provision for Bluetooth connectivity as well via a daughterboard. Our design, employing the new technologies that we have dubbed "iDCC", seemed like an ideal candidate for this experiment. Figure 2 shows the mock-up enclosure prepared early in the project and figure 3 depicts the view when train reprogramming is required. If two locomotives respond to a single knob, one requires reprogramming. The enclosure is inverted and the locomotive placed on a short track attached to what is normally the bottom of the controller for a few seconds while the locomotive is automatically reprogrammed.

4 The Market Space

The model train industry is a lucrative niche market worldwide. In New Zealand alone the market is estimated to consists of approximately 9000 consumers, based on exhibition attendance. Research gathered previous to this marketing teams involvement indicates that the market is vast, and ready for a user-friendly controller such as this. Ad-

[1] This count is for "billable" hours, used in calculating consultation rates. Statistics suggest the average engineering academic puts in rather more hours.

[4] Donald A. Norman, *The Design of Everyday Things*, 2002, ISBN 978-0465067107, originally published as *The Psychology of Everyday Things*, Basic books, 1988.

[5] *Model Rail*, Bauer Consumer Magazines, October 2010.

[6] Kickstarter web site, "A funding program for creative projects", `http://www.kickstarter.com/`, retrieved August 2012.

[7] Wortham, Jenna, "A Few Dollars at a Time, Patrons Support Artists on the Web", The New York Times, August 24, 2009.

[8] Chris Thomson & Ben Ryan, *Genie—Motion control time lapse device*, `http://www.kickstarter.com/projects/syrp/genie-motion-control-time-lapse-device` retrieved June 2011. See `http://syrp.co.nz` for the recipient company.

Encouraging Success – measure, mentor and engage

Gerard B. Rowe[1], Martin Shepherd[2] and Chris R. Smaill[1]
[1]Department of Electrical and Computer Engineering and [2]Faculty of Engineering
The University of Auckland, New Zealand
{gb.rowe, ma.shepherd and c.smaill}@auckland.ac.nz

Abstract: Five interventions are described that have the collective aim of improving preparedness for tertiary level engineering studies, the transition from high school to university, and subsequent academic success. A combination of a pre-admission test and first-year support via peer mentoring assist with shortcomings in prior preparation and facilitate the transition to university. Engagement and curriculum mastery are assisted by a combination of peer marking, supplemental instruction and teaching-assistant induction.

Keywords: Pre-admission quiz, peer mentoring, peer marking, supplemental instruction, teaching -assistant accreditation

1 INTRODUCTION

Recruiting suitably prepared students, managing their transition into tertiary study and ensuring their retention throughout their programmes are complex issues with which many institutions are currently grappling. This paper provides an overview of five initiatives currently employed at the University of Auckland to address these issues. Specifically, prior to beginning their studies all students are expected to complete a Ready for First-Year Quiz, the results of which inform both the lecturers and the students themselves about the student's prior preparation. A supportive peer-mentoring system (the Part-One Assistance Centre) is available to assist part-one students throughout their first year[1]. Engagement with course material (and attendance at tutorials) is encouraged by the use of peer-marking exercises within tutorials. The transition from Part One to Part Two is further enhanced by supplemental instruction (delivered to academically at-risk students) in special "Foundation Tutorials" in gate-keeper part two courses. The student laboratory and tutorial experience (throughout their degree) is enhanced by a formal training (and accreditation system) for teaching assistants.

2 DISCUSSION

A brief background is provided for each of these five initiatives, including details of the methodology and links to relevant references. (These contain more extensive quantitative evaluation of the success of the various interventions, including course results over several years, and feedback from questionnaires and student focus groups).

[1] Maori and Pacific students through the Tuakana Tutorial and Mentoring Programmes and to women students via women-only tutorials at Part 1 plus further on-request individual or group tutoring (as appropriate) via the Women in Engineering Network.

2.1 Ready for First Year Quiz

In 2009 The University of Queensland (UQ) introduced (Kavanagh et al (2009)) an on-line competency test (PFFY – Preparing for First Year) for all new first-year engineering (FYE) students. This PFFY test was designed to help address concerns about a perceived lack of prerequisite knowledge and growth in attrition exacerbated by a lower entrance standard due to an increased intake quota. Despite a more stable admissions climate, the ability to select students from the highest-achieving cohort of school leavers entering university, and high levels of retention among FYE, The University of Auckland (UofA) – based upon the experience of UQ – identified potential benefits from introducing its own 'Ready for First Year Quiz' (RFYQ) in 2011. This RFYQ was attempted by the students following their acceptance of their admission offer and prior to Orientation week. After two cycles, the RFYQ has proved to be a vehicle for raising students' awareness of their individual areas of weakness and/or lack of knowledge, as well as providing explicit references to which courses will assume or require this knowledge, and providing resources for review as relevant [Shepherd et al, 2011].

2.2 Part-One Assistance Centre

In an attempt to support the transition and retention of first-year students in its Bachelor of Engineering (Honours) degree, the Faculty of Engineering has for the last 5-6 years provided supplemental academic support to first year students through the coordination of a 'drop-in' peer mentoring service (known as 'the Part-One Assistance Centre'). This initiative draws from two long-held assumptions. First, that the transition from high school to the academic rigours of university study can be a difficult process for some new students, compounded by an array of social, cultural and emotional pressures felt by undergraduates in unfamiliar surroundings (James, Krause and Jennings, 2010). Entry-level '100' courses tend to be large and potentially impersonal learning environments in marked contrast to many students' recent

high-school experiences (Exeter et al., 2010) and many students consider dropping out in the first few weeks and months (Krause, 2005). The second assumption follows on from the first; in this new learning space undergraduate students find solace in their peers and often form study groups and many universities have sought to harness this with peer-tutoring or mentoring schemes with senior students taking on the role of teachers in undergraduate courses (Beasley, 1997).

The Part-One Assistance Centre attempts to help first-year Engineering students make a successful academic transition to university study by drawing from the 'insider' knowledge of part-two and -three engineering students who are only one-to-two years further advanced in their studies. In doing so new engineering students have a simple way of engaging with university teachers in a low-key manner. This engagement should not be underestimated: a peer mentor may be one of the few individuals involved in teaching that a new student will get to know and could be an important factor in helping new students stay on course (Russell and Slater, 2011).

The service provided bears many similarities to other peer-mentoring initiatives in existence throughout Australia, New Zealand and the wider academic world: mentors are selected from high-achieving students in their second or third year in the discipline who have achieved high grades in the courses for which they are providing assistance; the service is available to all students and promoted at the beginning of large first-year lectures and in first-year course outlines, and so on. The mentors are paid a nominal hourly rate for their efforts but appear to be motivated as much by their own service-ethic as they are by immediate financial considerations. (Anecdotally, mentors have at times indicated they would earn more through private tutoring but acknowledge other benefits as motivating them to apply for and accept the position.) The Faculty routinely receives 130-150 applications for 20-30 available positions from those invited to apply each semester.

2.3 Peer Marking

The first year of the four-year Bachelor of Engineering (Honours) program at the UofA is taught entirely in-house by the Faculty of Engineering. The first-year engineering course ELECTENG 101, Electrical and Digital Systems, is compulsory for all 700-plus first-year students. Typically two thirds of these students will subsequently choose a specialization other than electrical and computer engineering at the end of year one. Clearly, many of the students enrolled in ELECTENG 101 may have no particular interest in matters of an electrical nature, and may even have chosen not to study this aspect of physics at high school.

ELECTENG 101 involves students in three lectures and one tutorial each week of the twelve-week semester. As there are two lecture streams, a typical lecture is given to about 350 students. Two lecturers give the lectures, with each lecturing both streams for about half the semester. Tutorials are smaller, typically of size 40, and involve several academics. Each tutorial is managed by one academic and one teaching assistant, normally a postgraduate student. In addition to tutorials, informal drop-in clinics (the Part-One Assistance Centre discussed in Section 2.2) are also provided. Laboratory exercises provide further learning experiences, as do various online resources. While completion of all laboratory exercises is required, there is no attendance requirement for lectures or tutorials. Online assignments provide early feedback to students about their progress, and this feedback is augmented by that from two written tests.

It is difficult to promote active learning in the large lectures, although this is attempted. The small size of tutorials makes them more amenable to active learning. However, tutorial attendance almost always declines through the semester, and this effect has been more pronounced in recent years. For example, in 2007 and 2008 some tutorial streams were, in the latter part of the semester, being attended by only 10% to 20% of the enrolled students: most students did not make the most of the educational opportunities provided. This situation worried the course's lecturers, who were aware that many students perceived the course as difficult and as somewhat of a 'gate-keeper'. Without active involvement in tutorials, too many students were failing to meet the required standard. It was considered that something needed to be done to boost significantly attendance at and involvement in tutorials. The action chosen was the implementation, in 2009, of peer-marking exercises in tutorials (Smaill et al 2010, 2011, Carter et al, 2011).

2.4 Supplemental Instruction

As developers of course-concept inventories and researchers of threshold concepts will attest, a pass in a particular first-year course, or indeed the entire first year of study, is no guarantee that a student has mastered all of the relevant technical material or achieved a sufficient level of academic independence. On the contrary, the student may still have significant conceptual misunderstandings or learning difficulties. If these remain undetected and uncorrected, the future academic success of that student is threatened.

For some years, considerable engineering-education research effort has been directed internationally toward improving student success in first-year programmes by evaluating strategies which have ranged from curriculum redesign to interventions aimed at increasing academic and social integration. Investigations of students' academic preparedness, student study practices and engagement with their own learning have been useful in identifying students at risk of failure. The provision of supplemental instruction for at-risk students has been found to be effective in improving the grades of such students. While this remains a very active research area, similar attention is now starting to be focussed on the second and subsequent years when students encounter core discipline-specific courses.

An intervention aimed at improving student success in three key second-year electrical-engineering courses covering, respectively, the topics of circuit theory, electronics and electromagnetics is reported in Rowe et al (2009, 2010). These courses can be considered to contain the theoretical underpinnings for a degree program in

Electrical Engineering, and at this institution they have proved for many years to be "gate-keeper" courses in terms of progression through the degree. For the last two years, at-risk students have been identified, principally, by a combination of diagnostic testing and early course-work results. These students have then been invited to join a special, weekly, "Foundation Tutorial" in which both technical and learning obstacles are addressed. Two years of data collection suggest that such an intervention does produce a significant improvement in the success of participating students.

2.5 Teaching-Assistant Accreditation

When the Faculty was last accredited by IPENZ in 2010, the accreditation exercise highlighted that the Faculty had no transparent, auditable process to demonstrate that students employed in teaching-support roles – whether TAs, lab demonstrators or markers - had been trained and were consequently 'fit for purpose'. There was also a sense that the level of training offered could be very task-specific and varied from one Department to another (and possibly from one course to another within the same Department). This was subsequently confirmed by an initial fact-finding exercise conducted in March 2011.

As a result of this exercise a working group was established to develop and pilot a formal accreditation programme for the professional development of all students employed in teaching-support roles. A structure was devised for a single full-day programme that would incorporate professional development sessions on tutoring, teaching in labs, and marking and formative assessment. The sessions would be generic and so equally relevant and applicable for all disciplines. The objectives for the programme would be to provide advice and guidance that would be relevant for the actual tasks relating to each role, but also to link the relevance of the role to the wider context of the Faculty's teaching objectives and provision.

The programme was to be seen as a complement to, rather than an alternative for, existing Department or in-course training of TAs, lab demonstrators and markers. All students employed as TAs, lab demonstrators and markers were to be invited (and were expected) to complete the programme – not just those without any prior experience of a teaching-support role. To reinforce the importance of the programme and the value of participation in it, the programme would commence with a formal welcome from the Dean followed by a 'Teaching Excellence' panel session comprising members of the Faculty's academic staff who had been recognised for their teaching expertise. Participants would also receive a formal certificate and citation on completion of the programme.

3 RESULTS

A detailed quantitative evaluation of the individual success of the various initiatives is provided in the references associated with each initiative. Brief details are provided in the following sections, with the aim of this paper being to provide a ranking of the effectiveness of these initiatives for two situations – firstly discounting cost and secondly considering a rough cost/benefit analysis.

3.1 Ready for First Year Quiz

Given that UofA's main motivation for adopting the RFYQ was to raise FYE students' awareness as to their individual weaknesses and encourage them to seek redress prior to the start of semester, the initiative has been relatively successful, mirroring the experience at UQ.

Increased enrolment was observed in pre-semester calculus and physics preparation courses: 'Superstart Mathematics' noted an increase from 90 students in 2010 to 124 in 2011 (a 37% increase), and 'Launch Into Physics' experienced a 90% increase from 20 enrolments in 2010 to 38 in 2011. In addition, a survey of FYE students (N=438 responses from a cohort of 620) conducted at the end of first semester indicated that 34% agreed or strongly agreed that the quiz had been helpful in preparing them for their first year studies, 20% disagreed or strongly disagreed, 40% were 'neutral', and 6% answered 'N/A'. Whatever the student response, the fact remains that simply by participating in the tests, students would have reviewed knowledge and thereby been better prepared for their first semester. It is understood that as the tests are performed in the students' own time and space, there is scope for collusion and using textbooks or the web to answer questions. However, the tests are not for assessment and, it is argued, in the process of finding the correct answer, the students are effectively reviewing knowledge.

While the survey indicated that only 15% agreed or strongly agreed that they were motivated to access resources and review knowledge as a consequence of the quiz (53% disagreed or strongly disagreed and 26% were neutral), the high participation rate (88%) was seen as a positive outcome in raising awareness of what students could expect in their FYE courses.

Of those that were inspired to follow up the RFYQ with pre-semester preparation, the most common response (16%) was to review school notes, past exam papers and text-books, with others enrolling in University preparation programs, purchasing and reviewing course text books early, and/or accessing course materials on-line. Responses also indicated that students accessed material on subjects for which no prior learning was a pre-requisite – e.g. electricity, biology and chemistry – which appears to justify the motivation to include unfamiliar discipline problems in the RFYQ. However, it should be noted that students were not aware that this knowledge was not assumed. Unlike UQ, UofA is unable to ascertain whether those that prepared in advance for first year performed better than those that did not. Nor is there any obvious impact upon FYE student performance: some Semester 1 courses recorded slightly higher grade distributions whilst others were slightly lower.

Review of the inaugural implementation of the RFYQ initiative indicated that it was important to understand the FYE cohort's academic background, and to gather student feedback on the usefulness of the quiz in preparing them for first-year studies. The PFFY achieves this

very effectively by incorporating additional questions, and then following up with end-of-course surveys, and ad hoc focus groups or polls taken during orientation and early in the semester. Interestingly, both UofA and UQ noted that many students appeared to have forgotten much about the quiz by the end of first semester and the discipline areas it covered. Thus, there is a need to generate and record feedback at a number of intervals following completion of the tests to get a true understanding of how the tests can be best employed, to better understand student motivations, and improve the initiative.

3.2 Part-One Assistance Centre

In a survey conducted at the end of their first semester for one of the Faculty's required first-year courses (ENGGEN 140), students responded accordingly to the following question (Table 1):

> *Q. I have found the Part-One Assistance Centre to be positive in helping me to work through any problems I encountered with my Part I courses*

Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree	N/A
5	19	136	320	51	90

Table 1: ***Student response to ENGGEN 140 survey on effectiveness of first-year support initiatives, including the Part-One Assistance Centre***

Part-two engineering students have great potential to induct first-year students into the culture of the Faculty and to give practical advice on specific topics the new engineering students are grappling with. However, developing an effective peer mentoring scheme that meets the needs of the key 'stakeholders' (students, mentors and academic teaching staff) requires thoughtful 'behind-the-scenes' management. For example, while always actively supported and recognised as an excellent resource by part-one lecturers, there was some concern (arising from a lack of knowledge) about the type and level of academic assistance provided by the mentors as well as a lack of awareness about the mentors' general experience and effectiveness in tutoring Part 1. Originally the mentors were given minimal training prior to their assignment with no performance monitoring throughout the remainder of the semester. Nor was there any formal guidance on just how much assistance the mentors should provide on actual assignments for each course. The mentors were generally left to their own devices with only fellow mentors and the Part-One Assistance Centre Coordinator as possible channels of support and/or feedback. It would also have been easy for the mentors to question the value of their service given the lack of interaction with first-year teaching staff.

In 2010, the Faculty moved to remedy this sense of 'disconnection' and so enhance the quality of service and general support provided at Part One by adopting:

- the introduction of formal training sessions, under the supervision of The University's Centre for Academic Development, prior to and halfway through each semester, which directly engage each new cohort of part-one mentors with part-one course coordinators;
- the creation of an on-line forum in the university's learning management system (CECIL) in which mentors provide daily reports on their interactions with students, receive advice directly from part-one lecturers (and other mentors) on how best to advise first-year students, and share and access resources that help inform the assistance the mentors provide; and,
- a formal end-of-year ceremony, presided over by the Dean and marked with the conferment of certificates and citations of service to acknowledge the value and importance of the part-one mentors to the transition and academic success of first year students.

The training (Shepherd and Brailsford, 2012) has helped to ensure that the students are equipped with more than just their individual subject and discipline knowledge to help new first-year students, having the skill, patience and ability to listen, communicate and guide students who may not have a similar grasp of the course material as they have. The sessions focused heavily on learning strategies and discussions of the likely problems they would encounter mentoring students and the techniques they could employ in dealing with them. The mentors have embraced the training and regularly report how it has benefited their own learning and academic performance, raising their awareness of their own learning styles and alternatives they can employ when 'tried and trusted' methods fail them. While the peer-mentoring scheme was initially developed to help 'at-risk' students, it is becoming clear from feedback that the mentors are gaining hugely from the experience of playing the role of the teacher.

The creation of the on-line forum has been the most important and valuable innovation. The daily reports immediately gave reassurance to teaching staff about the quality and level of support being provided. Teaching staff were extremely impressed with suggested ways of approaching problems that they had not considered previously. More importantly, the daily reports have provided extremely valuable feedback on the academic progression of first year students – indicating whether students are at the right point of intellectual development, whether students are getting 'tripped up' or 'bogged down' in unexpected areas or by concepts which were assumed understood. For the mentors, the online forum led to regular positive reinforcement of their assistance from teaching staff, but also a much wider community of support with fellow mentors sharing their own ideas on how best to answer particular student concerns. As a result there has been a much greater sense of purpose among the mentors. They can see their contributions are valued, and they also know that any thorny problem they encounter is likely to have already been covered and resolved in an earlier forum.

3.3 Peer Marking

In an effort to improve student achievement via boosting tutorial attendance and engagement, peer-marked as-

signments were set in a compulsory first-year engineering course. The peer marking took place under the guidance of a staff tutor. Students received a small number of marks for doing an assignment and marking another's, thus creating an environment in which students were required to interact with the tutorial material both from the perspective of one solving the problem, and also from the perspective of a marker.

Four years of experience with this initiative have produced very encouraging results. Attendance at tutorial sessions improved very significantly during peer-marking sessions. Tutors and teaching assistants noted that peer marking appeared to help student engagement and learning. Survey results indicated that the great majority of students considered peer marking had made them think more deeply about the material and also about how a solution is communicated. Examination performance improved over earlier years, and results from a follow-up diagnostic test provided evidence that circuit-theory concepts were retained better into the following year. In consequence, peer marking is to be retained as a standard component in ELECTENG 101 and is now becoming a standard feature of other first-year engineering courses at the University of Auckland.

3.4 Supplemental Instruction

Supplemental instruction (in the form of Foundation tutorials) has been run in three courses across four semesters, namely: *Circuits & Systems* and *Electromagnetics* (for which two years' data is available) and *Electronics*. In each course, at-risk students were identified by a combination of diagnostic testing and analysis of early course-work results. Typically a (not-for-credit) diagnostic test would be administered without warning in the second lecture of the course. Such a test would usually be based on key concepts, and use the same principles as underpin most concept inventories. Students with very low marks (thus identified as being at-risk of failure) were then contacted by email, advised of their "at-risk" status and invited to attend an extra weekly tutorial to help improve their chances of successful course completion. Additional evidence from early course-work results was used to instigate a second round of invitations.

As a first attempt to assess the value of the tutorials, the at-risk students were divided into two groups. Students who attended approximately 50% or more of the available tutorials were classified as regular attendees, with the remaining students classified as non-regular attendees. By comparing the percentages of both groups who ultimately passed the various courses (Rowe et al 2010) a much higher pass rate was observed amongst regular attendees. Specifically, there is evidence that at-risk students who attend all (or most) of the Foundation tutorials tend to pass the course. On the other hand, at-risk students who attend no Foundation tutorials have a higher probability of failing the course. This data therefore provides support for the conjecture that attendance at the tutorials significantly improves the chances of a student passing the course. A possible confounding factor, which has not been eliminated at this stage, is that the regular attendees may simply be better students overall. In the next phase of this research project it is planned to perform a t-test on each group's grades in previous courses to see if the means are statistically different.

The research was further informed by interviews with student attendees and academic tutors. The student interviews probed three issues, namely: why students initially seen by the university as academically able were being identified as at-risk just one year later, whether the Foundation tutorials enhanced learning, and what recommendations the students had for improving their learning.

All but one of the interviewees confirmed that Electrical and Computer Engineering was their chosen specialty i.e. they hadn't been forced into this discipline because their chosen discipline was over-quota. Not surprisingly, a wide range of issues were identified as contributing to their at-risk status. Some could be considered related to their ability to take responsibility for their own learning, such as time management, insufficient practice, poor mathematics and physics background, poor motivation, shallow learning, and a mismatch between their learning style and the predominant style of course delivery.

The items of interest that arose from review of the transcripts of the interviews conducted with the tutors were their perception of the students' ability, the appropriateness of the classroom activities and learning spaces, their perception of the degree of success of the intervention, and their opinions of the usefulness of the tutor training and course-coordinator briefings they had received. The tutors, all academic high-achievers, were surprised at the low technical level of the misunderstandings hindering the students and at the extremely passive approach adopted by many of the students.

3.5 Teaching Assistant Accreditation

The students were distributed into three groups. The students remained in the same group throughout the day and attended the tutoring, teaching in labs and marking sessions in sequence. The sessions were interactive and generated a real buzz as students continued to interact as they moved between sessions and came together for lunch or the general plenary. The students were obliged to complete self-reflection forms for each of three professional-development sessions and a general plenary (in which the three groups all came together) to review the main lessons learnt in advance of their roles.

The students who attended engaged well with the programme, providing positive feedback on the ideas, tips and concepts they took away from each session as well as praising the effectiveness of individual sessions. Several experienced TAs also commented that they had wished the programme was available when they had first been appointed. A review of the semester-one event concluded that the TA Accreditation Programme had worked well and offered a sustainable model, albeit with some minor modifications.

Two changes to the existing programme were implemented (for the second semester) following the review. Firstly, an additional panel was included immediately preceding the plenary session – comprising acknowl-

edged 'star' tutors to share their experience and provide tips to new recruits. Secondly, the final plenary session was structured around the use of a learning styles instrument (VARK) in which all members of the audience participated.

In order for the programme to have maximum benefit for the Faculty, it would need to be offered twice per year – prior to the start of each semester – and rolled out to all Departments. Given that attendance is unpaid and non-compulsory, it is important that the Dean and HoDs give full support to the promotion of the event – not only among students but (perhaps more importantly) among colleagues to ensure the importance and value of participation is strongly and consistently recognised.

4 OVERVIEW

In terms of effectiveness in lifting student performance and not worrying about cost greatly, we see the Part 1 Assistance Centre and Peer Marking as being well ahead of the rest, which we would rate as roughly equal. On a bang-for-your-buck basis, we see Peer Marking as being out in front. We think it makes a big difference for very little extra investment. Then, roughly equal, would come RFYQ and the TA Accreditation Programme. These also produced quite good rewards for little extra cost to implement. By contrast, we see the Part 1 Assistance Centre and Supplementary Instruction as very good initiatives, but comparatively quite costly to implement. The cost (essentially teaching assistant (TA) salary costs) is roughly fixed for a given approach, and the benefit depends on how many students actually avail themselves of the service. (All approaches require extra investment of lecturer time to support the initiative and train the TAs.)

5 CONCLUSIONS

Five initiatives designed to address issues associated with recruiting of suitably-prepared students, managing of their transition into tertiary study, and ensuring their retention throughout their programmes are described. Results from several cycles of implementation give confidence that these initiatives produce sufficiently good outcomes to warrant the investment required.

6 REFERENCES

Beasley, C. (1997). Students as teachers: The benefits of peer tutoring. In Pospisil, R. and Willcoxson, L. (Eds), *Learning Through Teaching*, p21-30. Proceedings of the 6th Annual Teaching Learning Forum, Murdoch University, February 1997. Perth: Murdoch University.

Carter, L.J., Abdulla, W., Rowe, G.B. and Smaill, C.R. "Large-class peer-marked assignments for improving second-year student performance in Electrical and Computer Engineering", *Proceedings of the 22nd Annual AAEE Conference, Fremantle, Australia, December, 2011.*

Exeter, D., Ameratunga, S., Ratima, M., Morton, S., Dickson, M., Hsu, D. and Jackson, R. (2010). Student engagement in very large classes: the teachers' perspective. *Studies in Higher Education, 35*(7), 761-775.

James, R., Krause, K., and Jennings, C. (2010) *The first year experience in Australian universities: Finding from 1994 to 2009*. Melbourne: Centre for Studies in Higher Education, The University of Melbourne.

Kavanagh, L., O'Moore, L., Samuelowicz, K. (2009). Characterising the first year cohort knowledge. In K.E. Sparks & M. Simonson (Eds.), *Proceedings of the 20th Annual AAEE Conference, Adelaide, Australia, December, 2009.*

Krause, K. (2005). Serious thoughts about dropping out in first year: trends, patterns and implications for higher education. *Studies in Learning, Evaluation, Innovation and Development, 2*(3), 55-68.

Rowe, G.B., Smaill, C.R., Godfrey, J.E., Carter, L.J., Guillemin, B.J., Andrews, M., Abdulla, W.H. 'Dealing with the Tail: Remedial Tutorials for Second-year Electrical-engineering Students', *Proceedings of the 20th Annual AAEE Conference, Adelaide, Australia, 6-9 December, 2009.*

Rowe, G.B., Smaill, C.R., Godfrey, J.E., Carter, L.J., Guillemin, B.G., Andrews, M. and Abdulla, W. "Supplemental instruction: Foundation Tutorials for second-year electrical-engineering students", *Proceedings of the Annual Conference of the American Society for Engineering Education, Louisville, Kentucky, USA, 20-23 June 2010.*

Russell, B., and Slater, G. (2011). Factors that encourage student engagement: insights from a case study of 'first time' students in a New Zealand university. *Journal of University Teaching and Learning Practice, 8*(1). Retrieved March 7, 2012, from http://ro.uow.edu.au/jutlp/vol8/iss1/7

Shepherd, M., Kavanagh, L., McLennan, A., O'Moore, L. "Ready for first year? The use of pre-teaching diagnostic tests to prompt greater preparation and engagement among first year engineering cohorts at The Universities of Auckland and Queensland, *Proceedings of the 22nd Annual AAEE Conference, Fremantle, Australia, December, 2011.*

Shepherd, Martin and Brailsford, Ian, "Tightening the Safety Net - Assuring the Quality of the Peer Mentoring Experience for First Year Teaching Staff, Students and Peer Mentors in the Faculty of Engineering", *The 15th International First Year in Higher Education Conference, Brisbane, 26-29 June, 2012.*

Smaill, C.R., Rowe, G.B. and Carter, L.J. "Improving student learning through peer marking in a first-year engineering course", *Proceedings of the 21st Annual AAEE Conference, Sydney, Australia, 5-8 December, 2010.*

Smaill, C.R., Rowe, G.B. and Carter, L.J. "Peer marking – Does it really improve student learning?" *Proceedings of the Annual Conference of the American Society for Engineering Education, Vancouver, Canada, June 2011.*

Teaching introductory embedded systems with the UCFK4

Dr M. P. Hayes
*Department of Electrical and Computer Engineering, University of Canterbury
Private Bag, Christchurch, NEW ZEALAND.
Email: michael.hayes@canterbury.ac.nz

Abstract: A development system for introductory embedded system teaching based on the ATmega32u2 microcontroller is presented. The system is intended for simple games and includes a dot-matrix display, five-way navigation switch, and infra-red communications. It can be programmed using USB and open-source tools. The hardware and software design decisions are discussed.

Keywords: Embedded systems teaching, microcontrollers, C programming

1 INTRODUCTION

Teaching introductory embedded systems is not easy; students need an understanding of many disciplines, including computer architecture, systems programming, concurrency, control, and electronics. More importantly they need to be able to integrate aspects of these disciplines. Furthermore, the knowledge-core of the subject is immature and technology has been rapidly changing. Unlike most academic subjects with a disciplinary identity, embedded systems has a thematic identity and Grimheden suggests that a project-oriented approach is preferable to teach the subject [1].

Over the years many platforms have been developed for teaching embedded systems, but these have become rapidly obsolete due to advances in microprocessor technology. This paper discusses yet another platform, developed for teaching introductory embedded systems at the University of Canterbury. The challenge was to design a device that was simple to comprehend and yet stimulating for the students. We wanted a low-cost device that the students could keep, show their family and friends, and use for other projects. This requirement ruled out many of the approaches used at other universities, including commercial development systems [2–8], modular systems [9, 10], router based systems [11], and other custom systems [12].

The Atmel ATmega8 microcontroller has been used for teaching introductory[1] embedded systems at the University of Canterbury for ten years, in the form of kitsets we supplied to the students. The ATmega8 was chosen since it was well supported by GCC and could be programmed by bit-bashing a serial interface using a standard PC parallel port without any specialised hardware. With the demise of the PC parallel port, an interim solution was found using a bit-bashed USB interface[2]. This was a clever hack but violated the operating specifications for the ATmega8. Although, it worked well, we needed to preprogram the microcontrollers with a bootloader.

The impetus for a new teaching platform came with a change in the course structure and the offering of the course to Computer Science students, most of who had no soldering experience. The key requirements were:

- Have simple hardware but support multiuser games
- Use a microcontroller with a well supported open source toolchain
- Use a microcontroller with simple peripherals
- Be powered through USB (or battery when standalone)
- Be programmable via USB
- Be cheap

In addition, a simple pushbutton and LED were deemed mandatory to simplify teaching of port I/O. We had also found that a LED dot-matrix display gave the students insight into timing, multiplexing, and simple bit-mapped graphics. It also allowed them to write simple graphical games which we found motivated them during assignments. To facilitate better user interaction, we decided to incorporate a navigation switch and infra-red communications.

[1] ARM microcontrollers are used for more advanced courses.

[2] `www.obdev.at/products/vusb/bootloadhid.html`



Figure 1: UCFK4 simplified schematic.

In this paper, I start by discussing the hardware and software developed for the new platform (the UCFK4) in sections Section 2 and Section 3, respectively. I then discuss some of the lessons learnt in Section 4 and draw my conclusions in Section 5.

2 Hardware

The UCFK4 is based around an Atmel ATmega32u2 microcontroller and includes a reset push button, a general-purpose pushbutton, a five-way navigation switch, a green power LED, a blue LED, a seven by five dot-matrix display, an infra-red LED, a 36 kHz infra-red receiver, and a USB connector. The USB connector provides 5 V to run the microcontroller and to allow programs to be up-loaded. A simplified schematic of the UCFK4 is shown in Figure 1 and a photo of an assembled UCFK4 is shown in Figure 2.

2.1 *Microcontroller*

The UCFK4 is based around an ATmega32u2 microcontroller. This was chosen since the AVR architecture is well supported by GCC, it has simple peripherals, and it has a USB interface. From a teaching perspective, its main disadvantages are that it does not have an internal analogue to digital converter and that the three timer-counters are inconsistent. While the ATmega32u2 has 32 kB of flash memory (this is more than enough for student assignments[3]) it has only 1024 bytes of static RAM. Nevertheless, this was found only to be a limitation when students tried to allocate unnecessarily large arrays. The cheaper ATmega8u2 and ATmega16u2 have only 512 bytes of RAM and were considered likely to frustrate the students.

[3]To put this in context, a sophisticated space invaders style program for the UCFK4 with tunes and sound effects takes about 8 kB.



Figure 2: The UCFK4 with the ATmega32u2 microcontroller (centre), the navigation switch (top-right), the dot-matrix display (centre right), infra-red receiver (bottom right), infra-red LED (bottom centre), and mini USB-B connector (top left).

2.2 *LED displays*

The LED matrix display has five columns and seven rows. Each of the columns is driven by a MOSFET to supply sufficient current if multiple rows are simultaneously displayed. Turning on a specific LED in the matrix requires driving the corresponding column and row low. In addition, there are two additional LEDS: a green LED for lab instructor sanity, illuminated when power is present; and an active high blue LED driven from an I/O pin.

2.3 *Switches*

The five-way navigation switch shares I/O pins with the LED matrix columns. To read the switch requires the LED matrix to be momentarily disabled while the shared I/O ports are reconfigured as inputs with pullups. In addition, a simple active-high pushbutton was added since it is conceptually easier for the students to use for their introductory labs. A current limiting resistor was included in case the I/O port is mistakenly programmed as a low output.

2.4 *Infra-red communications*

The UCFK4 has an infra-red LED and an infra-red receiver. This allows wireless communication and also

allows the UCFK4 to be controlled from an infra-red remote-control.

The infra-red LED is wired between the UART transmit data (TXD) pin and a timer output pin (OC0B). Thus the timer can be programmed to automatically generate the desired modulation frequency. A modulated infra-red packet can be generated simply by writing to the UART.

A TSOP6136TR infra-red receiver detects infra-red light modulated at 36 kHz. This frequency was selected since it is a common frequency for TV infra-red remote-controls. The infra-red receiver is interfaced to the UART receive data (RXD) pin. This simplifies wireless communication between two UCFK4 devices; the students do not have to poll (or use interrupts) to decode the received packet. The use of the UART also introduces the students to simple serial interfaces.

2.5 *External interfacing*

The mini USB-B connector provides 5 V to power the board and communication with the AVR DFU bootloader for programming of the ATmega32u2 microcontroller. Thus no additional hardware is required for programming, such as a USB to JTAG device [13].

The UCFK4 expansion connector allows other external devices, such as piezoelectric tweeters, to be interfaced with the ATmega32u2 microcontroller. This is useful for other UCFK4 projects. In addition, since the LED matrix is not a surface mount component it is inserted in a surface mount socket. Thus, it can be removed to provide additional interfacing for other UCFK4 projects.

2.6 *PCB assembly*

The earlier versions of the UCFK were distributed to the students as kitsets that they had to construct. However, since the ATmega32u2 only comes in a 32 pin TQFP package, it was decided to have the UCFK4 commercially assembled. For a limited production run of 250 units, the unit cost was approximately NZ $45 (the price was dominated by the assembly costs and the dot-matrix display). Although less than half the price of a text book, this makes the device more expensive than other more sophisticated devices, such as the Raspberry Pi[4].

3 SOFTWARE

A goal of the course is to teach the students to write modular software and to show them how abstractions can hide unnecessary gory details and make their software reusable.

As teachers we expect students to write good code but most of them are beginner programmers (the ENCE260 students have only had one term of tuition in the C programming language and a course on introductory programming in Python). To give the students ideas and show them a consistent programming style, many demo application programs using the UCFK4 were provided to them. Several versions of each program were provided, using increasing levels of hardware abstraction.

[4] www.raspberrypi.org

```
#include "pio.h"

#define LED_PIO PIO_DEFINE (PORT_B, 5)
#define BUTTON_PIO PIO_DEFINE (PORT_D, 7)

void main (void)
{
    system_init ();
    pio_config_set (LED_PIO, PIO_OUTPUT_LOW);
    pio_config_set (BUTTON_PIO, PIO_INPUT);

    while (1)
    {
        if (pio_input_get (BUTTON_PIO))
        {
            pio_output_high (LED_PIO);
        }
        else
        {
            pio_output_low (LED_PIO);
        }
    }
}
```

Listing 1: A simple demonstration of PIO abstraction.

3.1 *Device drivers*

Two levels of device drivers were provided to the students: simple hardware abstractions for the ATmega32u2 I/O ports, timers, and EEPROM; and generic device drivers for the external hardware of the UCFK4. The latter included drivers for the LED, pushbutton, navigation switch, LED matrix display, and infra-red communications. The generic drivers were written to be independent of the ATmega32u2.

The key hardware abstraction is for port I/O (PIO) since this minimises the microcontroller dependence and simplifies retargetting of the device drivers to other microcontrollers. A simple demonstration of C code using this abstraction is shown in Listing 1. The macro `PIO_DEFINE` enumerates a port I/O pin as a 16-bit integer and this is used as a handle for the other port I/O functions. While optimisation is of no consideration for this course, there is no performance penalty when compiled with optimisation.

3.2 *Tiny GL*

A tiny graphics library was provided for the students. This allows them to draw points, lines, and text (stepped

```
#include "system.h"
#include "led.h"
#include "pacer.h"

#define FLASH_RATE_HZ 20

void main (void)
{
    system_init ();
    led_init();
    pacer_init (FLASH_RATE_HZ * 2);

    while (1)
    {
        pacer_wait();
        led_toggle ();
    }
}
```

Listing 2: A simple LED flasher using a paced loop.

and scrolling). Students can create their own fonts using ASCII art and these are converted by a Python script into a header file defined data structure.

3.3 Test scaffold

Debugging embedded systems software is difficult, even with a debugger. So to help the students debug their programs a test scaffold was provided. This replaced all the ATmega32u2 specific device drivers with native versions that ran on the host computer. For example, the dot matrix display was simulated on the terminal with an exponential filter for each pixel to provide persistence. Similarly, the navigation switch was mapped to the arrow keys. A real-time timer with a userspace interrupt was used to refresh the display and sample the keyboard.

Each application could be compiled against the test scaffold using a separate makefile. This allowed the the students to insert `printf` statements for debugging. The makefile was automatically generated by a Python script inspecting the student's program files.

3.4 Multitasking

Students find the concept of concurrency difficult, for example, how to refresh the display while waiting for data from the UART. To help them with this they are introduced to two patterns: paced loops and a simple run-to-completion scheduler. Examples, of these are shown in Listing 2 and Listing 3. These patterns allow simple concurrency without introducing race conditions.

```
#include "system.h"
#include "led.h"
#include "task.h"

#define FLASH_RATE_HZ 20

static void led_flash_task (void)
{
    led_toggle ();
}


int main (void)
{
    task_t tasks[] =
    {
        {.func = led_flash_task,
         .period = TASK_RATE / FLASH_RATE_HZ},
    };

    system_init ();
    led_init ();

    task_schedule (tasks, 1);
    return 0;
}
```

Listing 3: A simple LED flasher using a run-to-completion scheduler.

3.5 Dependency graphs

To help the students visualise the dependencies between the many files comprising their programs, a Python script was written that would generate a number of graphs from a given makefile using graphviz[5]. These graphs include a module dependency graph, a file dependency graph, a makefile dependency graph, and a call graph.

For example, the code for a simple demo application is shown in Listing 4. This simply turns the LED on while the button is pressed. It uses four modules: device drivers for the LED, button, and I/O ports and a system module for overall system configuration (such as the watchdog timer and clock oscillator).

The dependency between the modules is shown in Figure 3 (all the device driver modules depend on the system module but this has been removed for clarity). The callgraph is shown in Figure 4, with the functions grouped by file. Students often have difficulty understanding the compilation process. They seem to assume that the compiler is omniscient. The makefile dependency graph (see Figure 5) is useful to show that each C file is compiled independently. This helps to explain why header files are important.

[5] www.graphviz.org

```
#include "button.h"
#include "led.h"
#include "system.h"


int main (void)
{
    system_init ();

    led_init ();
    button_init ();

    while (1)
    {
        if (button_pressed_p ())
        {
            led_on ();
        }
        else
        {
            led_off ();
        }
    }
    return 0;
}
```

Listing 4: A simple demo application `ex4.c` (with comments stripped for compactness).



Figure 3: `ex4` module dependencies. The system module has been removed to reduce clutter.

3.6 Documentation

All the example applications and device drivers were commented using Doxygen mark-up. The generated documentation could be accessed through a web-broswer. In addition, a wiki was used for generic documentation. This could be edited by the students and was useful for describing how to set up the toolchain on the students' computers.

4 Discussion

The chief disadvantage of the ATmega32u2 is that it does not include an analogue to digital converter (ADC). It does have internal comparators so an ADC can be constructed with an external capacitor. A similar device with internal ADCs has 48 pins rather than 32.



Figure 4: Call graph for the `ex4` program. The PIO functions are implemented as macros and are not shown.



Figure 5: Make dependency graph for the `ex4` program. The leftmost nodes are phony targets; for example, the program target builds and downloads the program to the UCFK4.

The Arduino[6] has become a popular and cheap platform for microcontroller development. However, when the UCFK4 was developed, the Arduino did not have a USB bootloader and this was one of the main requirements.

Teamwork is important for engineers and so assignments using the UCFK4 were conducted in pairs. The students had to write simple multi-user games of their choice;

[6] `www.arduino.cc`

examples included battleships, tetris, rock-paper-scissors, pong, hangman, and maze games. The greatest difficulty the students had was with the infra-red communications; this introduced them to the concepts of an unreliable communication channel and race-conditions.

Prior to the assignment the students performed a number of labs to familiarise themselves with I/O ports, timers, and paced-loops. In each case the students had to program the I/O ports directly to read the pushbutton and drive the LED. They then wrote simple C modules to abstract the hardware dependence. Finally, they rewrote their modules to use the I/O port and timer abstraction modules while maintaining the same interface. The philosophy was to briefly show the students how the peripherals are controlled by memory-mapped registers but to get them to use hardware abstractions.

5 Conclusions

The UCFK4 hardware and software has been in use for two years and from a teacher's perspective it has helped teach introductory embedded systems. Further research is required to formally assess the learning outcomes for the students.

6 Acknowledgements

The author would like to thank the Digital Systems Laboratory technician Philipp Hof for his assistance with designing and manufacturing the UCFK4.

References

[1] M. Grimheden and M. Törngren, "What is embedded systems and how should it be taught?—results from a didactic analysis," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 4, no. 3, pp. 633–651, 2005.

[2] T. Hamrita and R. McClendon, "A new approach for teaching microcontroller courses," *International Journal of Engineering Education*, vol. 13, pp. 269–274, 1997.

[3] J. Dolinay, P. Dostalek, V. Vasek, and P. Vrba, "Teaching platform for lessons of embedded systems programming," in *Proceedings of 13th International conference on Automatic Control, Modelling & Simulation ACMOS 11*, pp. 158–160, 2011.

[4] D. Franklin and J. Seng, "Experiences with the Blackfin architecture for embedded systems education," in *Proceedings of the 2005 workshop on Computer architecture education: held in conjunction with the 32nd International Symposium on Computer Architecture*, p. 3, ACM, 2005.

[5] M. Benjamin, D. Kaeli, and R. Platcow, "Experiences with the blackfin architecture in an embedded systems lab," in *Proceedings of the 2006 workshop on Computer architecture education: held in conjunction with the 33rd International Symposium on Computer Architecture*, p. 2, ACM, 2006.

[6] L. Huettel, A. Brown, K. Coonley, M. Gustafson, J. Kim, G. Ybarra, and L. Collins, "Fundamentals of ECE: a rigorous, integrated introduction to electrical and computer engineering," *Education, IEEE Transactions on*, vol. 50, no. 3, pp. 174–181, 2007.

[7] D. Maskell and P. Grabau, "A multidisciplinary cooperative problem-based learning approach to embedded systems design," *Education, IEEE Transactions on*, vol. 41, no. 2, pp. 101–103, 1998.

[8] J. Song, X. Mu, H. Xu, and M. Yoder, "Learning and practicing fundamentals of electrical and computer engineering through building and programming a microcontroller with multiple peripherals," in *Frontiers in Education Conference (FIE), 2010 IEEE*, pp. S2F–1, IEEE, 2010.

[9] M. Wang, V. Callaghan, M. Lear, and M. Colley, "Teaching next generation computing skills; the challenge of embedded computing," *Intelligent Campus*, 2011.

[10] V. Callaghan, "Buzz-boarding; practical support for teaching computing based on the internet-of-things," in *1st Annual Conference on the Aiming for Excellence in STEM Learning and Teaching, Imperial College, London & The Royal Geographical Society*, pp. 12–13, 2012.

[11] D. Brylow, "An experimental laboratory environment for teaching embedded hardware systems," in *Workshop On Computer Architecture Education: Proceedings of the 2007 workshop on Computer architecture education*, vol. 2007, pp. 44–51, 2007.

[12] D. Bolanakis, E. Glavas, and G. Evangelakis, "A multidisciplinary educational board system for microcontrollers: Considerations in design for technically accurate custom-made platforms," in *Information Technologies and Applications in Education, 2007. ISITAE'07. First IEEE International Symposium on*, pp. 391–395, IEEE, 2007.

[13] R. Beneder, M. Kramer, and P. Brejcha, "A two stage approach for a cost-effective and versatile debugging unit and starter platform," in *Mechatronics and Embedded Systems and Applications (MESA), 2012 IEEE/ASME International Conference on*, pp. 330–335, IEEE, 2012.

Increasing Student Engagement in First Year Engineering

Craig. A. Watterson, Will. N. Browne, Dale. A. Carnegie,
School of Engineering and Computer Science
Victoria University of Wellington
Wellington, New Zealand
Craig.Watterson@vuw.ac.nz

***Abstract*—The paper details the changes undertaken to the introductory Bachelor of Engineering (BE) paper Foundations of Engineering (ENGR101) at Victoria University, Wellington, New Zealand (VUW) to improve first year student engineering experience. The paper also reports on student reactions to these changes based on information gained from surveys. It finds the introduction of active learning in 2011 increased the students passing with grade B from 45.3% to 50%. Furthermore, it was complementing the active learning with related and transparent academic learning in 2012 that increased this pass rate to 63.4%**

Keywords: curriculum, engineering education, retention, teaching and learning communities.

INTRODUCTION

This paper will discuss the changes made to Engineering 101 - Foundations of Engineering (ENGR101) to improve student retention and engagement and the results of those changes. In order to do this some understanding of the background to these changes is necessary.

New Zealand has six universities that offer a four year Bachelor of Engineering degree (BE). Victoria University of Wellington (VUW) offers a degree with specializations in Electronic and Computer Systems Engineering (ECEN), Network Engineering (NWEN) and Software Engineering (SWEN). Students are required to obtain a B (65%) average across the core courses in their first year, Part I, in order to be able to enter Part II and progress further in the degree program.

In New Zealand, the latest government Tertiary Education Strategy clearly articulates the need for accountability in terms of increasing student performance in tertiary education [1].

Since the inception of the BE at VUW in 2007, nearly sixty percent of the students have failed to advance past Part I of the degree, instead opting for the less demanding BSc equivalents. Between 2007 and 2012 we conducted grade data analysis. This comparison of student first year tertiary grades and their final year secondary school performance indicated that there were major gaps in levels of knowledge and learning abilities amongst most students. Crucially, the data revealed that prediction of success in first year tertiary engineering based on secondary school results was only possible for the ECEN major and even then problematic [2].

This is a significant concern since the New Zealand Government began in 2012 to link 5 percent of government funding for tertiary providers to student performance and dropout rates [3]. The Tertiary Education Commission (TEC) now measures the performance of Tertiary Education Organisations (TEO's) (NZ universities) against the following Educational Performance Indicators:

- Qualification completion rate: this measures successful qualification completions as a proportion of total enrolments in a given year.
- Course completion rate: this measures successful course completions as a proportion of course enrolments in a given year.
- Retention (completion/continuation): this measures the proportion of students in a given year that complete a qualification, or re-enrol at the same organisation in the following year.
- Progression: this measures the proportion of students who progress to a higher level of study after completing a qualification in the previous year.

Importantly, the TEC states: 'Longer term, it may be possible to include other outcomes, such as students' success in the labour market' [3], making it clear that future funding may depend on our ability to ensure that not only we pass a greater number of students, but those graduates then obtain employment in their field of study. Failure to take stock of these changes would place Engineering at VUW in a very difficult situation.

I. BACKGROUND

It is well understood in education that there are large transitionary factors involved in students making the step from secondary school to first year university and one of the most significant is the change required in the students methods of learning [4].

The National Certificate of Educational Achievement (NCEA), the predominant qualification framework for students entering tertiary study in New Zealand was designed to improve retention at secondary school level and reward excellence in the very top students. However, it provides a mismatch in terms of

learning style and environment to that which students face in tertiary engineering. For example, Hume's research into the NCEA science curriculum and student experience revealed that assessment for qualifications is driving the senior school classroom. School staff place a huge emphasis on students achieving success in qualifications, which directly impacts on the classroom curricula and the methods teachers use to teach course content [5]. This places emphasis on encouraging a memorization and regurgitation style of learning behaviour that poorly prepares the student for tertiary study where deep understanding and insight is required.

Most first year engineering courses in Part I of Engineering at VUW, particularly those taught in mathematics and physics require a largely independent learning style where information is provided to students, predominantly through lectures and laboratory work requiring the learner to adapt to conceptual learning. A recent New Zealand study, Thomas, et. al. [6] in 2010, noted that students entering tertiary study in mathematics faced a very different style of learning to the one they are used to in many Secondary Schools. The report stated: 'There is some disjunction in teaching emphasis and style between schools and tertiary institutions. Examples include a greater emphasis on applications and modelling, practice of examples (e.g., textbook questions), communication skills, teacher-student discussion, collaborative work and factual knowledge in school, while tertiary institutions value conceptual thinking, rigour and problem solving. In turn, students found the tertiary courses deeper and more challenging.'

International research also suggests that there is a major difference in the learning styles of engineering students compared to the teaching methods of engineering lecturers [7]. Felder and Spurlin found that in ten English speaking countries a consistent learning style was found for engineering students that differed from that of academic staff [8]. They noted that the approach used to teach engineering courses emphasized lectures over active engagement, which suits the more reflective and verbal learners over more active and visual learners. This form of learning concentrates on theoretical abstractions and mathematical models rather than experimentation and engineering practice, which in turn suits intuitive learners over sensing learners.

They also noted that engineering classes are predominantly taught in a self-contained manner that does not make the connections between material in other courses or relate to students own experiences, favouring a sequential learner over a global learners. By increasing the range of instruction to reach all types of learner, performance between the student learner groups decreased [8].

In summary, the literature points to a growing gap in the learning needs of students and what they experience. This factor was apparent in the results of the surveys detailed in section II below.

II. Method – Base Line

In April 2010 we conducted a survey of first year students and many of the comments included comments relating to the core engineering course Engineering 101 - Foundations of Engineering (ENGR101). Here is a sample:

- Practical and theoretical components seem disconnected (ENGR101). Lack of textbook – requires comprehensive notes (ENGR101). Lack of discussion/report on task-based labs, projects completed without direction or reflection (ENGR101, car challenge)
- 1st lecturer was good for ENGR. Current one is boring.
- Looking for more practical content in class.
- Link what we do in the labs to the material going on currently in the courses.
- Some lecture content feels rushed and not properly explored.
- Some assignments are not clearly announced. I knew nothing of the ENGR101 LAB prework till in the lab.

The end of year, first year student survey conducted in November 2010 contained similar findings:

- I found that Engr101 is particular was not what I was expecting, nor what I had heard about the course -- expected more hands-on/practical work.
- Teaching was generally good. Course content was good. Workload was a little high. Should involve more hands-on experience in first year.
- Lack of support for students struggling in some areas. Fairly dull second trimester courses (Comp103), came out of Engr101 not feeling like I learned anything or gained anything from it.
- Should have more practical work.
- ENGR101 put me off of BE. Made me wonder if other engineering courses are going to be anything similar and whether I should just switch to computer science.
- Engr101 felt like a bit of a waste of time.
- Interestingness of the ENGR*** courses.
- Unorganised. No clear outline of where course is headed/what is needed to be learnt.
- I don't think they explain why we need to do the courses we do. Marking sometimes seems to be unfair.
- Requires a lot of maths
- More practical, hands on. More emphasis placed on physics, calculus and electronics.
- Parts of the physics papers 1st year, did not see the relevance to ECEN.
- Can be difficult to schedule Part 3 [three coherent courses outside a student's specialization] because it can end up doing away with stuff you're interested in just so you can satisfy a less important criteria
- Most of the [interesting] courses (like Nwen401... 406) are only in the 4th year.

Significantly the surveys indicated that the students generally felt that the existing BE first year and in particular the ENGR101 course structure provided insufficient opportunities to do practical engineering. As such the course was not meeting the expectations of the students or providing a supportive

teaching and learning community crucial to student success [9] [10].

Students also identified that Part I of the BE felt like it contained a small engineering component. In part, this is directly attributable to the nature of the first year structure where all engineering students take only the following three papers (which can be called Engineering papers), out of a total eight: ENGR101, COMP102 Introduction to Computer Program Design and COMP103 Introduction to Data. All other first year papers (with the exception of SWEN102 Introduction to Software Modeling, which is compulsory for SWEN majors only) are taught by the physics or mathematics departments.

Focus group interviews were conducted in May 2011 with eight first year engineering students (one female and seven male) at Victoria University of Wellington. The interviewers from the Centre for Science and Technology Educational Research (CSTER) Waikato University used semi-structured interview protocols and analysed the findings in accordance with qualitative analysis principles. They noted:

"The nature of the first year of the course was the source of some discussion. One first year student was aware of the overall course structure, and so understood that the first year was mainly maths and physics and that engineering wouldn't really be dealt with until second year. Other students felt frustrated by this, and were impatient to get into engineering and resented all the 'maths and stuff' they have to do."

These findings strengthened the information gained from the surveys and focus groups indicating that relevance and context were very important; students want to feel like they are studying the field they have chosen. Clearly the expectation of engineering students does not match the reality of what must be learned in tertiary study in order to become an engineer. The issue of educating potential students in what an engineer must learn was dealt with separately and falls under a series of recruitment initiatives designed to raise awareness of engineering requirements amongst secondary school students, parents, teachers and careers advisors and lies outside the scope of this paper.

III. Method – 2011 Improvements

Informed by the surveys, focus groups and literature we decided to trial some changes in the core ENGR101 paper. As mentioned above we were unable to substantially influence most of the first year courses as many of the students' core first year courses are not taught by the Faculty of Engineering.

Modifications to ENGR101 involved changing the teaching to meet the expectations of the students without compromising on the content of the course. The hands-on, practical work was formerly based on laboratory sessions, where students followed a set of instructions to achieve a known result. Laboratories were distinct and could be completed 'successfully' without understanding, learning or extension of the presented material. The laboratories were also timed to ensure that the 'slowest' student completed, e.g. three laboratory sessions to build a PC, install an operating system and set up networking between them, where competent students can achieve all of these in one session. Thus, laboratories were identified as an important area in need of change.

Change involved teaching and learning practice and emphasized explicit styles of learning rather than simply course content. The changes to the ENGR101 course were designed to assist in transforming a student's learning style (from consuming to generating knowledge), without creating a prohibitive gap in their learning experience. Woodhead differentiates between education as 'transforming people' (commonly the modern secondary school approach) and initiation into a body of knowledge' (conventionally the approach at University) [11]. The former places the emphasis on the student as a consumer of knowledge, where this knowledge has been modularized and layered, such that they can choose topics that they enjoy at an attainable level. The latter emphasizes complete knowledge of the subject, such that the student may utilize/generate new knowledge. Further, whereas the former emphasizes the middle layers of Bloom's taxonomy [12], the latter requires all levels.

To enable students to start taking responsibility for their own learning, students first needed to be able to group fundamental related ideas and details together in their memory (as 'Chunks') and then be able to recall the necessary information in order to put it into practice [13]. Time management and an attitude of striving to pass assignments at the first attempt were also important skills to foster. The course content was revised to ensure subjects were relevant and digestible; although topics were split up they were related and flowed on from each other naturally. Study skills were integrated into the curriculum in liaison with the University's Student Learning Support Service. A custom textbook was created that as well as containing the standard academic information also contained study advice, which ranged from adapting to university life, time management, report writing, problem-solving techniques, revision guidance and exam practice. Lectures on the benefits of engineering were introduced on topics such as, sustainability and reliability, and Engineers without Borders.

To show that engineering requires a complete and broad taxonomy of learning, subjects were integrated together, such that the students had to learn to decouple their 'chunks' of knowledge into integrated projects, thus learning to apply fundamental engineering knowledge in projects.

The next section details the autonomous vehicle challenge (AVC) which is the main capstone project for the course. In the AVC students constructed an autonomous robotic vehicle and are introduced to a large number of diverse learning processes and a team based active learning exercise. This replaced the majority of 'recipe'-like individual laboratories, although a few core experiences were kept, e.g. students still build a fully networked PC from scratch, but now in only one session with time/help available outside of the standard laboratory hours to complete the task is needed.

IV. Autonomous Vehicle Challenge

The main objectives of the ENGR101 AVC project were:

- To set a practical engineering task that requires design compromises.
- To encourage students to take responsibility for their own learning and become independent learners.
- To show that engineering encompasses a wide range of learning and includes subjects such as physics and mathematics.
- To introduce students to relevancy of all aspects of engineering to real world situations.
- To encourage students to feel part of a community of engineers
- To make students aware that support help exists and encourage its uptake.
- To improve grade performance.

An important part of the AVC was peer collaboration where students work together to identify what they knew and compensate for gaps in knowledge. If there was a delay between a student realizing they had a gap in their knowledge and a tutor being available, the student often did not seek out help. This was evidenced by poor attendance at the help desk and peer assisted tutoring system. When questioned, students held the mistaken belief that they could "always catch up" or "could repeat any failed assignment" despite lecturers and senior students explaining to them that this was not the case.

Help was most successful when a tutor or lecturer was immediately available to answer a student's problem as it occurred to them. Thus laboratory sessions with tutors and lecturers being available for up to 15 minutes after the introduction of new material, had a high level of engagement.

In 2011 we decided to set up Peer Assisted Study Support (PASS) groups to help students. PASS groups are study groups run by a good student who did the course the previous trimester.

To introduce students to practical and relevant aspects of engineering without compromising the academic content or reducing the threshold concepts, laboratories were aligned with teaching and theoretical content was introduced with guidance provided on how it formed the basis of practical engineering systems. An important aspect of this was the development of the Autonomous Vehicle Challenge (AVC) project.

V. Results 2011

The success of the modifications to ENGR101 can be seen in the general comments of surveys conducted in 2011 and again in specific questions surveyed on the AVC in 2012.

19 out of 38 respondents in end of year survey in 2011 noted that ENGR101 was the course they felt most involved in and majority of the comments related to the AVC project. Here are some sample comments:

- ENGR 101 The practical aspect kept me very interested
- ENGR 101 had a practical group component which encouraged involvement and team work
- ENGR 101 because there was more practical and community than any other course
- ENGR 101 Encouraged group work and team buildings through a series of problem solving activities
- ENGR101 Group atmosphere and tasks
- ENGR101 the labs got everyone involved in the autonomous vehicle challenge
- ENGR101 as the large 5 week project required a large level of team work and interaction with other members of the class and combined ideas and skills from all other subjects which are part of the degree
- ENGR 101 because of the in lecture class discussions and because the robot assignment forced us to work as a team to overcome the challenges

We started PASS with Math 151 (An introduction to linear algebra) in 2011 and had a poor uptake, with only 2 students out of 112 attended PASS. This suggested that although students could be engaged through active learning (e.g. AVC project), they had still not transitioned to typical tertiary learning where they are required to take independent responsibility for their own learning.

VI. Method – 2012 Improvements

In 2012 team based tutorials were introduced where the first half of the tutorial was constructive (construction of a steam powered robot) and the second half was on a related academic subject. Team based learning and problem solving were encouraged in both halves.

Interestingly, these tutorials were first opened to the middle performing students only, due to resource limitations, although other students could apply to attend. These 'invite only' tutorials were seen as desirable and well attended, whereas previously the 'open to anyone' tutorials were not well attended.

In 2010 and 2011 the first two weeks worth of lectures in the trimester where introductory, e.g. welcome messages, professional considerations and soft skills, such as note taking. In 2012 more academic content was introduced in the first two weeks, with an 'epiphany' test towards the end of week two. This was marked and returned to the students on the Friday morning of week two, so that they could seek pastoral help – students may change or drop courses in the first two weeks of the trimester without penalty.

These activity-led tutorials were in contrast it with the University centric Peer Assisted Student Support (PASS) Scheme used for the equivalent mathematics course in the first trimester. PASS tutorials are like traditional tutorials where a senior student is paid to lead the students through practice questions. With the advent of a Pastoral Care staff member in 2012 we saw an increase in the uptake of students attending PASS groups for MATH151, with 28 students attending. The academic results of PASS were however disappointing with only 6 out of 28 students (21%) receiving a B or better grade for Math151.

In trimester 2 2012, we have again run PASS groups for ECEN220, MATH142, MATH161, SWEN224 and have had improved attendance though it is too early for the academic results.

VII. RESULTS 2012

Introducing academic content, reinforced by a test, served two purposes. Firstly, it identified students who were poorly prepared for tertiary level education enabling them to access advice or change to more appropriate routes. Ultimately, 20 students have left the BE early rather than wait for the more dire consequences of failing end of year exams. This also encouraged students to seek pastoral help and get advice on other sources of assistance, such as PASS. Secondly, it encouraged good working practices from the start of the course, which continued throughout the trimester.

Activity based tutorials had consistently high attendance, i.e. an average of 40 students per week instead of two in a term. As the trimester progressed it was noted that students wanted more time to be spent on the second half of the tutorial and continued with the problem-solving groups for other subjects after the tutorials.

In trimester two of 2012 we conducted specific survey questions aimed at seeing if the changes made to ENGR101 and specifically the implementation of the AVC project had met our objectives. We had 82 of 152 first year BE students; 73 Male and 9 female complete the surveys. There was no statistical difference between both groups.

Students were asked to rate their experience of the AVC in ENGR101 during Trimester 1, and asked students on a scale of 1 to 10, where 1 = totally disagree and 10 = totally agree, to indicate how much they agreed with a set of statements. Table I. below details the response data and questions asked.

TABLE I. AVC QUESTIONS AND RESPONSES

Statistic	Mean	Variance	Standard Deviation
Participating in the AVC project made me feel that I was studying Engineering	8.24	5.08	2.25
It helped me see how Physics was related to Engineering	6.4	6.38	2.53
It helped me see how Maths was related to Engineering	6.73	6.26	2.5
It helped me see how Software and Hardware are related	8.95	3.14	1.77
It helped me see the real world applications of Engineering lecture material	7.97	4.65	2.16
My interest for engineering increased as a result of the AVC project	7.21	4.28	2.07
I enjoyed working in a team for the AVC project	7.85	5.24	2.29
I preferred the 6 week AVC project to the weekly labs on individual topics	7.44	6.53	2.55
I would prefer the AVC assessment to be fully competition based	5.26	6.45	2.54
I felt that the AVC written report provided me with valuable skills	6.23	5.72	2.39
I would prefer projects like this over written assignments and tests in other courses	7.74	6.72	2.59

These surveys also contained a question on which course at first year students felt they were most involved in and why. 47 out of 79 respondents noted that ENGR101 was the course they felt most involved in and majority of the comments related to the AVC project. Here are some sample comments:

- ENGR 101 because of the team coordination it required
- ENGR 101 the labs were interactive, the avc allowed us to meet lots of people engineering events were held and there was more than one lecture at a time
- ENGR 101 working as a team for AVC. Variety of lecture material. Tasks for outside of the class
- ENGR 101 because I took part in activities that were outside of the lecture times
- ENGR 101, the labs and AVC challenge involved me interacting with other people which helped me learn.
- ENGR 101- due to the practical labs and AVC felt more involved actually being able to use my skills

The responses showed that the AVC project was successful in engaging students with engineering. The teamwork, interactive nature and expression of skills learnt in lectures were all considered beneficial. Although students preferred this work over written assignments and tests, they were not opposed to written reports.

A comparison of retention rates in Table II below, with previous years, shows a marked improvement in grade performance.

TABLE II. ENG101 GRADE PERFORMANCE

GRADE	2010-STUDENTS	2011-STUDENTS	2012-STUDENTS
A+	17	5	16
A	10	6	26
A-	6	13	16
B+	6	9	22
B	10	19	17
B-	7	11	8
C+	7	5	12
C	2	6	2
D	9	13	11
E	12	17	23
TOTAL STUDENTS	86	104	153
TOTAL STUDENTS WITH B PASS	39	52	97
% PASS WITH B	45.34%	50%	63.39%

VIII. Discussion

These pedagogic changes led to significant continuation improvements, up from 45.3% in 2010 to 63.4% in 2012.

Attendance at PASS was sporadic and many students wanted to attend just to 'obtain answers to the assignments', rather than gain the skills needed to complete the assignments themselves. The enquiry based style of ENGR101 tutorials were seen as favorable as it encouraged learning rather than "pass'ing" the course.

Surveys indicated that where students understood not only threshold concepts for the subject, but its relevance to engineering and the purpose of the teaching method, then engagement and success was high. The opposite was also true: in the AVC 2012 a networking challenge was introduced, which although a relevant/engaging activity was deemed too hard and so was poorly rated. Conversely, the presentation skills activity, which was not well motivated and seemed trivial, did not score highly.

IX. Conclusion

Although secondary level study did not prepare the majority of our students for the rigors of tertiary level study, satisfactory academic performance could be improved by active learning complemented by academic learning that was relevant to engineering, well-motivated and whose purpose was understood by students. The transition to tertiary learning occurred late in the first trimester as students began to own their learning experience once there was demonstration of what was required.

References

[1] Tertiary Education Strategy 2010-2015, Ministry of Education, New Zealand Qualifications Authority. Available online: http://www.minedu.govt.nz/theMinistry/PolicyAndStrategy/TertiaryEducationStrategy.aspx

[2] Dale A. Carnegie, Craig Watterson, Peter, Andreae, Will N. Browne, 'Prediction of Success in Engineering Study', IEEE Engineering Education Conference: Educon, Morocco, conference proceedings 2012

[3] Tertiary Education Commission, New Zealand Government: http://www.tec.govt.nz/Funding/Policies-and-processes/Performance-linked-funding/

[4] Krause, K., Hartley, R., James, R., & McInnis, C. (2005). The first year experience in Australian Universities: Findings from a decade of national studies. Melbourne, Victoria: University of Melbourne, Centre for the Study of Higher Education.

[5] Anne Hume, PhD Thesis: Carrying out a Practical Investigation: The Student Experienced Curriculum, University of Waikato, 2006. P.341 available online: http://researchcommons.waikato.ac.nz/bitstream/handle/10289/3991/thesis.pdf?sequence=1

[6] Thomas, M., Klymchuk, S., Hong, Y Y., Kerr, S., McHardy, J., Murphy, P., et al. (2010). The transition from secondary to tertiary mathematics education. Wellington: Teaching and Learning Research Initiative.

[7] Richard M. Felder, & Linda K. Silverman, Institute for the Study of Advanced Development, Learning and Teaching Styles In Engineering Education, Engr. Education, 78(7), 674–681 (1988)

[8] Felder, Richard M. and Joni Spurlin, 2005. Applications, Reliability, and Validity of the Index of Learning Styles (Acrobat (PDF) 223kB Feb2 06), International Journal of Engineering Education, v. 21, n. 1, pp. 103-112, see p.110.

[9] Clark, J., Dodd, D., & Coll, R. K., "Border crossing and enculturation into higher education science and engineering learning communities", Research in Science & Technological Education – Higher Education Science and Engineering Learning Communities, Vol.26(3), 2008, pp.323 – 334.

[10] Parkinson, T. J., Hughes, H., Gardner, D. H., Suddaby, G. T., Gilling, M., and MacIntyre, B. R., Enga*ging Learners Effectively in Science, Technology and Engineering: The pathway from secondary to university education*, Ako Aotearoa – The National Centre for Tertiary Teaching Excellence, 2011.

[11] Woodhead, C., A Desolation of Learning-Is This The Education Our Children Deserve?, Pencil-Sharp Publishing Ltd, 2009.

[12] Bloom B. S. (1956). Taxonomy of Educational Objectives, Handbook I: The Cognitive Domain. David McKay Co Inc.

[13] Awad, E. M., Knowledge Management, Pearson Education India, 2004, p.72.

Pose Tracking using Ultrasonic Sensors in a Small Scale Environment

Raymond ShuHe Liu, Chee Kit Wong and Patrick P.K. Lim

Industrial Research Ltd
24 Balfour Road, Parnell, Auckland, New Zealand
Email: {r.liu, k.wong, p.lim}@irl.cri.nz

Abstract: In this paper, we present a pose (position and orientation) tracking system for the localization of a tabletop platform for upper limb rehabilitation using ultrasonic sensors. Trilateration techniques are applied to time-of-flight measurements between two ultrasonic transmitters mounted on the mobile device and two receivers fixed on the tabletop. The advantages of this system are: 1) ease of setup, 2) does not require calibration, 3) not affected by slippage and hence, no accumulative errors and 4) no maintenance is required. However, the design of the system is not straightforward as a number of issues must be addressed during the design process. Experimental results demonstrate the accuracy and robustness of the system when operating within the specified workspace.

1 INTRODUCTION

An important aspect of an upper limb rehabilitation device is the ability to track a patient's arm movements as they perform a task. Being able to do so, the patient's ability can be ascertained and monitored over time. For this reason, tracking the device's location and orientation is very important.

For this application, traditional approach of dead reckoning is unsuitable due to inaccuracies in encoder measurements from wheel slippage. Whilst object tracking using vision systems such as [1] works, it requires setup and calibration and hence, not particularly suitable for users with limited limb movement. Furthermore, performance of such systems is affected by environmental conditions such as lighting, which in most circumstances, cannot be controlled.

An alternate approach has been described in [2] where two optical mice are used to track the location and orientation of the rehabilitation device. Whilst similar systems have also been proposed for mobile robots in [3, 4], optical mouse sensors have two major disadvantages. Firstly, the physical installation of these two mice is very difficult and needs to be accurate and durable - the mouse sensors must have very good contact with the operating surface and accurately positioned with respect to the device and to each other, making it a challenging task to mount the two mice. Secondly, the mouse sensors can slip or misread within the movement, this will cause the localization and orientation error accumulated in memory. And finally, the optical mouse sensors are not compatible with all surfaces.

Various localization approaches exist, however most involve tracking in larger scale environments and hence setup of the system is not straightforward. For example, in [5] an infrared system which uses differential phase-shifts to find the mobile-robot position in the intelligent spaces; an ultrasound method has been proposed for accurate positioning in [6]; and the DS-CDMA and SDMA wireless local positioning system for airport (indoor) security [7].

In this paper, we describe a novel pose tracking system based on ultrasonic sensors for table top operations. This system not only overcomes aforementioned problems, but also has a high degree of resolution within a 1500mm x 800mm workspace. This is because the speed of an ultrasound in air is sufficiently slow to allow the time-of-flight of a signal to be accurately measured between the transmitter and receiver.

Sections 2 and 3 provide details of the hardware implementation; and computation of pose respectively whilst Section 4 discusses some of the critical design considerations. Section 5 presents the experimental results of the system.

2 SYSTEM DESCRIPTION

The pose tracking system comprises of a microcontroller, signal generation and timing circuitry, high voltage driving circuitry, ultrasonic transducers and pre-amplifier circuitry. Figure 1 below provides an overview of the relationship between each of these components.



Figure 1. System overview

2.1 *Arduino Mega 2560 board*

The Arduino Mega 2560 development board is responsible for all arithmetic calculations, as well as reporting the location and orientation of the mobile device to the PC via RS232

connection. It is equipped with an 8-bit AVR microcontroller, which is pre-burned with a boot loader, thus allowing the process of uploading new code without using an external programmer.

2.2 Transceivers SRF005

The SRF005 ultrasonic range sensor module measures the time-of-flight of the traveling ultrasonic signal in air between the transmitter and receiver, and outputs the result in the form of a pulse signal where the width is proportional to the signal time-of-flight.

The SRF005 ultrasonic module is connected to the Arduino development board via the trigger input pin. When the system is ready to perform a measurement, the Aduino board outputs a trigger signal, and accordingly, the SRF005 will output a burst signal to the transmitter. At the receiver end, it detects the burst signal and outputs a pulse to the Arduino microcontroller.

2.3 Ultrasonic Transmitters and Receivers

Certain synthetic polymer films have piezoelectric properties [8], also referred to as piezopolymer films or simply piezoelectric films (PVDF). Piezoelectric films are small, flexible, lightweight and inexpensive compared to electrostatic transducers. However they have low sensitivity which means as transmitters, they need a high driving voltage, and as receivers they need a pre-amplifier to reduce the noise.

PVDF ultrasonic transmitter and receiver for air ranging has been designed by Fiorillo [9, 10] and further characterized by Wang and Toda [11]. They have a wide frequency bandwidth, low resonance Q and excellent impact resistance. These characteristics offer unique solutions in many applications such as two-dimensional positioning, digitizer, object detection, and distance measurement.

Compared to ceramic transmitters and receivers, PVDF devices have superior horizontal beam directivity qualities, providing the transmitter coverage of 360° and the receiver ±150°. Furthermore they have a very low resonance Q value. Low Q value makes the signal rise and decay times much faster than conventional ceramic ultrasonic sensors which improves the accuracy of the time-of-flight measurement. Hence, PVDF ultrasonic transmitters and receivers was a logical selection for the work proposed in this paper.

2.4 Transmitter Driving Circuit

As mentioned previously, one of the drawbacks of PVDF ultrasonic transmitter is the high driving voltage, which is around 300Vp-p. However the SRF005 module can only output 5Vp-p. Therefore, a driving circuit is needed to bring the low voltage to 300Vp-p. In addition, it is important to match the electrical resonance frequency with the mechanical resonance frequency in order to achieve maximum transmitter efficiency.

Here, a driving circuit USDE-01 by Measurement Specialities is used to drive the 40 kHz PVDF transmitters. Since the equivalent circuit of the PVDF transmitters is capacitive, a transformer is used to create the electrical resonance. The high voltage side of the transformer is connected directly to the piezoelectric film of the transmitter.

A further consideration is the number of bursts emitted by the transmitter driving circuit, as it affects the maximum detected distance [12]. A single pulse drive signal is recommended for distance measurement applications such as ultrasound digitizers. A 12 pulses or higher signal can be used to maximize the detectable distance for object detection applications. However, a higher number of pulses increases the time between measurements and hence decreases the refresh rate of the system. For this application, bursts of 8 pulses are used on the driving circuit to achieve a suitable detectable distance, whilst maintaining a reasonable refresh rate.

2.5 Receiver Pre-amplifier

Ultrasound traveling through the air is converted by the piezoelectric film receiver into electrical signals, which is then fed to the pre-amplifier. A pre-amplifier US40KA-01 by Measurement Specialities is used because the signal received is very weak and noisy.

One purpose of this pre-amplifier is to reduce the noise interference with a bandwidth filter. It removes frequencies outside the desired operating band of the system, therefore increasing the signal to noise ratio (SNR). The other purpose is to amplify the receive signal for further processing, this pre-amplifier has the gain of 31.

3 LOCATION & ORIENTATION TRACKING

This section discusses the application of the basic principle of distance computation using time-of-flight data for location and orientation tracking through trilateration.

3.1 Distance Calculation

Step one is to determine the distance between receivers and transmitters, which is based on the time difference between the transmitted and received pulses.

When the Arduino microcontroller sends a trigger signal to the SRF005 module, a burst signal is transmitted and the time is recorded as the start time. The microcontroller then checks for an incoming signal from SRF005 transceiver by using one of the onboard external interrupts. When a signal is received, the end time of the sound wave is recorded.

The actual time of the sound wave can be calculated from the start and end times. However in practice, there is a fixed delay (offset time) due to the electronics and algorithm of microprocessor in the transceiver. Therefore the actual time of flight in air is:

$$t_{air} = t_{end} - t_{start} - t_{offset} \qquad (1)$$

With the speed of sound, s at 343.26 m/s at room temperature 20 °C, the distance can be calculated by:

$$d = s \times t_{air} \qquad (2)$$

3.2 Position Tracking

To determine the position of a mobile platform, a setup of two ultrasonic receivers (fixed in space) and one transmitter (mounted on the mobile platform) is arranged as shown in Figure 2. The trilateration algorithm is applied here to calculate the mobile transmitter's location.

The locations of the two receivers are fixed, with the left receiver used as the reference point and the origin in a two dimensional coordinate system. Based on this configuration, we have constrained the work space to positive y coordinates thereby simplifying the calculations. Since the positions of the receivers are fixed, their separation (*b* in Figure 2) is known. The other two distances, *a* and *c*, can be calculated by time-of-flight measurements as previously described.



Figure 2. Workspace with two stationary receivers and one mobile transmitter with their respective x- and y-coordinates in mm.

As shown in Figure 2, the two receivers and one transmitter links together to become a triangle $\angle ABC$. Let *a*, *b* and *c* denote the distances between them. Since *a*, *b* and *c* can be measured, our goal is to determine the coordinate x1, y1 of the transmitter within the workspace. Using the sides and angles relationship in right triangle $\angle ABD$, we have two equations:

$$x_1 = c \times \cos\beta \qquad (3)$$

$$y_1 = c \times \sin\beta \qquad (4)$$

In the equation (3) and (4), distance c is known. Hence, to determine x_1 and y_1, we need to calculate $\cos\beta$ & $\sin\beta$ first. Using Cosine Law on triangle $\angle ABC$, we derive

$$\cos\beta = \frac{b^2 + c^2 - a^2}{2bc} \qquad (5)$$

Now if we substitute equation (5) into equation (3), we have an equation with three known distances.

$$x_1 = \frac{b^2 + c^2 - a^2}{2b} \qquad (6)$$

Use the basic relationship between the sine and the cosine in Pythagorean trigonometric formula, $\sin^2\beta + \cos^2\beta = 1$, then rearrange to make the β as a subject, we have:

$$\beta = \sin^{-1}\left(\sqrt{1 - \cos^2\beta}\right) \qquad (7)$$

Substitute equation (7) into equation (4), we have:

$$y_1 = c \times \sqrt{1 - \cos^2\beta} \qquad (8)$$

Finally, substitute equation (5) into equation (8) to describe y_1 in terms of the known *a*, *b* and *c* values:

$$y_1 = c \times \sqrt{1 - \left(\frac{b^2 + c^2 - a^2}{2bc}\right)^2} \qquad (9)$$

3.3 *Orientation Tracking*

To track orientation, a second ultrasonic transmitter is required on the mobile platform. The position of each transmitter is tracked using the same approach as described in the previous section.

However it is not possible to enable both transmitters at the same time, as their signals will interfere with each other. Therefore a simple switching circuit was built to drive the two transmitters alternately, as shown in Figure 3 below. This approach has the added benefit of reducing the component costs by sharing a single driving circuitry for both transmitters.



Figure 3. Eletronic switches for drive two transmitters

For the switching electronics, two Phototriac Couplers were used as they have very high isolation voltage between the driving sources to the load. With this characteristic, the low voltage microcontroller signal can be used to alternately trigger the switches.

Figure 4 depicts the relationship between the components of the tracking system. (x_1, y_1) is the coordinate for the left transmitter and (x_2, y_2) is the coordinate for the right transmitter. These two transmitters have been placed next to each other, and mounted a fixed distance apart.



Figure 4. Workspace with two stationary receivers and two mobile transmitters and their respective x- and y- coordinates in mm.

When the coordinates of the two transmitters are known, the orientation of the mobile platform, α can be calculated using the formula:

$$\alpha = \arctan\left(\frac{\Delta y}{\Delta x}\right) = \arctan\left(\frac{y_1 - y_2}{x_1 - x_2}\right) \qquad (10)$$

3.4 *Potential Error Affecting Distance Measurements*

3.4.1 Accuracy of Distance Measurement in Practice

The ultrasonic transceivers operate at 40 kHz and based on the speed of sound (343.26 m/s) at 20 °C, the length of the waveform is about 8.6 mm.

Since the transmitters output 8 bursts at a time, we expect the receiver to detect 8 pressure peaks of the ultrasound waves. These pressure peaks are 8.6mm apart. In practice, if the SRF005 misses even one pressure peak, then the distance measurement has 8.6mm error.

3.4.2 Distortions Due to Environmental Variations

The speed of sound varies with temperature, humidity and air pressure, and so does the outcome of any distance measurement procedure relying on sound waves [13]. At a constant temperature, the ideal gas pressure has no effect on the speed of sound, because pressure and density (also proportional to pressure) have equal but opposite effects on the speed of sound, and the two contributions cancel out exactly [14].

The humidity has a small but measurable effect on the speed of sound (causing it to increase by about 0.1%-0.6%), because oxygen and nitrogen molecules of the air are replaced by lighter molecules of water [12]. In our application, we have ignored the effects of humidity.

In the Earth's atmosphere, the most important factor affecting the speed of sound is the temperature [12], as can be seen by their relationship in Equantion 11:

$$C_{air} = (331.3 + 0.606 \times \vartheta) m \cdot s^{-1} \qquad (11)$$

C_{air} - Speed of sound in air $m \cdot s^{-1}$

ϑ - Air temperature in degrees Celsius (°C).

If we use the above formula to compare the speed of sound at 10 °C and 25 °C, the difference is 9.09 ms^{-1}, or approximately, 0.18% change in speed per degree Celsius change in temperature.

Therefore if we want to have more accurate ultrasonic measurements between the transmitter and receiver, temperature compensation needs to be incorporated in the design. For our system, we can assume operation within room temperature, and hence there will be minimal fluctuations. Hence, temperature compensation was not considered.

4 CRITICAL DESIGN

4.1 Shielding and Low-Noise Requirements

As mentioned before, the ultrasonic transmitter driving circuit is running at a high voltage of 300Vp-p which produces interference to other nearby equipment. In particular, the PVDF receivers are extremely sensitive to noise. Therefore it is important to limit electromagnetic interference as much as possible. Below are some of the measures that formed an integral part of the design process:

- All the connection cables are as short as possible to minimize impedance, which reduces the risk of signal degradation.
- All the connection cables are shielded. Shielding protects the signal by reducing electrical noise from a noisy environment and also reducing the electromagnetic radiation that may interfere with other devices.
- The transmitter driving circuit and receiver unit are mounted in an aluminum enclosure and grounded.
- Use low noise power source such as linear regulated power supplies to minimize the power supply interference.

4.2 Installation of transmitters and receivers

There is a limitation on the vertical beam directivity of the transmitter and receiver units. They are set at the same level to achieve the best signal strength. Also the receiver has a limitation on the horizontal beam directivity. Therefore the zero horizontal beam angles need to face the transmitters.

5 EXPERIMENTAL RESULTS

All the experimental results have been taken at room temperature about 20 °C. The workspace is 1500 mm by 800 mm and is the most effective coverage area for these two ultrasonic transmitters and two receivers.

5.1 Behaviors of the signals in different stages

The measurements are captured by an oscilloscope to monitor the behavior of signals at different stages. Figure 5 shows the results of the time-of-flight bursts between one transmitter and one receiver.

The transmitter was placed approximately 170mm in front of the receiver and the direct propagation time between the transmitter and receiver was measured. The plot in Figure 5 represents the 8 cycle 40 kHz burst signal from the SRF005 transceiver. The beginning of this signal corresponds to the time-of-flight activation.

This burst signal goes into the transmitter driving circuit, it is amplified from 5Vp-p to about 300Vp-p as shown in Channel 2 by using a x10 probe. The high voltage drives the 40 kHz PVDF transmitter and transmits as ultrasound though the air. As this signal is attenuated when it is captured by the ultrasonic receiver, a pre-amplifier is required to amplify and remove any unwanted noise.



Figure 5. Behaviour of the signals at different stages

Channel 3 shows the received signal after the pre-amplifier. The first large signal is a pickup of the burst signal similar to Channel 2, and the second small signal is the actual signal received by the receiver. We can see most of the noise on the receive signal has been removed and the signal has been amplified by the pre-amplifier with a gain of 31. Therefore the end of the sound wave travel through the air can be defined by the SRF005 transceiver.

Channel 4 is the output from SRF005. After sending the low voltage burst signal, it sets the output high. When it detects the burst at the receiver, it sets the output low. The pulse width is the time-of-flight detected by SRF005 transceiver. However the actual time-of-flight is detected by the Arduino. We used the beginning of the burst signal at transmitter to be the start time, and the beginning of the received burst signal at the receiver to be the end time, as described in Section 3.1.

5.2 Test result between one transmitter and one receiver

The location and orientation precision is highly dependent on the accuracy of distance measured between the transmitter and receiver. To test the accuracy of the measurements, the time-of-flight of the ultrasonic signals between one transmitter and one receiver are measured at selected locations through the oscilloscope and their corresponding distances computed (Equation 2). The results are recorded in Table 1. The test has been taken at room temperature, approximately 20 °C and the speed of sound is 343.26 m/s.

TABLE I. COMPARING THE ACTUAL AND MEASURED DISTANCE BETWEEN ONE TRANSMITTER AND ONE RECEIVER

Actual Distance (mm)	Time-of-Flight measured by oscilloscope (us)	Measured Distance (mm)	Error (mm)	Error (%)
100	301.65	103.54	3.54	3.54
200	592.30	203.31	3.31	1.66
300	880.30	302.17	2.17	0.72
400	1170.60	401.82	1.82	0.46
500	1468.60	504.11	4.11	0.82
600	1748.60	600.22	0.22	0.04
700	2026.50	695.62	4.38	0.63

The recorded results show a maximum percentage error of approximately 3.54% at 100mm. However these results only represent the range from 100 to 700mm and the accuracy depends on the measured time on the oscilloscope.

5.3 Localisation and orientation performance

Two transmitters are mounted together on the mobile unit and moved within the workspace. We use the equations in Section 3 and output the results to the PC HyperTerminal. The HyperTerminal records the results and a graph plot is generated as shown in Figure 6.



Figure 6. Plot of location and orientation of mobile platform as reported by the ultrasonic tracking system

The co-located transmitters are moved in a pre-defined path. The location and orientation are recorded at an update rate of 20 times per second. As we can see when these co-located transmitters move along the path, we are able to approximate their location and orientation in real time.

However the update rate has been limited by only a single transmitter activating at a time and the calculation of orientation required the distances between two transmitters and two receivers. Because there are too many position results in figure 6, therefore we only picked five positions to demonstrate the localisation and orientation.

TABLE II. COMPARE THE ACTUAL AND MEASURE COORDINATE AND ORIENTATION

	Real coordinate			Measured coordinate			Error		
	x mm	*y mm*	β (°)	*x mm*	*y mm*	β (°)	*x mm*	*y mm*	β (°)
1	-200	500	90	-198	4967	88	-2	4	2
2	200	500	135	199	498	134	1	2	1
3	0	300	45	1	302	45	-1	-2	0
4	200	100	-90	198	102	-87	2	-2	-3
5	-200	100	-45	-203	103	-43	3	-3	-2

In Table II, if we compare the real coordinate and the measured coordinate, we can see the maximum error of x and y coordinate is about 4 mm, and the orientation β is about 3°.

The result in here represented the performance of the actual system. We compared them to the results in Table I, and then we can confirm the measurements in our system as good as the manually attained results through the oscilloscope.

6 CONCLUSIONS

In this paper we have presented a method using two ultrasonic transmitters and two receivers to determine the location and orientation in a pre-defined 2D workspace. The approach uses trilateration techniques to locate multiple targets mounted on the platform so that its pose can be determined. Whilst the implemented system is easy to setup and use, the design process is not straightforward as several design issues must be considered, for the system to work effectively. Test results reflect the accuracy and consistency of the system when used within the effective coverage area of the ultrasonic devices.

7 REFERENCES

[1] A. Peattie, A. Korevaar, J. Wilson, B. Sandilands, X. Q. Chen and M. King, "Automated variable resistance system for upper limb rehabilitation," in Australasian Conference on Robotics and Automation (ACRA), Sydney, Australia, 2009

[2] C. K. Wong, K. Jordan and M. King, "Robotic Arm Skate for Stroke Rehabilitation," Industrial Research Limited, IEEE International Conference on Rehabilitation Robotics, Zurich, Switzerland, 2011

[3] A. Bonarini, M. Matteucci and M. Restelli, "Automatic Error Detection and Reduction for an Odometric Sensor based on Two Optical Mice", in IEEE International Conference on Robotics and Automation, Barcelona, Spain, 2005

[4] L. Mudrová, J. Faigl, J. Halgašík, T. Krajník, "Estimation of Mobile Robot Pose from Optical Mouses" Research and Education in Robotics - EUROBOT 2010 Communications in Computer and Information Science, Vol. 156, 2011, pp 93-107

[5] E. M. Gorostiz, J. L. L. Galilea, F. J. M. Meca, D. S. Monzú, F. E. Zapata and L. P. Puerto, "Infrared Sensor System for Mobile-Robot Positioning in Intelligent Spaces," Department of Electronics, University of Alcalá, Alcalá de Henares, Madrid, 28801, Spain. 18 May 2011

[6] A.R. Jimenez and F. Seco, "Ultrasonic localization methods for accurate positioning," Instituto de Automatica industrial. CSIC, Spain, september 2005.

[7] S. A. Zekavat, H. Tong and J. Tan, "A Novel Wireless Local Positioning System for Airport (Indoor) Security," Dept. of Electrical and computer Engineering, Michigan Tech University, Houghton, MI, 49931

[8] H. Kawai, "The piezoelectricity of poly(vinylidene _uoride)," Japanese Journal of Applied Physics, 8:975.6, 1969.

[9] A. S. Fiorillo, "Design and characterization of a PVDF ultrasonic range sensor,". IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control, 39(6):688.692, November 1992.

[10] A. S. Fiorillo, "PVDF ultrasonic sensors for location of small objects," Sensors and Actuators A . Physical, 42(1.3):406.409, 1994.

[11] H. Wang and M. Toda, "Curved PVDF airborne transducer," IEEE Trans. on Ultrasonics,Ferroelectrics, and Frequency Control, 46(6):1375.1386, November 1999.

[12] Measurement Specialties Incorporated: Drive Electronics for PVDF Transmitters USDE-01; Rev C, 30 MAR 2001

[13] O. Cramer, "The variation of the specific heat ratio and the speed of sound in air with temperature, pressure, humidity, and concentration," Journal of the Acoustical Society of America, Vol. 93, No.5, May 1993, pp.2510–2516.

[14] D. A. Bohn, "Environmental Effects on the Speed of Sound," Rane Corporation, Mukilteo, WA 98275 USA, J. Audio Eng. Soc., Vol. 36, No. 4, 1988 April

AWSAM - Miniaturised Wireless Nodes for Structural Health and Environmental Monitoring

Ameer Ivoghlian, Kevin I-Kai Wang, Zoran Salcic

Department of Electrical and Computer Engineering
University of Auckland
Auckland, New Zealand
kevin.wang@auckland.ac.nz

Abstract: **This paper presents two generations of miniaturised wireless sensor nodes targeted at earthquake simulation applications. The nodes are based on an ultra-low power SoC core and incorporate Accelerometers and other MEMS sensors. The nodes are powered by a 3V coin-cell or a 3.7V lithium-ion polymer battery pack. The design emphases miniaturisation and low power consumption throughout. The first generation nodes are used in experiments modelling liquefaction utilising a single 3-axis accelerometer. Improvements implemented on the second generation node include the addition of a gyroscope and magnetometer, allowing acceleration data to be referenced to an absolute orientation.**

Keywords: wireless sensor networks, low power node, structural health monitoring

1 INTRODUCTION

With recent advancements in embedded and wireless technologies, wireless sensor nodes, which are physically small and consume very low power, are becoming reality. The miniaturised sensor nodes can be easily embedded into various applications while communicating with each other through wirelessly. Low power consumption ensures long operating life expectancy without human attention. These features satisfy the requirements for many applications in practice, such as body area networks for personal health monitoring [1] and wireless networks for building and natural environment monitoring [2].

Among various research topics, a worthwhile application of wireless sensor nodes is in civil engineering research. Such research is often focussed on the behaviour of objects or materials under a variety of external forces or conditions. In this project, we have targeted earthquake simulations as a potential application of wireless sensor nodes. This area is topical in New Zealand at present and a fair amount of work is being done in this arena. The ability to do away with awkward cabling and bulky interfaces is certainly attractive.

A typical wireless sensor node consists of one or more microcontroller(s), radio transceiver(s), a power source, and various sensing and actuating elements. Additional data storage can be integrated to support acquisition of longer data sequences. Modern SoC components allow integration of multiple functionalities into a single sensor node.

In this paper, two generations of low power, miniaturised wireless sensor nodes designed and developed at the University of Auckland, targeting earthquake simulation applications are presented. The first generation of Auckland Wireless Sensor and Actuator Mote (AWSAM-1) have been used in physical experiments with promising results. Room for improvement has been identified, and a second generation node (AWSAM-2) has since been designed and fabricated. This paper outlines the hardware for both AWSAM-1 and AWSAM-2. The design decisions are discussed. The software architecture is also touched upon with particular focus on reusability of software. Some experimental results are presented and the performance of the node is evaluated. Further works are also discussed.

2 RELATED WORKS

Wireless sensor nodes are an active and growing area of research. In one category there are nodes designed with a great deal of processing power while still being an embedded solution. The Imote [3] and the Imote2 [4] by Intel are examples of these. The Imote2 is based on the XScale PXA271 and is capable of running a Linux operating system. It provides I^2C, UART, USB, even AC97 audio [4]. Then, there are devices in the ultra-low-power category. The Telos [5], TinyNode [6], and Mica2Dot are examples of these. These types of nodes are typically capable of running a lightweight operating system such as TinyOS, but can be limited by their interfaces. The Mica2Dot, while based on the low power ATmega128L and being no bigger than a US quarter coin, provides only digital IO, ADC, and UART. The functional specifications of the node depend highly on the target application.

Those specifications are a moving target due to changes in technology such as the growth of system-on-a-chip (SoC) solutions, and the constant redefining of the term “ultra-low power”.

3 AWSAM-1

Physical size, power consumption, and flexibility were the primary goals when designing AWSAM-1 [7]. These characteristics are common among related works and are also well suited to the earthquake simulation target application. Having a physically small module allows the node to more closely resemble a natural particle. Furthermore, the node is relatively light by association, and as such has a minimal impact on the test results when attached to a model structure.



Figure 1. AWSAM-1 and NZ 50c coin.

2.1 Hardware

The hardware for the node consists primarily of a SoC microcontroller, which integrates a low power microprocessor and a radio core, a power supply, and digital sensors, including a 3-axis accelerometer and a temperature sensor. There are also some additional peripherals which are discussed later in this section. The overall node dimensions are 28mm by 35mm, which is slightly larger than a New Zealand 50c coin, as shown in Fig. 1. It should be noted that the integrated antenna on the top of the node occupies more than a quarter of the node's size.

2.1.1 Microcontroller and Radio SoC

The heart of AWSAM-1 is the CC430 SoC from Texas Instruments. It has been selected primarily due to its low power capabilities and its relatively small component footprint. Measuring less than 8mm by 8mm in total PCB footprint size, the CC430F5137 includes an ultra-low power MSP430 microcontroller and a CC1101 sub-1GHz radio transceiver. The 16bit MSP430 microcontroller offers a 6-channel 12bit ADC, two 16-bit timers, hardware Real-Time Clock (RTC), SPI, I^2C, and 30 GPIO pins. It operates with just 160µA/MHz in active mode, and 2.0µA in standby mode (RTC running). This feature-set is quite ideal for the target applications.

In order to achieve lower power consumption, the CC430 has a relatively small program flash (32KB) and data memory (4KB) in comparison with other common microcontrollers. While this may limit the implementation of feature rich communication protocol stacks or some of the larger real-time operating systems, it is sufficient for the software that is designed for this ultra-low-power application.

The radio core is based on the CC1101 sub-1GHz IC from Texas Instruments. It can operate from a 2.0V supply and draws as little as 15mA in RX mode (915MHz, 250kBaud) and 17.6mA in TX mode (915MHz, 0dBm).

Using 915MHz on the unlicensed ISM band gives us a few advantages over the 2.4GHz band commonly used in wireless sensor networks. Firstly, the lower frequency allows the radio waves to propagate through and around obstacles more easily. There is also the matter of bandwidth crowding. With the growing popularity of wireless communications operating in the 2.4GHz band, interference is a very real possibility [8]. Finally, at 915MHz we can achieve 2.6 times longer transmission distance than with a 2.4GHz radio at the same power output level (in free space). This is shown using the simplified free-space path loss formula [9]:

$$FSPL = \left(\frac{4\pi d}{\lambda}\right)^2 = \left(\frac{4\pi d f}{c}\right)^2$$

2.1.2 Sensors

Two sensors were incorporated into AWSAM-1. The most important sensor for earthquake simulation experiments is the accelerometer. In addition to this, a temperature sensor is also included in the design to allow access to the node temperature should the application require it.

The MMA8451Q accelerometer from Freescale Semiconductor has been chosen primarily due to its small physical size and low power consumption. It is a 3-axis, 14-bit, capacitive accelerometer capable of sensing up to ±8g at up to 800Hz. It communicates over I^2C which makes interfacing it with the MCU straightforward. It additionally has two programmable interrupt pins, allowing events to be triggered on the MCU. It measures 3mm by 3mm and consumes between 6µA and 165µA depending on the mode and sampling parameters.

The temperature sensor on the node is a TMP102 temperature IC from Texas Instruments. It has a 12-bit resolution and an accuracy of ±0.5°C. It also communicates over I^2C allowing it to share the same bus with the accelerometer.

2.1.3 External Flash

Having the ability to store acquired measurements data on the node before sending it over the radio can be useful on wireless sensor nodes. It can support more efficient radio communication by sending more data at a time. This reduces the overhead associated with radio start-up and shutdown. It can also be useful when data processing is to be done on the node itself. Moreover, in the case where a radio link is temporarily lost, larger data storage allows the node to operate alone for a longer duration.

Taking into account the size of a communication protocol stack and the application sitting on top, there is very little room left for data storage on the microcontroller. That is why external flash memory has been included in the design The AT45DB081D from Atmel is an 8Mb NOR based flash memory with an SPI interface. It supports a minimum of 100,000 program/erase cycles, and data retention of 20 years.

Most importantly, it consumes 7mA of current in active-read mode and just 15µA in deep-power-down mode.

2.1.4 Power Supply

Owing to the wide supply voltage range of the MCU and peripherals, the node can be driven directly from a 3V CR2032 lithium coin-cell battery. The battery is mounted on the back side of the node in an enclosure, resulting in a total node thickness of 10mm. The particular variety of CR2032 battery that is being used with this node has a nominal capacity of 225mAh with a cut-off voltage of 2.0V (the lowest voltage required to operate the radio on the CC430).

2.1.5 Auxiliary

In addition to the core components described above, the node also contains a few additional components that are a necessary part of the design. A 32.768 kHz crystal is incorporated to support the functioning of the hardware RTC. This can be used to append a time-stamp to collected data. A chip antenna measuring 7mm by 2mm has been chosen to reduce the overall size of the radio front-end (a quarter wavelength monopole would be 82mm in length at this frequency). This is implemented alongside a discrete balun, and matching circuitry.

Expansion of the node's hardware is achieved through a 20-pin connector. This connector allows a daughter board to be stacked on top of the rest of the circuit. It provides access to the power supply rail, ADC channels, digital IO pins, SPI, and I^2C for the interfacing of additional sensors.

2.2 Software

The software for the AWSAM-1 is written in C using Code Composer Studio, an Eclipse based development environment from Texas Instruments. In order to keep the node flexible in terms of usage, the focus was to design a software framework that would abstract away direct hardware control and make developing applications for different purposes easy.

2.2.1 Drivers

Writing code to interface with the sensors and external memory can take a large proportion of the programming time. To alleviate the effort when targeting the node to an application, a set of drivers has been written for all the peripherals.

A set of functions has been written for each sensor that provides a common application interface (API) for configuring the sensor and retrieving acquired data. As the sensors share a common I^2C interface, all sensor specific drivers sit on top of an I^2C driver. This modular design approach speeds up driver development for future expansion boards.

2.2.2 Wireless Protocol

The limited size of the internal flash on the CC430 somewhat limits the range of communication protocols that can be used. However, the target applications of these wireless sensor nodes do not require support for complex routing or mesh topologies. A simple star network with an application programming interface will suffice.

The SimpliciTI wireless protocol stack developed by Texas Instruments has been chosen for use with these nodes. It supports star and peer-to-peer network topologies and its API consists of just five commands. An extended subset of the API also provides support for frequency agility and payload encryption. Both of these features contribute to the robustness of the network. Frequency agility allows the nodes to switch channels to avoid crowded or noisy frequencies, while encryption prevents rogue nodes from interfering with legitimate communications. Most importantly, the stack can occupy just 10KB of program memory, leaving plenty of room for hardware drivers and application code.



Figure 2. AWSAM-1 mounted on structure.



Figure 3. Laminar box on shake table.

3 LIQUEFACTION EXPERIMENTS

AWSAM-1 has been put to work in an experiment designed to prove a numerical model for liquefaction in sand [10].

3.1 Test Jig

The test set-up consists of a sand filled laminar box on top of a one dimensional shake table as shown in Fig. 3. All the layers are constrained in one horizontal dimension and are allowed to move freely over each other like a deck of cards in the other horizontal dimension. The aluminium frames pictured in Fig. 3, slide with relative ease due to the use of bearings between each layer. A semi-rigid model structure is placed on top of the sand which represents a building.

Two wireless sensor nodes were used in the test. One was placed within the sand. The other was mounted on top of the model structure. Both nodes were enclosed in an ABS case for protection against dust. In addition to the wireless sensor nodes, draw-wire displacement sensors were mounted against the model structure and the aluminium framing in order to verify the node performance. A node can be seen mounted to the structure in Fig. 2. Models of the September 2010 and February 2011 earthquakes in Christchurch, New Zealand were scaled and used to excite the shake table below the laminar box.

3.2 Results

The raw acceleration data, while useful for earthquake simulation research, is difficult to directly compare to the motion of the structure. In order to verify the accuracy of the sensor node, the acceleration values were transformed to displacement using a simple algorithm based on double integration.



Figure 4. AWSAM-1 comparison to draw wire displacement graph.

The node attached to the structure was chosen for comparison due to its position relative to the displacement sensor. As shown in Fig. 3, the draw wire displacement sensor is attached directly below the wireless sensor node. Their position relative to each other cannot change throughout the test. The node embedded in the sand however, can move relative to the draw wire mounting point and as such, cannot be used for direct comparison. Nonetheless, it is still useful for evaluating the degree to which the node rotates over the course of a test.

As illustrated in Fig. 4, the transformed data from the wireless sensor node closely tracks that of the draw wire displacement sensor.

3.3 Analysis

The increasing inaccuracy that can be seen near the end of the test may be due to the accumulated error that is common in such transformation algorithms [11]. Further analysis of the algorithm is required to verify this.



Figure 5. AWSAM-1 and AWSAM-2 comparison.

4 AWSAM-2

While the earthquake simulation experiments did highlight AWSAM-1's performance and validate its practical usage, it also identified something to be desired.

4.1 Improvements

The density of the sand did well to keep AWSAM-1 in place during the test, and the model structure kept the orientation consistent. This behaviour may change however, when using other materials or test set-ups. If the node begins to rotate or tumble when subjected to force, the software may be unable to compensate for this. The use of a single accelerometer limits the capability of the node to handle angular acceleration. The decision was made to develop a second generation wireless sensor node, AWSAM-2, with two additional sensors; a gyroscope and a magnetometer, in order to gather more precise information, taking into consideration the sensor orientation.

4.2 Hardware

While redesigning to add the two additional sensors, we took the opportunity to make improvements in other areas such as the power supply and MCU. The design of AWSAM-1, its similarities, and differences to AWSAM-2 are outlined in the following sub-sections and illustrated in Fig. 5.

As illustrated, the node size has remained unchanged between revisions. This is made possible through the removal of test-points, and by compressing the circuit layout.

4.2.1 Reused Components

The MMA8451Q accelerometer used in AWSAM-1 has again been included in this design. Experimental results discussed in the previous section proved its usefulness in the target application.

The radio front end on AWSAM-2 remains unchanged. The chip antenna, the discrete balun, and the matching circuitry use the same components and are laid out in the same way. The same applies to the 26MHz crystal driving the radio. The same 20-pin expansion connecter is used, as is the 32.768 kHz RTC crystal. The temperature sensor is also carried across into this design.

4.2.2 Gyroscope

To support the sensing of the node's angular velocity, AWSAM-2 includes the L3G4200D 3-axis gyroscope from STMicroelectronics. It can measure up to ±2000 degrees per second with 16-bit resolution.

Like the accelerometer used in AWSAM-1, the L3G4200D features two interrupt pins. Communication with this sensor is possible via SPI or I^2C. I^2C has been chosen to maintain a consistent interface across as many sensors as possible. The L3G4200D draws 6.1mA of current when actively sensing, and 5µA when powered down.

4.2.3 Magnetometer & Accelerometer Combo

The addition of a gyroscope described above can help the node acquire a sense of orientation. This method however, is based on a dead-reckoning approach and as such, is subject to accumulative errors. An example of this can be seen when extracting displacement data from the accelerometer in AWSAM-1 (Fig. 4). A magnetometer can help, at least in the horizontal plane, by providing an absolute reference of orientation to the magnetic poles. The LSM303DLHC 3-axis magnetometer and accelerometer combo from STMicroelectronics has been chosen for this purpose.

The LSM303DLHC provides 3-axis' of magnetic field sensing at up to ±8.1gauss. In addition to this, 3-axis' of acceleration sensing are available at up to ±16g. All sensor readings are taken with 16-bit resolution, and can be communicated to the system via I^2C. Once again, two interrupt pins are also available. The LSM303DLHC consumes 110µA when active and just 1µA when sleeping. This is quite an improvement over the MMA8451Q ±8g accelerometer which contains only one sensor. Should the accelerometer in the LSM303DLHC match or exceed the performance of the MMA8451Q, it should be possible to remove the latter in future generations in order to achieve a more compact design.

4.2.4 External Flash

In the course of carrying out the earthquake simulation experiments, it was predicted that, for some applications, 8Mb of flash memory may not be enough to store a complete data-set for a single test, especially when sampling at high sample rates. The decision has been made to pre-empt this problem and increase the size of the flash. With a minimum physical node size still a goal, a flash module has been chosen that had more capacity whilst being no larger physically.

AWSAM-2 comes equipped with an AT45DB161D from Atmel. A 16Mb NOR based flash memory with an SPI interface. In fact, this module is almost identical to the one used in AWSAM-1. The package remains the same, as does the pin-out.

4.2.5 Microcontroller and Radio SoC

The variant of CC430 has been changed between generations. A CC430F6137 has been chosen in place of the CC430F5137. The most significant differences for this application are the addition of 14 additional GPIO pins, and the availability of 2 extra ADC channels, which provide additional flexibility in accommodating extra sensing and actuating capabilities for satisfying the requirements of new applications.

4.2.6 Power Supply

The CR2032 coin-cell battery was used with AWSAM-1 because of its small size and the fact that it wouldn't need any power conditioning. While this was convenient to use, it did have some drawbacks. The CR2032 variety that was used has a nominal discharge rate of 0.2mA. During periods of continual data transmission, the node can exceed this level by almost two orders of magnitude. The consequence of this is that the battery can "run flat" far sooner than would be anticipated with 225mAh of charge. It quickly became a common exercise to continually replace the batteries when experimenting over multiple days.

The decision has been made to implement a rechargeable battery solution into AWSAM-2. A lithium-ion polymer battery has been chosen due to its high energy density, its high maximum discharge rate, and its ease of recharging with off the shelf chargers.

Due to the fact that lithium-ion polymer batteries can have voltages up to 4.23V when fully charged, it was necessary to include a voltage regulator between the battery and the rest of the circuit. The TPS79733 linear regulator from Texas Instruments can operate with an ultra-low 1.2µA quiescent current which fits in nicely with the low power design constraint. It also features a power-good output signal. Connecting this digital pin to the MCU will allow the node to power down when the battery voltage drops below 3V. This step protects the lithium-ion polymer battery from over-discharging and also improves the system's robustness by triggering an alert when the node is about to power down.

4.3 Power Consumption

Fig. 6 displays the current consumption profile for AWSAM-2. The radio and external flash memory can be identified as the major current consumers. This will need to be considered during the software design stage. These two components should be kept in an idle state as much as possible.

Component	Current Consumption (µA)			
MSP430 [CC430] (MCU)	2300 *Active*			1.2 *LPM3*
CC1101 [CC430] (RADIO)	17600 *Transmit*	15000 *Receive*	1700 *Idle*	1 *Sleep*
AT45DB161D (FLASH)	12000 *Write*	7000 *Read*	25 *Standby*	15 *Sleep*
MMA8451Q (ACCELEROMETER)	165 *800 Hz*	24 *100 Hz*	6 *1.56 Hz*	1.8 *Sleep*
L3G4200D (GYRO)	6100 *Sampling*	1500 *Standby*		5 *Sleep*
LSM303DLHC (MAGNETOMETER)	110 *Sampling*			1 *Sleep*
TMP102 (TEMPERATURE)	85 *Communicating*	7 *Sampling*		0.5 *Sleep*
TPS79733 (REGULATOR)	1.2 *Supplying*			

Figure 6. Current consumption per component.

5 SENSOR FUSION

Additional sensors in the AWSAM-2 provide a valuable source of raw data. It would be useful however, to merge this raw data together on the node itself and provide the user with a more meaningful data set compared with just raw data. This form of sensor fusion is made possible with the improvements made in AWSAM-2.

5.1 Algorithm

Algorithms based on Euler Angle Transformation are directly applicable to AWSAM-2 and its applications. Using this approach, the position of the node can be determined relative to fixed reference frame, established at the beginning of the experiment [12]. The gyroscope can be used to correct for the nodes rotational motion by adjusting the accelerometer reading accordingly. The resultant data set that is transmitted to the user is in reference to the position of the axis at the beginning of the test and the rotation of the node is transparent.

5.2 Impact on power consumption

One side effect of fusing the sensor data is that the resultant data set can be smaller. With respect to AWSAM-2, the size of the data set can be reduced by half. The consequence of this is that a reduced amount of data needs to be stored in the external flash and transmitted over the radio. As can be seen in Fig. 6, these two components are the major energy consumers, and reducing the time during which they are active can have a major effect on prolonging the battery lifetime.

6 CONCLUSIONS AND FURTHER WORKS

This paper presents the design, implementation, and applications of two generations of miniaturised wireless sensor nodes developed by the Embedded Systems Group at the University of Auckland. The effectiveness of AWSAM-1 has been demonstrated through realistic earthquake simulation experiments carried out in the Department of Civil Engineering. AWSAM-2 has been developed to further improve on performance and capabilities to fulfil the requirements of the target applications.

AWSAM-2 offers valuable features over AWSAM-1 and the result for the researcher is a richer data set from which to extract information from. Future work relating to this work will initially involve developing drivers and applications for AWSAM-2 and eventually putting it to use in Civil Engineering experiments. This is needed to prove the practicality of the use and creation of wireless sensing and actuating networks for structural health monitoring and other applications in building environments.

REFERENCES

[1] T. Dishongh and M. McGrath, Wireless Sensor Networks for Healthcare Applications. Boston: Artech House, 2010.

[2] E. Jafar et al., “Design of miniaturized wireless sensor mote and actuator for building monitoring and control,” Proc. 17th Int. Conf. on Telecommunications 2010, Doha, Qatar, 2010.

[3] L. Nachman, R Kling, R Adler, J Huang, V Hummel, “The Intel® Mote Platform: A Bluetooth*-Based Sensor Network For Industrial Monitoring”, IPSN track on Sensor Platform Tools and Design Methods for Networked Embedded Systems 2005.

[4] L. Nachman, J Huang, J Shahabdeen, R Adler, R Kling, "IMOTE2: Serious Computation at the Edge," Wireless Communications and Mobile Computing Conference, 2008. IWCMC '08. International , vol., no., pp.1118-1123, 6-8 Aug. 2008.

[5] J. Polastre, R. Szewczyk, D. Culler, "Telos: enabling ultra-low power wireless research," Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on , vol., no., pp. 364- 369, 15 April 2005.

[6] H. Dubois-Ferriere *et al*, "TinyNode: a comprehensive platform for wireless sensor network applications," Information Processing in Sensor Networks, 2006. IPSN 2006. The Fifth International Conference on , vol., no., pp.358-365, 19-21 April 2006.

[7] K. Wang *et al*., “Miniaturized wireless sensor node for earthquake monitoring applications,” presented at the 7th Int. Symp. Industrial Embedded Systems, Karlsruhe, Germany, 2012.

[8] W. Guo W. Healy, M. Zhou, "Impacts of 2.4-GHz ISM Band Interference on IEEE 802.15.4 Wireless Sensor Network Reliability in Buildings," Instrumentation and Measurement, IEEE Transactions on , vol.61, no.9, pp.2533-2544, Sept. 2012.

[9] A. Goldsmith, “Path Loss and Shadowing,” in Wireless Communications. New York: Cambridge Univ. Press, 2005, pp. 32.

[10] IPENZ. (2012, October 18). Liquefaction Fact Sheet [Online]. Available: http://www.ipenz.org.nz/ipenz/forms/pdfs/ChChFactSheets-Liquefaction.pdf

[11] M. Lele, J. Gu, "Evaluation of solid state accelerometer sensor for effective position estimation," Intelligent Control and Automation (WCICA), 2011 9th World Congress on , vol., no., pp.959-964, 21-25 June 2011.

[12] E. Akeila, Z. Salcic, A. Swain “Smart Pebble for Monitoring Riverbed Sediment Transport,” IEEE Sensors Journal, Vol. 10, pp.1705-1717, 2010.

How to protect against RSA side-channel attacks using redundant number hardware

I. L. Manuel*, A. Bainbridge-Smith†

* Email: isaac@isaacmanuel.com

† Department of Electrical and Computer Engineering, University of Canterbury
Private Bag 4800, Christchurch, NEW ZEALAND.
Email: andrew.bainbridge-smith@canterbury.ac.nz

Abstract: RSA has been one of the most commonly used encryption systems since it was published in 1978. Its security relies on the intractability of the Integer Factorisation Problem. Cunningly, cryptanalysts have found non-mathematical weaknesses in RSA implementations. Attacks aimed at these vulnerabilities are known as side-channel attacks. They attempt to find the private key by physically monitoring the system. This paper presents a modified repeated squaring algorithm that uses the binary signed digit number system to protect against these attacks. It is intended for use in FPGA implementations of RSA.

Keywords: RSA, Side-channel attacks, Redundant Number Systems, FPGA, Encrpytion

1 Introduction

RSA is a commonly used encryption system. One of its most important uses is in the Diffie-Hellman key exchange [1] of the secure socket layer protocol (SSL). When RSA cryptosystems are implemented they can suffer from unintentional weaknesses. If physical aspects of the system can be monitored by attackers then the system may be vulnerable to side-channel attacks. This report investigates the use of the binary signed digit number system to protect against these attacks and provides a system level design of a RSA cryptosystem for implementation on a field programmable gate array (FPGA).

This report is structured in the following way. Sections 2, 3, and 4 introduce the RSA cryptosystem, side-channel attacks, and redundant number representations respectively. Section 5 presents the design of the modified RSA cryptosystem. Then Section 6 discusses the limitations of the design. The report is concluded in Section 7.

2 The RSA cryptosystem

The RSA cryptosystem was published 1978 by Rivest, Shamir, and Adleman [2]. Its widespread use can be attributed to its asymmetric nature - the key used to encrypt a message is different from the key used to decrypt it. Therefore, knowledge of the encryption key by a third party does not compromise the security of the encrypted data. This is the one-way trapdoor requirement: the key held in secret must not be easily derivable from the key that is made public. We call these the public and private keys. Example 1 demonstrates how the RSA cryptosystem works and is illustrated in Figure 1.

Example 1 (RSA Cryptosystem) *Bob wants to send a secret message to Alice. Bob encrypts the message with Alice's public key to get the ciphertext. The ciphertext is sent over an insecure network to Alice. Alice decrypts the message using her private key. However, the communications channel is insecure and the ciphertext is read by an eavesdropper known as Eve. Eve runs an algorithm to factor Alice's RSA modulus, N, into its constituent primes, p and q. Fortunately, Alice generated a key pair with a large modulus and Eve is growing old as she waits for her algorithm to complete. This example is illustrated in Figure 1.*

The encryption process is based on modular exponentiation[2]. The modulus, N, is the product of two randomly generated large primes, p and q.

$$C \equiv M^e \pmod{N}, \quad \text{(encryption)} \tag{1}$$

$$M \equiv C^d \pmod{N}, \quad \text{(decryption)} \tag{2}$$

where M, C, e and d are the plaintext, ciphertext, encryption exponent, and the decryption exponent respectively.



Figure 1: The RSA Cryptosystem

The RSA public key is comprised of the encryption exponent and the modulus,

$$\text{key}_{public} = \{e,\ N\}, \tag{3}$$

and the private key is comprised of the decryption exponent and the modulus,

$$\text{key}_{private} = \{d,\ N\}. \tag{4}$$

For the decryption algorithm to produce the original message, e and d must satisfy

$$ed \equiv 1 \pmod{\phi(N)}, \tag{5}$$

that is e is the modular multiplicative inverse of d, mod $\phi(N)$.

Euler's Totient, $\phi(n)$, is defined for all positive integers. It is the count of all integers, k, $0 \leq k < n$, where $\gcd(k, n) = 1$: the count of positive integers less than n that are co-prime to n [3]. It is given by,

$$\phi(n) = n \prod_{p_i|n} (1 - \frac{1}{p_i}), \tag{6}$$

where $p|n$ denotes primes, p, that are divisors of n. This can be factorised for the two prime case as used in RSA:

$$\phi(N) = (p-1)(q-1). \tag{7}$$

The security of RSA comes from the intractability of the Integer Factorisation Problem; it is hard to factor an integer into its constituent primes [4]. As long as no polynomial time prime factorisation algorithm exists RSA will be mathematically secure [5] [6]. If an attacker could factor N into p and q, Euler's Totient could be used to inverse the encryption exponent of the public key. This would give the decryption exponent. This provides an attacker with the full knowledge of the private key as N is known through the public key.

Example 2 (Breaking the RSA cryptosystem) *Eve the eavesdropper from Example 1, has managed to factor Alice's RSA modulus into p and q. Eve uses these primes to calculate Euler's Totient, Equation 7. She uses this result to find the modular multiplicative inverse of Alice's encryption exponent, $d \equiv e^{-1} \pmod{\phi(N)}$. This is Alice's decryption exponent. Eve can now decrypt any ciphertext that is encrypted using Alice's public key.*

3 Side-channel attacks

Side-channel attacks target the implementation of a RSA cryptosystem. RSA encryption and decryption are implemented by repeated squaring, Algorithm 1. The timing and power analysis attacks exploit the difference in execution when the current exponent bit b_i is a 1 or a 0.

Algorithm 1 Computes $y = x^b \pmod{N}$

RepeatedSquaring(x, b, N)

```
y := 1
for all b_i in b do
    if b_i = 1 then
        y := y × x (mod N)
    end if
    x := x^2 (mod N)
end for
return y
```

3.1 Timing Analysis

Timing attacks were first proposed by Kocher [7]. At the time attacks against RSA were directed at its mathematical basis. This attack attempts to find the private key by carefully timing the response of the RSA implementation. Suppose that a repeated squaring algorithm is used to implement modular exponentiation for decryption. The computation time will be dependent upon C, the ciphertext, and d, the decryption exponent. The attacker can select known ciphertexts C_j, where $j = 1, 2, \ldots$, and time the RSA system giving times $T(C_j)$. The attacker can then make timing estimations for $C_j^d \bmod N$ with the bit d_i of d, where $d = (d_0, d_1, \ldots, d_i, \ldots)$, set to 0 and 1, and compare these against $T(C_j)$ [8]. Kocher made the observation that the variance will be smaller for the correct guess of d_i.

The attacker, who knows the bits $(d_0,\ d_1,\ \ldots,\ d_{i-1})$, can repeat this attack to find bit d_i until the entire exponent is found. In each repetition a new message is used and a new timing measurement is recorded. These can be used to estimate the probability that the current guess of the exponent is correct. This attack has be refined by Brumley and Boneh [9], and further by Aciicmez, Schindler, and Koc [10].



Figure 2: An example of power analysis measurements [11].

3.2 Power Analysis

The power attack attempts to trace the execution of the modular exponentiation algorithm. Suppose the modular exponentiation algorithm is implemented using Algorithm 1. The attacker monitors the power consumption of the RSA hardware. The power required for the multiplication operation is higher than for the squaring operation [12]. Therefore the sections of higher power consumption correspond to the multiplication operations and the lower power consumption sections correspond to squaring. These are lines 4 and 6 of Algorithm 1 respectively. When there is a multiplication then a squaring operation it implies the bit is a 1. When there is no multiplication the bit must be a 0 [12]. An example of power analysis measurements are shown in Figure 2. It can be seen that an attacker can effectively read the private key exponent from the measurements as d_i will be 1 where there is a multiplication and squaring.

3.3 Preventing Attacks

The side-channel attacks differentiate between a 1 and 0 bit in the exponent by the multiplication step, line 4. A simple prevention is to introduce a dummy multiplication when the current exponent bit is zero. When analysed by power or timing analysis every exponent bit will appear as a 1. However, the introduction of the dummy multiplication reduces the execution speed of the decryption algorithm.

4 Redundant Number Representations

Redundant number systems have some properties that can be utilised to confuse cryptanalysts. This section introduces the Binary Signed Digit number system (BSD) and some properties that can be used to provide security for RSA implementations.

4.1 Binary Signed Digit

BSD is a redundant number representation. It uses the symbol alphabet $\{1, 0, \bar{1}\}$. The symbol $\bar{1}$ represents the value -1.

This representation is similar to binary which uses the alphabet, $\{1, 0\}$. A binary number is represented by a string of n characters, $(b_{n-1}, \ldots, b_2, b_1, b_0)$ with the values $(2^{n-1}, \ldots, 2^2, 2^1, 2^0)$.

Likewise, a number in BSD is represented by a string of n characters, $(r_{n-1}, \ldots, r_2, r_1, r_0)$. The positions have the same magnitude as the corresponding binary representation. Example 3 shows a representations for eleven in both binary and BSD.

Example 3 (BSD Number System) *Both binary and BSD can represent eleven with the string* 1011*,*

$$1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0 = 11_{10}.$$

However, BSD can also represent eleven with the string $110\bar{1}$*,*

$$1 \times 2^3 + 1 \times 2^2 + 0 \times 2^1 + (-1) \times 2^0 = 11_{10}.$$

4.2 Properties of Redundant Number Representations

Redundant number representations allow a number to be written in more than one way. This provides the ability to hide the binary representation of the key from side-channel attacks using random binary signed digit representations.

Example 4 (Multiple BSD Representations) *With* $n = 3$ *BSD characters, the number three can be represented in the following ways:*

$$011 \rightarrow 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0 = 0 + 2 + 1 = 3_{10},$$

$$10\bar{1} \rightarrow 1 \times 2^2 + 0 \times 2^1 + 1 \times 2^0 = 4 + 0 + 1 = 3_{10},$$

$$1\bar{1}1 \rightarrow 1 \times 2^2 + (-1) \times 2^1 + 1 \times 2^0 = 4 - 2 + 1 = 3_{10}.$$

The second useful property of redundant number systems is the ability to represent numbers with a minimal hamming weight. A representation with a minimum Hamming weight has minimal amount of non-zero symbols. Zero is the additive identity,

$$x + 0 = x, \tag{8}$$

and it absorbs multiplication,

$$x \times 0 = 0, \tag{9}$$

therefore zero symbols are useful for speeding up arithmetic operations [13]. The Hamming weight, $H(R)$, of a representation, $R = (r_{n-1}, \ldots, r_2, r_1, r_0)$ is,

$$H(R) = \sum_{i=0}^{n-1} \delta(r_i), \tag{10}$$

where,

$$\delta(r_i) = \begin{cases} 0 & \text{if } r_i = 0 \\ 1 & \text{if } r_i \neq 1. \end{cases} \tag{11}$$

When BSD is used in multiplication, the multiplication will be faster if the specific representation has low Hamming weight. Example 5 provides a discussion using the number fifteen.

Example 5 *Consider multiplication by fifteen. Fifteen can be represented by* 1111. *An alternate representation is* $1000\bar{1}$. *Multiplication using the second representation will be faster as there are only two non-zero bits in the representation as opposed to four in the first. It is faster as a factor of 0 in a multiplication guarantees a product of 0, Equation 9.*

4.3 Finding BSD Representations

BSD Representations of a given number can be generated using the method in Kamp [13]. The recursive algorithm generates a representation tree that starts with the value to be represented at the root node. During each step a new symbol is appended to the representation starting at the least significant symbol and appending towards the most significant symbol. After each step the value of the symbol is removed, leaving the value remaining to be encoded. This is repeated until the value remaining is 0. Figure 3 and Example 6 provide a trace of this algorithm.



Figure 3: A representation tree for the number seven.

Example 6 (Generating the number seven in BSD) *Figure 3 shows a trace for generating all BSD representations of seven. The dotted arrows represent infinite loops.*

This produces the following representations:

$$0\ldots1\bar{1}\ldots00\bar{1},$$

$$0\ldots1\bar{1}\ldots0\bar{1}1, \quad \text{and}$$

$$0\ldots1\bar{1}\ldots11,$$

where an ellipsis represents an arbitrary length sequence of the preceding symbol. In theory these sequences could be infinite, however the hardware implementation will restrict the bit length of the representations.

5 A BINARY SIGNED DIGIT CRYPTOSYSTEM



Figure 4: The Proposed RSA Cryptosystem.

Figure 4 shows the design of the proposed RSA cryptosystem to combat side-channel attacks. It takes encrypted messages from a insecure network like the internet, and decrypts it using a randomly selected BSD representation of the decryption exponent. The decryption step uses a modified repeated squaring algorithm, Algorithm 2. The function $ModularMultiplicativeInverse(x, N)$ returns the modular multiplicative inverse of x, mod N. Representations of the decryption key can be generated using the method in Kamp [13] as presented in Section 4.3. This system will hide information about the private key as each representation of the exponent will result in a different sequence of execution within the repeated squaring algorithm. Therefore, attempts at correlating data by an attacker will not reveal any information about the private key.

The modified repeated squaring algorithm requires the modular multiplicative inverse of the ciphertext, u, as an exponent symbol of $\bar{1}$ implies division. Therefore, when the current exponent symbol, r_i, is $\bar{1}$, the result is multiplied by u. This is useful as it allows the modulus operation to be used after every multiplication. A naïve implementation of the repeated squaring algorithm with a BSD exponent would be to divide when the exponent is a $\bar{1}$ and take the modulus only once before returning the result. The hardware to implement this naïve implementation would require registers twice the length of the modulus, N, to accommodate the growth of the result before the modulus operation. This would slow down the algorithm and require more expensive hardware.

Algorithm 2 Computes $y = x^r \pmod N$,using an exponent represented in Binary Signed Digit.

ModifiedRepeatedSquaring(x, r, N)

```
y := 1
u := ModularMultiplicativeInverse(x, N)
for all r_i in r do
    if r_i = 1 then
        y := y × x (mod N)
    else if r_i = 1̄ then
        y := y × u (mod N)
    end if
    x := x^2 (mod N)
    u := u^2 (mod N)
end for
return y
```

There are two methods to find the modular multiplicative inverse of the ciphertext. These methods are Euler's Theorem and the Extended Euclidean Algorithm [14].

Euler's theorem is the simpler of the two:

$$u = C^{\phi(N)-1} \pmod N, \tag{12}$$

where u is the modular multiplicative inverse of C. This method uses modular exponentiation like the decryption algorithm. $\phi(N)$ can be calculated using Equation 7. The major drawback of this approach is that it could be targeted using the same side-channel attacks that are used on RSA decryption. If an attacker could guess the exponent, $\phi(N) - 1$, then they could calculate $\phi(N)$ by adding one, and use this to get the private key. This is why $\phi(N)$ is calculated by Eve in Example 2.

As there is a critical weakness in the calculation of the modular multiplicative inverse with Euler's theorem, the Extended Euclidean Algorithm is the only feasible approach. Further, the implementation of the Extended Euclidean Algorithm was found to be faster then that of Euler's theorem with large primes.

6 Discussion

The proposed binary signed digit RSA cryptosystem can effectively protect against timing analysis as it will reduce the attacker's ability to make correlations for the timing attack. However, it will only provide limited protection against power analysis. Power analysis will reveal the location of non-zero bits in the representation of the decryption exponent. If there are a significant number of non-zero bits this should remain secure to brute force attacks. Consider a 2048 bit key. If 50% of the bits are non-zero then after power analysis 1024 unknown bits remain. In this particular situation the attack has reduced the security of the key to the equivalent of a 1024 bit binary private key. The worst case will occur when a minimum Hamming weight private key representation is chosen.

The modified repeated squaring algorithm requires the calculation of the modular multiplicative inverse of the ciphertext which must be squared with every loop. Both operations are additional to the original repeated squaring algorithm, Algorithm 1. These operations will reduce the execution speed of decryption. Conversely, when a BSD representation with a minimal Hamming weight is selected, it will speed up the computation as zero dominance in the exponent will result in fewer multiplications.

Another potential drawback is that the modified algorithm will not work for every ciphertext, C. For the modular multiplicative inverse of C to exist, the condition $\gcd(C, N) = 1$ must be satisfied. As N is the product of two primes, p and q, there is a chance that one of these primes will divide the ciphertext. Fortunately having $\gcd(C, N) \neq 1$ is unlikely to occur. Example 7 presents random primes for a 1024 bit key, with the probability that $\gcd(C, N) \neq 1$ of $5.3 \times 10^{-154}\%$.

As $N = pq$, there will be $(p-1)$ multiples q, and $(q-1)$ multiples of p less than N. Therefore the count of positive integers less than N that do not have a modular multiplicative inverse mod N is given by,

$$I_{Err}(N) = (p-1) + (q-1), \tag{13}$$

and more generally,

$$I_{Err}(n) = \sum_{p_i|n} \left(\frac{n}{p_i} - 1\right), \tag{14}$$

where $p_i|n$ indicates the primes p dividing n.

Example 7 *Given the following two primes for a 1024 bit key,*

p =1314413183426951221926094199371466960500662574317200603052950464552780095152369762014990305566325185422006702050378352478552367581915883654773477065606947 7 (155 digits),

q =1228850628609180410826264540765870996280335818631630987120576970337123311585677265823682463109274040305712727192882036398381954429295019558590530369501597 1 (155 digits),

the probability that a ciphertext chosen at random from a uniform distribution does not have a modular multiplicative inverse mod pq, is $5.3 \times 10^{-154}\%$*:*

$$P_{Err}(N) = \frac{(p-1) + (q-1)}{pq} = 5.3 \times 10^{-156} \quad (2 \text{ s.f.})$$

.

As the probability of the modular multiplicative inverse step failing is approximately $10^{-153}\%$ for a 1024 bit key, which is the smallest common key size, this implementation is suitable for practical applications.

7 Conclusions and further work

The proposed binary signed digit RSA cryptosystem can effectively protect against timing analysis but will only provide limited protection against power analysis. Power analysis will reveal the location of non-zero bits in the representation of the decryption exponent.

The modified repeated squaring algorithm requires the calculation of the modular multiplicative inverse of the ciphertext, and this value must also be squared with every loop. BSD representations with a minimal Hamming weight will speed up the computation as the zero dominance will reduce the number of multiplies. Taking the modular multiplicative inverse of the ciphertext will fail if the ciphertext is not coprime to the RSA modulus. Fortunately this will only occur with a percentage of approximately $10^{-153}\%$ with a 1024 bit key which is the smallest commonly used key size.

Overall the proposed design provides a new method for implementing a RSA cryptosystem that natively protects against timing attacks. However, as this implementation requires additional operations it is likely to perform slower than the repeated squaring algorithm with a dummy multiply.

References

[1] W. Diffie and M. E. Hellman, "New directions in crytography," *IEEE Transactions on Information Theory*, vol. 22, no. 6, 1976.

[2] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, February 1978.

[3] V. Shoup, *A Computational Introduction to Number Theory and ALgebra*. Cambridge University Press, 2 ed., 2009.

[4] R. Crandall and C. Pomerance, *Prime Numbers*. Springer Science+Business Media Inc., 2 ed., 2005.

[5] S. Cook, "The p versus np problem," *The Millennium Prize Problems*, pp. 87–104, 2006.

[6] D. M. Bressoud, *Factorization and Primality Testing*. Springer Verlag, 1989.

[7] P. C. Kocher, "Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems," *Advances in Cryptology*, pp. 104–113, 1996.

[8] M. Stamp and R. M. Low, *Applied Cryptanalysis: Breaking Ciphers in the Real World*. John Wiley & Sons Inc., 2007.

[9] D. Brumley and D. Boneh, "Remote timing attacks are practical," in *Proceedings of the 12th Usenix Security Symposium*, 2003.

[10] O. Aciicmez, W. Schindler, and C. K. Koc, "Improving brumley and boneh timing attack on unprotected ssl implementations," in *ACM Conference on Computer and Communications Security*, 2005.

[11] Y. Zhang, Y. Tang, and L. Liu, "A salting rsa cryptosystem against side-channel attack," *Journal of Computational Information Systems*, vol. 8, no. 6, pp. 2559–2565, 2012.

[12] S. Y. Yan, *Cryptanalytic Attacks on RSA*. Springer Science Business Media Inc., 2008.

[13] W. Kamp, *Alternate Number Systems for Optimising DSP Performance in FPGA*. PhD thesis, University of Canterbury, 2010.

[14] H. Beker and F. Piper, *Cipher Systems*. Northwood Publications, 1982.

Comparison of Time and Frequency Domain Methods for Analysis of Bearing Damage

H.J. Lawrence and T.C.A. Molteno

*Department of Physics, University of Otago
Box 56, Dunedin, NEW ZEALAND.
Email: hilaryl@physics.otago.ac.nz

Abstract: This paper compares RMS, kurtosis, crest factor, and demodulation as methods of bearing analysis, using data from a sample of new and damaged bearings. Ten new bearings were run for over 50 hours, and data were used to create a new bearing profile. This was used for comparison with 21 further bearings, where data were taken before and after seeded defects were applied to bearing races or rolling elements. Thresholds created using time domain measures of the new bearing data set can in most cases be used to distinguish a damaged bearing from a new bearing, and indicate whether a bearing has deteriorated when compared with itself earlier. Distinguishing damage type cannot be done in the time domain, and existing frequency domain methods are only successful in distinguishing damage type in some cases, with a higher success rate when damage is more advanced.

Keywords: Bearings, vibrations, time domain, frequency domain, seeded defect

1 Introduction

Rolling element bearings are used widely in industrial settings. Bearing failure can be expensive, both in terms of damaged machinery and lost production time. Predictive maintenance can be done by processing vibration or acoustic data detected at or near the bearing housing, and comparing a bearing's current condition with known values for an undamaged bearing.

The problem with condition monitoring is that bearings are usually in environmentally noisy settings[1]. In addition to this, the signal recorded at or near an undamaged bearing consists of several modes of vibration, from the bearing and surrounding structure. It can be difficult to detect any evidence of damage - when it arises - over and above background vibration [2]. As a bearing sustains more damage, the amplitude of vibrations increase. Measuring the RMS of the vibrations, and comparing them with previous RMS measurements is a reliable predictor of bearing condition [3], but is complicated by the fact that there may be other sources of increased vibrations.

A signal from a bearing consists of low frequency vibrations caused by shaft rotation. Any imbalance - including that from the normal operation of a motor, or an applied load - will create vibrations at harmonics of the shaft frequency [4]. Impacts between parts of the bearing structure, or rolling elements passing a fixed point (such as a stationary sensor) result in periodic impulses at structure resonance frequencies. These ringing events become more pronounced when a bearing is damaged - for example, when rolling elements roll over damage on one of the races [5]. The period (or frequency) of these impulses provide information about damage sources. However almost all damage frequencies are suppressed in Power Spectral Density (PSD) plots [6]. These frequencies can be recovered by demodulating the signal first [2].

This paper compares the performance of RMS, kurtosis and crest factor in the time domain, and demodulation in the frequency domain, as methods of bearing analysis. Analysis is done using data from a sample of new and damaged bearings, and thresholds and decision trees are used to classify bearing condition.

2 Overview of Time Domain Analysis

2.1 RMS

RMS is defined in Equation 1. An increase in RMS is associated with an increase in damage, although an increase in RMS does not always consistently correlate with an increase in damage [4].



Figure 1: Bearing tester.

$$\text{RMS} = \sqrt{\frac{1}{N_s}\sum_{i=1}^{N_s} x_i^2} \tag{1}$$

Where x_i is the amplitude of sample i, and N_s is the number of samples.

2.2 Kurtosis

Kurtosis is the 4th moment of the probability density function (PDF), and is defined in Equation 2. Kurtosis of an undamaged bearing is usually about 3, and kurtosis increases when damage causes the vibrations to become more chaotic. However in some cases kurtosis will return to lower values - even back to three - once the bearing becomes severely damaged [7]. Also, the same bearing, with the same level of damage can display a varying kurtosis[8].

$$K = \frac{\int (x_i - \bar{x})^4 P(x) d(x)}{\sigma^4} \tag{2}$$

2.2.1 Crest Factor

The crest factor (defined in Equation 3) normalises the peak signal amplitude to the RMS voltage, hence an increasing crest factor indicates an increase in the relative size of the largest vibrations. Crest factor trends upward with damage, but can also rise, then fall again [4].

$$\text{Cf} = \frac{\sup(x) - \inf(x)}{\text{RMS}} \tag{3}$$

Additional time domain measures, such as standard deviation and other moments of the PDF, signal mean, square root mean, or peak values normalised to measures other than the RMS are used in automated analysis, including that done using Neural Networks[9].

3 Demodulation and Frequency Domain Analysis

Frequency domain analysis is usually done using demodulated data [1], [10], [11]. The Fourier Transform of raw data can be used to identify resonance bands at higher frequencies, and low frequency data can indicate problems such as structure imbalance. The geometry of the bearing structure allows frequencies of interest to be identified. For example a burr on a rolling element will impact with both of the races, causing resonant ringing of a known period. However, impulses do not necessarily occur every period, and relevant spectral peaks usually do not show up on PSD plots. The High Frequency Resonance Technique demodulates the signal, allowing the defect frequencies to be recovered. Sometimes sidebands are present around defect frequencies at $\pm f_s$ (where f_s is the shaft frequency)[12]. First, the signal is bandpass filtered around the structure resonance frequency, next the signal is demodulated by taking the envelope of the signal. This can be done in Matlab by taking the FFT of the amplitude of the Hilbert Analytic [10]. Ericsson [1] compares demodulation, and periodisation, with the Wavelet Transform as part of a bearing classification process, and found that demodulation provided the lowest misclassification rate.

4 Data Gathering

The bearing tester is run using an Associated Electrical Industries AC motor with a frequency of 24.8 Hz. The motor is attached to a rotating shaft using an isolating dog to reduce vibration from the motor. The shaft is held in place by two 6204.2 single row ball bearings - the "test bearing" is housed furthest from the motor, as shown in Figure 1. A load is applied by applying torque to a bolt, which in turn applies a force to the bearing housing of the FAG 6006.2 load bearing, which is transferred to the test bearing via the shaft. An Agilent HEDS-9140 shaft encoder, and HEDS-5140 code wheel are mounted on the end of the shaft with associated circuitry. The shaft encoder outputs are connected to an FPGA, which uses the signals to count revolutions and part revolutions, and control the turning on and off of the motor, including a safety cutout mechanism.

4.1 New Bearing Dataset

Ten new bearings were each run for 56 hours, with data gathered in 20s blocks once an hour - after more regular data gathering during the first hour. The data gathering was semi-supervised, and not done overnight. To minimise the effect of starting each new day, the bearing tester

was turned off for a short time (5-20 minutes) after the last data were gathered for the day. A high load was applied, so this same load could be used with later damaged bearing experiments to increase the onset of the effects of damage. Bearing lifetime decreases with the cube of the load, and high loads are widely used in a laboratory setting [3],[6],[7],[11] as the natural lifetime of a bearing is in the order of several million revolutions [13]. This high load, along with multiple changes of bearing over a short time, caused some problems.

New bearing data were used to create a new bearing profile, used for thresholding. The load caused problems - such as slippage - when running three of the bearings (bearings N4, N7 and N10). The new bearing dataset included some data from these three bearings, carefully omitting any data from when problems presented in order to make the dataset as large as possible. The three problem bearings were not used when using new bearings data to provide 'typical' or 'control' data (for example, the six new bearings in Figure 4).

4.2 Seeded Defect Experiment Process

Ten test bearings were used to trial damage application and supply other experimental design information. Seven of these were used to trial the application of seeded defects. A die grinder with fine grinding stone was used to apply damage to the inner or outer race or ball. Some data gathering was done, comparing the bearing before and after damage in three cases, and just after damage was applied in the other five. One aim of these test runs was to ensure that damaged bearings could be distinguished from new bearings. Figure 2 shows time domain data for test bearings pre and post damage ('TM Undam' and 'TB Dam'), along with new bearing data at around 17 hours and around 50 hours ('NB Early' and 'NB Later'). All time domain measures tend to be higher - sometimes significantly so - in damaged bearings. Another aim was to determine how long after turn-on data could be gathered. New bearing data indicated that RMS was lower at the start of the day, but seeded defect experiments required that data be gathered early, before damage caused further damage. One undamaged, and three damaged test bearings were used to determine the ideal time to recored data.

Twenty-one bearings were used for the seeded defect experiments. Eighteen bearings had a single source of damage applied to the inner race, outer race or ball (making three sets of six bearings). The final bearings each had a different combination of two sources of damage applied. The experiment schedule was:

1. A small portion of the cage was damaged, in order to allow for access to bearing races and /or balls.

2. The bearing was thoroughly cleaned and re-lubricated with SKF LGMT 2/1 Bearing Grease.



Figure 2: RMS, kurtosis and crest factor for 9 test bearings and 10 new bearings. Bearings T2-T6 had data taken at 5 min intervals after application of seeded defects, T7 had data taken hourly with no damage applied and T8-T10 had data taken before and after the application of seeded defects. New bearing data from around 17 hours and 50 hours (3 hours from each period) are also displayed.

3. The bearing was pressed onto the shaft and run in for 17 hours (overnight).

4. Data were gathered for 60s at 17:25h, 17:30h and 17:35h and the bearing tester was turned off.

5. The shaft was removed, and the bearing was damaged while still on the shaft. Cleaning and re-lubrication were also done.

6. The shaft was returned, and the bearing tester run for 37 minutes with data gathered at 25, 30 and 35 mins.

Monitoring data was also gathered at scheduled intervals - on one occasion this was used as a very damaged bearing could not be run for the whole 37 minutes.

5 Time Domain Analysis of Results

The new bearing data were used to create threshold values for comparison with other bearings. RMS, kurtosis and crest factor were measured for all new bearings, after first decimating the data to a Nyquist frequency (F_{nyq}) of 10 kHz. The Matlab Decimation algorithm default filter is used, which is an 8th Order Chebychev filter, with a cutoff frequency of $(0.8)F_{nyq}$. Each bearing's data from 15-56 hours and the data set as a whole from 15-56 hours were plotted as Boxplots (see Figure 3). The first 15 hours of most bearings produced a-typical data due to bedding



Figure 3: RMS, kurtosis and crest factor data for each new bearing (N1-N10), and for the new bearing dataset as a whole (All NB). Boxes define upper and lower quartiles, whiskers indicate 1.5 x the interquartile range above the upper and below the lower quartile, '+' indicates an outlier.

in. RMS tended to be lower, and kurtosis more variable. Boxplots of data from 0-56 hours indicated the presence of more outliers than that from 15-56 hours, supporting this decision.

Thresholds were calculated using quantiles, and inter quartile range (IQR), which is a measure of the spread of data, and tested using test bearings data (shown in Figure2) to determine whether thresholds were reasonable. Thresholds were created for comparison with new bearings, and for a bearing to be compared with itself. This second set of thresholds were created using the quantiled data for individual new bearings (see Figure 3), and avoiding problem bearings (see Section 4.1). The bearing with the largest IQR for each measure (RMS, kurtosis, and crest factor) provided thresholds for that measure, all based on multiples of the IQR. It must be noted that the process of cutting open bearings and measuring damage has not been done at this stage, as there is some possibility bearings will be re-run. So, while 'undamaged' vs 'damaged' is quite clearly defined, any 'Level of Damage' is subjective at this stage - based on knowledge of what happened during damage application, inspection of the bearing data, and subjective observation of volume and pitch and measurement of the temperature during experiments.

5.1 Analysis of Seeded Defect Data

Results by damage type (see Figure 4) show that RMS, kurtosis and crest factor tend to be higher for damaged bearings, but clearly some bearings show the effects of damage more than others. Al-Ghamd [6] applied various sized seeded defects to bearings, and found that there was



Figure 4: RMS, kurtosis and crest factor for 27 bearings (3 readings each): six undamaged with data taken at around 50 hours, six of each of inner race, outer race and ball damage and 3 with two sources of damage (1:inner & ball; 2:outer & ball; 3:inner & outer), with data taken 25, 30 and 35 minutes after application of damage.

only a small increase in RMS in bearings when less damage was applied, and that kurtosis tended to be higher with damage, but sometimes returned to 3 when damage progressed. Actual damage has not yet been measured in case bearings need to be re-run, but initial results are consistent with our attempts to apply varying levels of damage.

Thresholding results for RMS, kurtosis and crest factor were combined using a weighted average favouring RMS over kurtosis over crest factor, as RMS is the more reliable measure of bearing condition [14]. Initial results indicate that 15 of the 21 bearings tested positive for damage, and a further five usually or always indicated an alarm for early damage. Only one damaged bearing (outer race defect #5 (OD5) in Figure 4) was classified as being in 'good' condition. Sometimes damage causes transient effects, particularly when the ball bearing is damaged. This can be seen in Figure 4 - there tends to be more variation in results for ball damage than other sources. Ball defect #2 (BD2) moved from 'good' to an 'alarm' state, and BD5 and BD6 varied in their levels of damage over the three data. Inner race defect #5 (ID5) also varied, having clear evidence of damage at 25 and 35 minutes, but presenting as 'good' at 30 minutes. This emphasises the need for regular monitoring of bearings especially when nearing the end of their rated lifetimes, or after showing possible signs of damage.

When using the thresholds for comparing a bearing with itself, sixteen bearings showed a clear worsening in con-

dition after damage was applied. A further four bearings (BD1, BD3, OD4 and OD5) showed an alarm condition for possible deterioration. BD1, BD3 and OD4 also went from 'good' to 'possible damage' condition. Only BD2 showed no significant worsening in condition.

Slight damage was done to all 21 bearings' cages before running in and seeding damage (see Section 4.2). This had some effect on bearing condition. Six of the 21 bearings presented with an alarm for possible early damage before seeded defects were applied. All of these bearings showed clear evidence deterioration after damage was applied, with all indicating positive for 'damage' and a 'worse' condition.

These results are only provisional - different weighting of the RMS, kurtosis and crest factor scores will be trialled, as will the effect of altering threshold levels. Scores already include thresholds for increasing levels of damage and increasing deterioration ('d', 'D', 'D+', 'w', 'W', 'W+'), however, these levels cannot be quantified until bearings are opened, photographed and measured.

6 Determining Damage Type

The analysis in Section 5 makes no attempt to measure damage type. Boumahdi [11] uses decision trees to attempt to determine damage type, first using only RMS, kurtosis, crest factor. He then adds thresholds to indicate the presence or absence of peaks at the first and second harmonic of inner race, outer race and ball damage frequencies, using demodulated data (see Section 2). Three versions of these trees were implemented using the Matlab classregtree function. Trees one and two uses the same measures as Boumahdi (time domain only, time domain plus thresholds). Tree three uses time domain measures, and frequency domain peak-to-floor amplitude ratios. These ratios were calculated by measuring the peak amplitudes at the first two harmonics of each damage frequency and the local lower quartile amplitude between each peak.

All three trees identified 83% of damaged bearings as damaged (with damage type not necessarily correct), and 90% of undamaged bearings as undamaged (see Table 1). Trees one and two could correctly identify damage type in less than half of all cases. Tree three identified damage type in 57% of all cases - 11 of the 21 bearings were always classified correctly (17 were always correctly classified as damaged). Misclassification were more likely in less damaged bearings - of the 15 bearings that tested positive for 'damage' using thresholding, ten always gave a correct result for damage type, and a further two did for 2 out of 3 data. Of those identified as damaged, most misclassification were as having ball damage. This is because of the decision tree, after separating undamaged and damaged bearings, then tested for presence of inner then outer race defects, giving remaining bearings a classification of ball damage. A test for presence of ball damage, then a result of 'undetermined' if this also fails would give a better indication of the effectiveness of the classification process. Results make it clear that time domain analysis alone cannot determine bearing damage type, and that demodulated frequency domain data is successful in determining damage type in only a small majority of bearings, and is rarely successful when damage is less advanced. Problems arise as peaks are not always detectable at damage frequencies - even if RMS and other measures indicate damage. Also, the frequency which balls pass over a fixed point (eg a stationary sensor) is the same frequency that they pass over a defect on the stationary outer race - so even new bearings produce impulses at this frequency. The amplitudes are lower with new bearings, but a noisier new bearing can be hard to distinguish from a bearing with slight damage to the outer race. Figure 5 shows two new bearings (N7, N8) and two with inner race damage (ID2, ID4). The N7 and N8 both have a peak at 76 Hz, corresponding to the ball pass frequency, but there is a 10 dB difference in the amplitude of this peak. Bearing ID2 shows a clear indication of inner race damage, with harmonics of the inner race damage frequency of 122 Hz, up to at least 1 kHz. Each peak has sidebands present at $\pm f_s$, although some of the upper sidebands are obscured by noise. Bearings ID2 and ID4 both test positive for significant damage 'D+' and significantly worse 'W+' compared with itself earlier, using thresholding. However, demodulated data from ID4 does not produce evidence of inner race damage.

	Tree 1	Tree 2	Tree 3
All:correct cond+type	64%	66%	72%
Dam:dam type Ided	40%	43%	57%
Undam:IDed as undam	90%	90%	90%
Dam:IDed as dam	83%	83%	83%

Table 1: Results of the three different classification trees, showing the percentage of all bearings correctly classified (row 1); damage type correctly identified (row 2); undamaged bearings correctly classified (row 3); and damaged bearings identified as damaged (row 4).

7 Conclusions and Future Work

Use of existing time domain methods allow the condition of bearings to be accurately measured in most cases. A bearing with a very small amount of damage might not show up as damaged, or may only raise an 'alarm' for possible damage. A full evaluation of thresholding cannot be done until bearings are cut open, and actual damage measured, however results so far are consistent with prior knowledge about the amount of damage that was applied. Determining damage type is more difficult. Existing



Figure 5: Demodulated data from two undamaged bearings and two bearings with damaged inner races. N8 and N9 show different sized peaks at the ball pass frequency of 76 Hz. ID2 shows harmonics of the inner race damage frequency of 122 Hz, with sidebands at $\pm f_s = 24.8$ Hz, ID4 fails to show evidence of inner race damage.

methods, such as taking measures from demodulated data can allow damage type to be correctly identified in over half of cases, with better results in bearings with a greater level of damage. The possibility of using Bayesian Inference, in particular Markov Chains, to perform frequency domain analysis is currently being investigated. The output of a suitable chain would not only indicate the type of defects, and would also give an associated probability with the results. An initial investigation is being done using the demodulated data in the frequency domain, with a simple model of the signal. This is to be replaced with a model of the state of the bearings, taking into account parameters like the motion of an undamaged bearing, defect size and position and friction coefficients pre and post damage.

References

[1] S. Ericsson, N. Grip, E. Johansson, L.-E. Persson, R. Sjöberg, and J.-O. Strömberg, "Towards automatic detection of local bearing defects in rotating machines," *Mechanical Systems and Signal Processing*, vol. 19, no. 3, p. 509, 2005.

[2] P. Stepanic, "A new approach to detection of defects in rolling element bearings based on statistical pattern recognition," *International Journal of Advances Manufacturing Technology*, vol. 45, pp. 91–100, November 2009.

[3] D. Mba, "Acoustic emissions and monitoring bearing health," *Tribology Transactions*, vol. 46, no. 3, pp. 447–451, 2003.

[4] T. Karacay and N. Akturk, "Experimental diagnostics of ball bearings using statistical and spectral methods.," *Tribology International*, vol. 42, no. 6, p. 836, 2009.

[5] H. Li, H. Zheng, and L. Tang, "Wigner-ville distribution based on emd for faults diagnosis of bearing," *Fuzzy Systems and Knowledge Discovery: Lecture Notes in Computer Science*, vol. 4223/2006, pp. 803–812, 2006.

[6] A. Al-Ghamdi and D. Mba, "A comparative experimental study on the use of acoustic emission and vibration analysis for bearing defect identification and estimation of defect size," *Mechanical Systems and Signal Processing*, vol. 20, no. 7, p. 1537, 2006.

[7] M. Elforjani and D. Mba, "Accelerated natural fault diagnosis in slow speed bearings with acoustic emission.," *Engineering Fracture Mechanics*, vol. 77, pp. 112–127, January 2010.

[8] M. Elforjani and D. Mba, "Assessment of natural crack initiation and its propagation in slow speed bearings," *Nondestructive Testing & Evaluation*, vol. 24, pp. 261–275, November 2009.

[9] Z. Xu, J. Xuan, T. Shi, B. Wu, and Y. Hu, "A novel fault diagnosis method of bearing based on improved fuzzy artmap and modified distance discriminant technique.," *Expert Systems with Applications*, vol. 36, pp. 11801–07, November 2009.

[10] Z. Karagulle and H. Kiral, "Simulation and analysis of vibration signals generated by rolling element bearing with defects," *Tribology International*, vol. 36, pp. 667,678, September 2003.

[11] M. Boumahdi, J.-P. Dron, S. Rechak, and O. Cousinard, "On the extraction of rules in the identification of bearing defects in rotating machinery using decision tree," *Expert Systems with Applications*, vol. In Press, Corrected Proof, 2010.

[12] S. Sassi, B. Badri, and M. Thomas, "A numerical model to predict damaged bearing vibrations," *Journal of Vibration and Control*, 2007.

[13] *NSK Rolling Bearings CAT No E1102i.* 1-6-3 Ohsaki, Shinagawa-ku, Tokyo 141-8560, Japan: ©NSK Ltd, 2005.

[14] X. Williams and T. Ribadeneira, "Rolling element bearing diagnostics in run-to-failure lifetime testing," *Mechanical Systems and Signal Processing*, vol. 15, no. 5, pp. 979–993, 2001.

Estimation of Indoor Local Average Power of Wireless Sensor Network

Kun Yu, Stevan Mirko Berber, Meng-Jung Ho

Electrical and Computer Engineering, University of Auckland
38 Prince Street, Auckland
kyu024@aucklanduni.ac.nz, s.berber@auckland.ac.nz, mho023@aucklanduni.ac.nz

Abstract: This paper presents results of a series of received power measurements that were carried out in the Telecommunication Laboratory at the University of Auckland with the aim to investigate the behaviour of local average power in indoor environment. The measurements are done at different distances, antenna positions and space resolutions. Theoretical calculations based on a Rayleigh fading and two-ray reflection model are conducted and compared with the results of measurement in order to determine if the outdoor propagation theory can be applied on indoor measurements. Based on the measurements and calculations the conclusion can be made that the Rayleigh fading and two-ray reflection model can be used to estimate the local average power under certain indoor situations. The obtained results can be used in determining the location of nodes in wireless sensor networks.

Keywords: Indoor Radio-channel Measurements; Local Average Power; Rayleigh Fading; Indoor Wireless Communications; Wireless Sensor Networks

1 INTRODUCTION

A wireless sensor network (WSN) is comprised of a collection of spatially distributed autonomous sensor nodes that work cooperatively and wirelessly to measure certain parameters and forward data to a certain point for further processing [1]. How to determine the location of nodes in a WSN is very important in many applications such as medical applications or stock control systems. There are several ways to calculate the location of a node, such as choke point concepts, grid concepts, long range sensor concepts, angle of arrival, time of arrival, inertial measurements or received signal strength indication (RSSI). RSSI method is based on the signal power level that is received by a sensor node. The distance between two nodes can be estimated based on the relationship between transmitted and received signal power because radio waves propagate according to inverse-square law. However, the simple two-ray propagation model cannot be applied due to the influence of the environment like reflections, refractions and absorption from walls. So a reliable method to measure and calculate the local average power in such kind of environment is necessary.

Various methods were developed to obtain local average power such as median method or mean method. In order to have a better understanding of indoor radio-channel behaviour so a better method to obtain local average power can be chosen or developed. For that purpose, a series of measurements was carried out, the measured data are processed and analyzed and the method to estimate local average power was proposed.

The two-ray ground reflection model is used in the calculation of theoretical indoor signal strength. In spite of that this model considers both the direct path and the ground reflection path; we used this model as a guide for real measurements instead of the less accurate free space model. The results of a two-ray model are plotted alongside with the real measurements to determine their similarity.

2 THEORETICAL STUDY

The local average power of the received signal can be found by taking into account two considerations [2][3]. The first consideration includes a proper determination of the interval (or length) of measurements. The second consideration includes a proper specification of the number of measurements in that interval that will guaranty acceptable level of confidence in the average power estimation.

The received signal can be expressed as a function of time

$$s_{Rx}(t) = r(t)e^{j\psi(t)} \tag{1}$$

where $r(t)$ is the envelope and $\psi(t)$ represents the phase of the received signal. The envelope is a function of time

and the distance from the receiver. As a function of the distance y, it can be expressed as a product of three terms

$$r(y) = m(y)r_0(y) = s(y)m_0(y)r_0(y), \tag{2}$$

Where $s(y)$ is the signal part, $m_0(y)$ represents slowly changing part that has log-normal distribution and $r_0(y)$ represents fast fading part that has Rayleigh distribution

$$f(r_0) = \frac{r_0}{\sigma^2} e^{-\frac{r_0^2}{2\sigma^2}}. \tag{3}$$

We also define $m(d)$ as a product of the first two terms to represent the slowly changing mean value of the received signal. We can estimate this local mean at the distance $d=x$ as

$$\hat{m}(x) = \frac{1}{2L}\int_{x-L}^{x+L} m(y)r_0(y)dy = m(x)\cdot\frac{1}{2L}\int_{x-L}^{x+L} r_0(y)dy, \tag{4}$$

where the interval $2L$ should be chosen in such a way to make this estimated value to be equal to the local mean $m(x)$. The normalized estimator can be represented as a new random variable that is an integral of the Rayleigh random variable expressed in this form

$$\bar{m} = \bar{m}(x) = \frac{\hat{m}(x)}{m(x)} = \frac{1}{2L}\int_{x-L}^{x+L} r_0(y)dy. \tag{5}$$

In our measurements these values form one realization of a discrete time stochastic process with the mean, autocorrelation function and variance calculated as

$$\eta_{\bar{m}} = E\{\frac{1}{2L}\int_{x-L}^{x+L} r(y)dy\} = \sqrt{\frac{\pi}{2}}\sigma, \tag{6}$$

$$\begin{aligned} R_r(\tau) &= E\{\frac{1}{4L^2}\int_{x-L}^{x+L} r(y_1)dy_1 \int_{x-L}^{x+L} r(y_2)dy_2\} \\ &= [\frac{\pi}{4} + (1-\frac{\pi}{4})\cdot J_0^2(\beta\tau)]\cdot 2\sigma^2, \end{aligned} \tag{7}$$

$$\sigma_{\bar{m}}^2 = \frac{\eta_{\bar{m}}^2}{2(2L/\lambda)}\int_0^{2L/\lambda}(1-\frac{x}{2L/\lambda})\cdot J_0^2(2\pi x)dx. \tag{8}$$

The variance value is decreasing when the interval $2L$ of measurements increases tending to zero when L tends to the infinity. The values of variance are not changing significantly between 20λ and 40λ. Therefore, there is no need to use an interval longer than 40λ. We also want to reduce the number of measurements in this interval but preserve the confidence in the local mean value estimation. The number of measurements needed to achieve the confidence that the measured mean is obtained with confidence 90% is found to be at least N = 36 [2][3].

In our measurements we obtained N = 350~400 samples inside the interval of 15λ to gain a higher accuracy. If the mean value is one, ie. $\eta_{\bar{m}}^2 = 1$, we may calculate the variance use $2L$= 15λ as follows

$$\begin{aligned} \sigma_{\bar{m}}^2 &= \frac{\eta_{\bar{m}}^2}{2(15\lambda/\lambda)}\int_0^{15\lambda/\lambda}(1-\frac{x}{15\lambda/\lambda})\cdot J_0^2(2\pi x)dx \\ &= \frac{1}{30}\int_0^{15}(1-\frac{x}{15})\cdot J_0^2(2\pi x)dx \\ &= \sqrt{\frac{1}{30}\int_0^{15}(1-\frac{x}{15})\cdot J_0^2(2\pi x)dx}, \end{aligned} \tag{9}$$

which result in this value of the theoretical variance

$$\sigma_{\bar{m}} = 0.1023. \tag{10}$$

The variance for 10λ is 0.122, for 20λ is 0.09[3][4]. In our measurement a variance of 0.1023 is accurate enough.

3 MEASUREMENTS AND RESULTS

Several basic methods of the received power measurements are presented in this section and the measured results are related to the theoretical model developed in Section 2. This includes measurement with the receiver antenna rotation, measurement without antenna rotation, and measurement with 5cm space (distance) resolution and measurement with 1cm space resolution. Indoor experiments were performed at 2.193 GHz using a signal generator (E4438C) as a transmitter (*Tx*) and a portable receiver PR100 (*Rx*).



Figure 1: Set up of measurements

Two methods were used in our measurement. One is distance dependent (DD) which is point by point measurement along the straight line from transmitter *Tx* to receiver *Rx*. The second way is distance and rotation dependent (DR) measurement. In this measurement, at each point along the line from *Tx* to *Rx* the antenna was rotated and the received power was recorded for one rotation of 360°, as shown in Figure 1. The radius of the circle was 100cm, which is around 15λ, gives a rotating distance of 46 λ.

3.1 Measurement without Rotation

A rail equipped with a step motor was used for these measurements. The measurements were made every 5cm along a direct line and the sample of around 350 meas-

ured values for each measurement point was recorded. The distance from the receiver antenna to transmitter antenna was changed from 110cm to 515cm in steps of 5cm.

As shown on Figure 2 the power decreases on the way with very significant fluctuations. This could be due to the interference in telecommunication lab such as the movement of people, the reflection of ground and walls. With the measurement like this it is very hard to determine the value of the local mean at a certain distance.



Figure 2: Point by point measurement from 110cm to 515cm

In order to eliminate interference another set of measurement was done during a period that there are no people at the telecommunication lab and the sample size was increased to 1000 to average out the fluctuation. The result is show in Figure 3.

Figure 3: Point by point measurement with 1000 sample size

As shown on the graph, the fluctuation of the values is still very big. This confirms that the presence of people in the telecommunication lab is not the main source of the interference, fading and reflection of the walls would be the main source. Further measurement is followed to find a better way to obtain local mean and the suitable model for this situation.

3.2 Measurement with Rotation

The rail is also used in this set of measurements and another step motor was installed on the rail so that the receiver antenna can linearly move alone the rail and rotate at each point for 360°. The graph, showing the measured mean values as a function of the distance, is presented in Figure 4.

As it shows on Figure 4, the fluctuation of the mean values obtained in this way are much smaller than the fluctuation of the mean values obtained without antenna rotation and the decreasing trend is more pronounced. Namely, by doing rotation, the power values measured are averaged in the space around the antenna and fast fluctuation of the received power due to the fading are eliminated.



Figure 4: Measurement with 1meter radius rotation

3.3 DD and DR measured powers comparison

Figure 5 and Figure 6 present the power received as a function of the antenna space location.



Figure 5: DD compare to DR at distance 160cm



Figure 6: DD compare to DR at distance 505cm

Figure 5 and Figure 6 compares the measurements at different distances. Figure 5 shows the measurements at 160cm distance. The change of power of the DR measurement is very big while there is almost no change for DD measurement. The mean value of DR at 160cm is -98.51dB and of DD is -92.06dB. The variance of DR is 40.01dB and of DD is 0.21dB.

When the distance is increased to 505cm, as shown in Figure 6, the change of power of the DR become smaller and the values for DD show some changes. The mean value of DR at 505cm is -103.21dB and of DD is -102.67dB. The variance of DR is 13.97dB and of DD is 3.88dB. This result shows that when the measurement point is closer to the transmitter the difference of DR and DD is larger and also that the DR's value is less stable.

It would be interesting to estimate the mean value of the signal received and present it as a function of the distance. This can be done by calculations or by filtering as it was done in reference [5]. This can be a subject of further investigation.



Figure 7: Variance of DD and DR

Figure 7 shows the variances of DD and DR values measured from the distance of 110cm to 515cm. As shown on the graph when the distance increases there is a trend that the variance of DR decreases while the variance of DD is increases. Namely, in DD measurements, when the receiver moves further away from the transmitter, due to the fading inside the channel, the signal fluctuations are increasing and the variance of the measured power increases. On the other hand, in the DR measurement, the measured power values are taken in the space around the point of measurements and the fluctuations of the power are slightly decreasing with a decreasing trend when the distance from the transmitter increases.

3.4 Measurement with 1cm and 5cm resolution

In order to find a better method to obtain the local mean values we have to have a better understanding of the behavior of the signal strength in indoor environment, so another two sets of measurements were made in the telecommunication lab. This time the distance from the transmitter antenna to the receiver antenna was changing from 50cm to 750cm and different antenna heights were used as well. The obtained results of measurements are plotted alongside the expected power values obtained by the theoretical two-ray model. The power of the two-ray reflection model was calculated using this equation [6][7]

$$P_R = \left(\frac{\lambda}{4\pi d}\right)^2 4\sin^2\left(\frac{2\pi h_r h_t}{\lambda d}\right) G_R G_T P_T. \qquad (11)$$

The first measurement is point by point (DD measurement) with 5cm distance between each measurement. The antenna height is 1.2 meter. After acquiring the results of measurements we calculated the power according two-ray reflection model in an ideal situation and for the same frequency. The comparison of two-ray model and our measurements is show on Figure 8.



Figure 8: 2.193GHz 1.5 meter height with 5cm resolution compares to two-ray model

As shown on Figure 8, there are significant similarity between the two-ray model and the measurement results. Especially, the minima positions coincide and the shapes of the graphs overlap at the reasonable level. The existing differences in the values measured and theoretically expected could be due to the fading existence inside the room. The pronounced difference between theory and the measurement at the distance greater than 600cm is due to the reflections from the wall at the other end of the lab. Namely, the receiver antenna was getting closer and closer to the wall after 600cm distance from the transmitter. In order to further minimize the effect of other reflections and acquire a smother graphs, another set of measurements was made with antenna height of 0.5 meter and 1cm distance between each point of measurements. The obtained results are also compared with the two-ray model as shown in Figure 9 below.



Figure 9: 2.913GHz 0.5 meter height with 1cm resolution compares to two-ray model

Figure 9 shows a much clearer change of the power and the waveform. The waveform follows the trend of two-ray model. The positions of minima for two-ray model and the minima obtain by measures are nearly overlapping. The power at the longest distances increased which could be due to reflections of the wall at the receiver end of the laboratory.

4 CONCLUSIONS

Two methods for local mean of the received signal power measurements and estimation are used and compared in our study. The measured values obtained by the distance dependent (DD) method had higher variance when the receiver was moving away from the transmitter due to the stronger influence of fading. The variance of the measured values obtained by the distance and rotation dependent (DR) method had slightly larger values when the distance was smaller. Overall, the DR measurement could give us a better result of local mean value due to the fact it measures the signal strength of an area instead of a single point. This method would also be implemented I wireless sensor network, which could be a future topic of the research in this area.

The measurements with 1cm and 5cm resolution and different antenna heights showed a good match with the theoretical two-ray model, especially the positions of the minima values coincide quite well. This shows that, even though the transmitter and receiver are in indoor environment with multiple reflections, the two-ray reflection model can still be used to see the basic behaviour of the received signal power. In future research we could add noise components or other variables to improve this model. The measurement methods are based on the theoretical study from Section 2, which was applied for outdoor measurements. Our study shows that the same theory, with possible improvements, can be used for indoor measurements.

5 REFERENCES

[1] Terplan, K., Morreale, P. *The telecommunications handbook*. CRC Press. Boca Raton, Fla, 2000.

[2] William C. Y. Lee, "Estimate of local average power of a mobile radio signal" *IEEE Trans. Vehic. Technol.* Vol34, Issue 1, 1985, pp. 22-27.

[3] William C. Y. Lee, Y. S. Yeh "On the estimation of the second-order statistics of log normal fading in mobile radio environment" *IEEE Transactions on Communications*, Vol22, Issue 6, Jun 1974, pp. 869-873.

[4] William C. Y. Lee, *Mobile Communications Engineering*. New York: McGraw-Hill, 1982, pp. 105.

[5] Tao Jiang, Nicholas D. Sidiropoulos, Georgios B. Giannakis, "Kalman filtering for power estimation in mobile communications" *IEEE Trans. On Wireless Communications*, Vol2, Issue 1, Jan 2003, pp. 151-161.

[6] Martin J. Feuerstein, Kenneth L. Blackard, Theodore S. Rappaport, Scott Y. Seidel, Howard H. Xia, "Path Loss, Delay Spread, and Outage Models as Functions of Antenna Height for Microcellular System Design", *IEEE Trans. Vehic. Technol.* Vol43, Issue 3, Aug 1994, pp. 487-498.

[7] Tri T. Ha, *Theory and Design of Digital Communication Systems.* Naval Postgraduate School, Monterey, California, 2010, pp. 197.

Compressively Sensed Hyperspectral Image Recovery using Total Variation Minimisation by Approximation

Duncan Eason*, William Lee†, Mark Andrews‡

Department of Electrical and Computer Engineering, University of Auckland
Private Bag 92019, Auckland Mail Centre, Auckland 1142, NEW ZEALAND.
Email: *duncan.eason@auckland.ac.nz, †william.lee@auckland.ac.nz, ‡m.andrews@auckland.ac.nz

Abstract: Compressive sensing (CS) theory states that any signal that is sparse in a known basis may be recovered from a small set of linear combinations of the signal. Capturing data a fraction of the size of the original signal may be beneficial when applied to hyperspectral imaging (HSI), an imaging technique that introduces spectral content in order to classify materials in a scene. Recently, total variation (TV) minimisation algorithms have achieved success in recovering images captured using CS. In this paper we present a novel implementation of the TV minimisation algorithm that includes a differentiable approximation of the TV norm. This new method compares favourably with other TV minimisation algorithms, and we show how it can be extended from monochromatic to hyperspectral CS image recovery.

Keywords: Compressive sensing, hyperspectral imaging, total variation minimisation.

1 INTRODUCTION

Hyperspectral imaging (HSI) is an extension of the more familiar monochrome or RGB imaging in that each pixel contains a densely sampled portion of the optical spectrum. Many practical problems in image enhancement, object recognition and target detection have resisted solution with monochrome or RGB image sources, but there is great promise that hyperspectral (HS) data will provide the richness of information content necessary to solve problems in these areas [1]. One of the major issues with HS processing is that the images are extremely large since every pixel is no longer a single number but potentially several hundred samples.

Over the past few years, there has been an increasing interest in the study of compressive sensing (CS), a new acquisition paradigm that captures linear combinations of a signal, where the number captured is related to the *sparsity* of the original signal's transformed representation [2]. Reconstruction of the original signal relies on promoting this sparsity, and promising results have been achieved in CS image recovery by minimising the total variation (TV) of the solution [3].

In this paper, we outline HSI processing and present the theory for Compressively Sensed Hyperspectral Imaging (CS-HSI) with focus on TV minimisation. We propose our own reconstruction algorithm and a method for handling HS data, with recovery results similar to the pioneering ℓ_1-Magic [4] and the state-of-the-art TVAL3 [5], as well as improved computational times.

2 BACKGROUND

2.1 *Hyperspectral Imaging*

A typical HS image pixel typically contains between 30 and 200 spectral bands (but occasionally many more [6]) covering the visible and near-infrared parts of the spectrum [1]. Figure 1 shows a common form of representation, with the x- and y-directions representing spatial information and the λ-direction representing the various spectral bands, with each element called a *voxel*.

HSI exploits that fact that all materials reflect, absorb, and emit electromagnetic radiation in distinctive patterns related to their molecular composition. For non-



(a) (b)

*Figure 1: (a) 'Peppers' test image—monochromatic. (b) 'Peppers' test image—*HS *representation.*

fluorescent material, each HSI pixel contains a *spectral signature* with contributions from the illumination spectrum, the intrinsic reflectance function of the material and the spectral response of the detector and associated optics. In principle, matching the measured reflectance function with known spectral signatures allows materials in the scene to be automatically determined [1].

2.2 Conventional Sampling

The Shannon-Nyquist sampling theorem has achieved almost dogmatic status in the signal processing community. It states that if all information in a signal is to be captured, sampling must be done at a rate twice the bandwidth. Due to the relationship between the number of measurements and signal bandwidth, often a prohibitively large number of elements is captured and compression is required [2].

2.2.1 Sparsity and Transform Coding

Consider an image[1] $\boldsymbol{x} \in \mathbb{R}^n$. Now suppose there exists a basis Ψ that produces a representation $\boldsymbol{s} \in \mathbb{R}^n$. This can be written as $\boldsymbol{s} = \Psi^{-1}\boldsymbol{x}$, and if $\boldsymbol{s}$ can be adequately represented by $K \ll n$ non-zero entries, $\boldsymbol{s}$ is said to be *K-sparse* [2]. Sparsity is defined as the number of non-zero elements in a vector, and expresses the notion of "information rate" or complexity of a signal [7]. A signal with a sparse representation (such as Figure 2) may be compressed using *transform coding*—transforming a signal into a sparse basis representation then discarding all but the K largest entries [2].

Techniques similar to this have been applied to compression of HSI data, but have met little success. Clearly this process is inefficient; the *whole signal* must be acquired *and* transformed, before most of it is thrown away anyway [2]. Can we not just *directly* measure a compressed version?



(a) (b)

Figure 2: Transforming to a sparse representation. (a) Original 'cameraman' test image. (b) Wavelet representation. Warm colours—high values; cold colours—low values. Many elements may be discarded prior to storage.

[1] For 2D images, we use $\boldsymbol{x} \in \mathbb{R}^n$ to denote a 1D ordering of n pixels.

2.3 Compressively Sensed Hyperspectral Imaging

Potential to alleviate HSI difficulties may lie in compressive sensing (CS), an exciting new sampling paradigm that captures a small[2] number of linear combinations of the entire signal. The number of measurements required is related to the sparsity of the original signal's sparse representation, i.e., the *complexity* of the original signal, rather than the bandwidth. The ground breaking work by Donoho [8] and Candès *et al*. [9] showed that from this small set of linear combinations, a compressible signal can be reconstructed *exactly*.

CS-HSI is advantageous in a number of ways: firstly, there is no need to ever store the full-resolution signal, decreasing the up-front memory demands of the acquisition device. Secondly, with no need for extra compression, computational demands on the acquisition device may be decreased, in turn decreasing power demands and potentially increasing battery life on portable devices. Additionally, because less data is captured, acquisition times for many imaging applications may be shortened.

Each CS-HSI system can be broadly split into two distinct tasks: *acquisition*, acquiring the compressed data; and *recovery*, reconstructing the original HS image. Mathematically, we can view each measurement (a linear combination of each voxel) as the inner product of the original HS image and a *collection vector* φ_j. Each measurement becomes

$$y_j = \langle \boldsymbol{x}, \varphi_j \rangle, \qquad j = 1, \dots, M \tag{1}$$

where[3] $\boldsymbol{x} \in \mathbb{R}^N$ and $M < N$. Alternatively, we can view the acquisition of all M measurement elements as a single matrix operation. Arranging φ_j^T into rows as a matrix $\Phi \in \mathbb{R}^{M \times N}$, acquisition can be viewed as

$$\boldsymbol{y} = \Phi \boldsymbol{x} \tag{2}$$

where $\boldsymbol{y} \in \mathbb{R}^M$. To ensure that HS images are mapped uniquely to the measurement space and vice versa, a condition called the *Restricted Isometry Property* (RIP) must be satisfied [2]. Developed in 2004, it has set the basis for all CS theory [10].

Recovery of $\boldsymbol{x}$ from $M < N$ elements involves solving an underdetermined system with infinitely many solutions [7]. While this seems problematic, the key lies in the original HS image having a sparse representation [2]. Recall $\boldsymbol{s} = \Psi^{-1}\boldsymbol{x}$, where Ψ is a sparsifying basis. To

[2] Small in this case refers to some fraction of the number of elements in the original signal.

[3] For 3D HS images, we use $\boldsymbol{x} \in \mathbb{R}^N$ to denote a 1D ordering of $N = n \times WL$ voxels, separated by wavelength, where WL is the number of wavelengths.

recover $\boldsymbol{x}$, we find the HS image consistent with the observations that produces the *sparsest* $\boldsymbol{s}$ [7]. This is called *sparsity promotion* and underlies all CS recovery. The classic inverse-problem method of ℓ_2 minimisation is inappropriate for this situation because it almost never produces a sparse solution [11]. Instead, provided Φ satisfies the RIP [12], a sparse solution can be found by solving

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \|\Psi^{-1}\boldsymbol{x}\|_1 \quad \text{s.t.} \quad \boldsymbol{y} = \Phi\boldsymbol{x} \tag{3}$$

where $\|\cdot\|_1$ denotes the ℓ_1 norm [11]. In the noisy case, when $\boldsymbol{y} = \Phi\boldsymbol{x} + \boldsymbol{n}$ and $\|\boldsymbol{n}\|_2 = \varepsilon$, the equality constraint can be relaxed[4] and a sparse solution found by solving

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \|\Psi^{-1}\boldsymbol{x}\|_1 \quad \text{s.t.} \quad \|\boldsymbol{y} - \Phi\boldsymbol{x}\|_2^2 \leq \varepsilon^2 \tag{4}$$

2.3.1 CS-HSI Data Acquisition

Despite recent advances in CS theory, large steps are still required to develop practical imaging systems. In particular, hardware that is capable of capturing spatial combinations of an image as well as distinct spectral information is still very much in the developmental stages. A *monochromatic* solution, dubbed the Single Pixel Camera (SPC), was developed by Rice University in 2006 and is shown in Figure 3 [13].

The SPC may be easily extended to a Hyperspectral SPC (HS-SPC) by replacing the photodiode with a spectrometer [2]. Light at each wavelength experiences the same pattern of reflections, and acquisition of an HS image with $\lambda = 1, \ldots, WL$ may be viewed as

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_1 \\ \boldsymbol{y}_2 \\ \vdots \\ \boldsymbol{y}_{WL} \end{bmatrix} = \begin{bmatrix} \Phi & 0 & \cdots & 0 \\ 0 & \Phi & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \Phi \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_{WL} \end{bmatrix} = \Theta\boldsymbol{x} \tag{5}$$

where $\boldsymbol{x}_i$ and $\boldsymbol{y}_i$ correspond to the ith wavelength of the original image and the measurements respectively [2].



Figure 3: SPC Schematic. Light is reflected by a Digital Micromirror Device with mirror orientations modulated by φ_j. Each mirror pattern directs light at a single photodiode, corresponding to one measurement y_j.

[4] $\boldsymbol{y} \neq \Phi\boldsymbol{x}$, so the equality constraint in (3) is no longer valid.

2.3.2 CS-HSI Recovery

To recover HS images from data acquired using the HS-SPC outlined in 2.3.1, (4) may be modified using (5) to provide a new optimisation problem

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \|\Psi^{-1}\boldsymbol{x}\|_1 \quad \text{s.t.} \quad \|\boldsymbol{y} - \Theta\boldsymbol{x}\|_2^2 \leq \varepsilon^2 \tag{6}$$

With Ψ^{-1} the discrete Daubechies-8 orthogonal wavelet transform, solving (6) produces numerically promising reconstruction results [12]. However, recent research has shown that *TV minimisation* is more successful in preserving edges than (6), essential in characterising HS images [3].

Recall that $\boldsymbol{x}$ is a 1D ordering of an HS image; ignoring voxels that would be in the bottom row or right hand column of any wavelength plane in the standard 3D HS representation, the TV norm for an HS image is

$$\begin{aligned} \|\boldsymbol{x}\|_{TV} &= \Sigma_i(|x_{i+1} - x_i| + |x_{i+h} - x_i|) \\ &= \|D_v\boldsymbol{x}\|_1 + \|D_s\boldsymbol{x}\|_1 \end{aligned} \tag{7}$$

where h is the height of the image, and D_v and D_s are matrices that produce vertical and horizontal variation respectively [12]. Examining (7) it should be apparent that the TV norm can be regarded as the ℓ_1 norm of the *spatial gradient* [11]. Solving

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \|\boldsymbol{x}\|_{TV} \quad \text{s.t.} \quad \|\boldsymbol{y} - \Theta\boldsymbol{x}\|_2^2 \leq \varepsilon^2 \tag{8}$$

finds the HS image with the sparsest spatial gradient, or *least variation* consistent with the observations. This is expedient as natural images tend to have sparse gradients [12]. Figure 4 compares the ℓ_2- and TV-minimisation reconstructions of a single wavelength 256×256 *Shepp-Logan Phantom* from 5481 CS measurements. TV minimisation recovers the image *perfectly* [9].

While this paper concerns pure TV minimisation, it is possible to solve a weighted combination of wavelet and TV minimisation, with algorithms of this nature currently being developed [14].

3 PROPOSED APPORACHED

The introduction of the wavelength dimension has greatly increased the dimension of the optimisation



(a) (b) (c)

Figure 4: TV minimisation recovery example. (a) Original image. (b) Recovered using ℓ_2 minimisation. (c) Recovered using TV minimisation.

problem and made speed and simplicity crucial in designing our algorithm. We have developed a gradient-descent based optimisation algorithm for solving (8) that is both fast and simple. Note that gradient-descent methods require a differentiable objective function, whereas the TV norm is clearly non-differentiable. We avoid this technical restriction by formulating a differentiable approximation of the TV norm.

3.1 Total Variation Approximation

We have developed a differentiable approximation of the TV norm by adding in a small scalar $\mu = 10^{-15}$ (an arbitrary small number), shown below

$$\|\boldsymbol{x}\|_{TVa} = \Sigma_i \left(\sqrt{(D_v \boldsymbol{x})_i^2 + \mu} + \sqrt{(D_s \boldsymbol{x})_i^2 + \mu} \right) \quad (9)$$

which is differentiable with a gradient

$$\nabla \|\boldsymbol{x}\|_{TVa} = D_v^T \Gamma_v D_v \boldsymbol{x} + D_s^T \Gamma_s D_s \boldsymbol{x} \quad (10)$$

where $\Gamma_\bullet$ – with $\bullet$ representing v or s – is diagonal with

$$(\Gamma_\bullet)_{ii} = \frac{1}{\sqrt{(D_\bullet \boldsymbol{x})_i^2 + \mu}} \quad (11)$$

3.2 TV Minimisation by Approximation

Utilising our differentiable TV approximation, we have developed an iterative gradient-descent based algorithm that finds the HS image with the minimum TV data by solving

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \|\boldsymbol{x}\|_{TVa} \ \text{ s.t. } \ \|\boldsymbol{y} - \Theta \boldsymbol{x}\|_2^2 \le \varepsilon^2 \quad (12)$$

The objective function was modified to include the residual[5] (with weight β on $\|\boldsymbol{x}\|_{TVa}$) giving

$$f(\boldsymbol{x}) = \tfrac{1}{2}\|\boldsymbol{y} - \Theta \boldsymbol{x}\|_2^2 + \beta \|\boldsymbol{x}\|_{TVa} \quad (13)$$

Clearly (13) is differentiable, and can be minimised by gradient descent by applying the recursive formula

$$\hat{\boldsymbol{x}}^{(k)} = \hat{\boldsymbol{x}}^{(k-1)} - \tfrac{1}{\alpha^{(k-1)}} \nabla f\left(\hat{\boldsymbol{x}}^{(k-1)}\right) \quad (14)$$

with $\alpha^{(k-1)}$ chosen at each step using the Barzilai-Borwein method [15]. After calculating a starting point using the pseudo-inverse $\Theta^\dagger$, steps can be computed and the process repeated until a specific stopping criterion is satisfied, providing the estimated image $\hat{\boldsymbol{x}}$.

Algorithm *TV Minimisation by Approximation (TVMA)*

1. initialise counter, $k \leftarrow 0$
2. calculate initial guess, $\hat{\boldsymbol{x}}^{(0)} \leftarrow \Theta^\dagger \boldsymbol{y}$
3. **repeat**
4. $\quad k \leftarrow k + 1$
5. $\quad \hat{\boldsymbol{x}}^{(k)} \leftarrow \hat{\boldsymbol{x}}^{(k-1)} - \left(1/\alpha^{(k-1)}\right) \nabla f\left(\hat{\boldsymbol{x}}^{(k-1)}\right)$
6. **until** stopping criterion is satisfied

3.3 Stopping Criterion

Experimental results have shown that determining an appropriate stopping criterion for this type of algorithm is difficult, primarily due to the iterative fluctuations in the solution caused by following two different gradients. A simple approach is to examine the relative change in the TV norm. For a small tolerance, tol, we terminate if

$$\frac{\left\|\hat{\boldsymbol{x}}^{(k)} - \hat{\boldsymbol{x}}^{(k-1)}\right\|_{TV}}{\|\hat{\boldsymbol{x}}^{(k-1)}\|_{TV}} \le \text{tol} \quad (15)$$

3.4 Setting β

It should be obvious that β is related to the power in the noise ε; a high β puts less weight on minimising the residual accounting for large noise, and a small β puts less weight on minimising the TV, tightening the residual. However, the literature provides no definitive mathematical relationship.

Our approach involves choosing an arbitrarily high[6] $\beta^{(0)}$, finding the optimal HS image using TVMA, then decreasing β and, starting from the previous solution, finding the optimal solution for the new β, then repeating until the residual inequality in (12) is satisfied. This new algorithm is outlined below.

Algorithm *TVMA with decrementing β (TVMAβ)*

1. initialise counter, $p \leftarrow 0$
2. choose $p_{\max}$, $\beta^{(0)}$, and $\eta \in (0,1)$
3. calculate initial guess, $\hat{\boldsymbol{x}}^{(0)} \leftarrow \Theta^\dagger \boldsymbol{y}$
4. **repeat**
5. $\quad p \leftarrow p + 1$
6. $\quad \hat{\boldsymbol{x}}^{(p)} \leftarrow \text{TVMA}\left(\hat{\boldsymbol{x}}^{(p-1)}, \beta^{(p-1)}\right)$
7. $\quad \beta^{(p)} \leftarrow \eta \beta^{(p-1)}$
8. **until** $\left\|\boldsymbol{y} - \Theta \hat{\boldsymbol{x}}^{(p)}\right\|_2^2 \le \varepsilon^2$ or $p > p_{\max}$

In line 6, $\text{TVMA}\left(\hat{\boldsymbol{x}}^{(p-1)}, \beta^{(p-1)}\right)$ denotes a run of TVMA with $\hat{\boldsymbol{x}}^{(0)}$ and β replaced by $\hat{\boldsymbol{x}}^{(p-1)}$ and $\beta^{(p-1)}$ respectively. The large initial β creates substantial steps in the $\nabla \|\boldsymbol{x}\|_{TVa}$ direction, particularly important when $\hat{\boldsymbol{x}}^{(0)}$ has high a TV. Successively decrementing β puts more weight on reducing the residual, ensuring the estimated image is consistent with the observations.

3.5 HS data handling

Efficient implementations of CS-HSI methods lead to their own challenges given the dimensions of typical problems. For example, to acquire an HS image with

[5] The residual is the difference between the recovered and observed observations, i.e., $\|\boldsymbol{y} - \Theta \boldsymbol{x}\|_2$.

[6] Producing a solution with a violated residual inequality.

256×256 pixels and 256 bands with a sampling ratio of $M/N = 0.5$ requires $\Theta \in \mathbb{R}^{2^{23} \times 2^{24}}$, which is clearly impractical to store and manipulate. Judicious use of algebraic rearrangements can help minimise storage demands and computation time. The function below computes the product $\Theta \boldsymbol{x}$ without explicitly assembling the matrix Θ (only *one* much smaller Φ) and is crucial to making CS-HSI recovery possible.

Function $\Theta_{HS}(\boldsymbol{x})$

1. Rearrange the input (consisting of n pixels and WL bands)

$$\begin{bmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_{WL} \end{bmatrix} \in \mathbb{R}^{n \cdot WL \times 1}$$

as

$$[\boldsymbol{x}_1 \quad \boldsymbol{x}_2 \quad \cdots \quad \boldsymbol{x}_{WL}] \in \mathbb{R}^{n \times WL}$$

2. $[\boldsymbol{y}_1 \quad \boldsymbol{y}_2 \quad \cdots \quad \boldsymbol{y}_{WL}] \leftarrow \Phi[\boldsymbol{x}_1 \quad \boldsymbol{x}_2 \quad \cdots \quad \boldsymbol{x}_{WL}]$

3. Rearrange the output according to

$$\begin{bmatrix} \boldsymbol{y}_1 \\ \boldsymbol{y}_2 \\ \vdots \\ \boldsymbol{y}_{WL} \end{bmatrix} \leftarrow [\boldsymbol{y}_1 \quad \boldsymbol{y}_2 \quad \cdots \quad \boldsymbol{y}_{WL}]$$

This function is also highly parallelisable, potentially allowing multiple wavelengths to be recovered simultaneously. The same concept can be extended to Θ^T (the transpose) and $\Theta^\dagger$ (the pseudo-inverse).

4 SIMULATION RESULTS

The performance of our proposed TVMAβ algorithm has been evaluated and compared to the following reconstruction algorithms: (1) min-TV with quadratic constraints from the ℓ_1-Magic package [4]; (2) anisotropic TV/L2 using TVAL3 [5]. ℓ_1-Magic was an early pioneer in TV minimisation solvers, and TVAL3 is one of the current state-of-the-art algorithms for CS image recovery. The HS-SPC architecture has been simulated with Φ equal to the Scrambled Block Hadamard Ensemble proposed in [16], and the measurements have been corrupted with white Gaussian noise with an acquisition SNR of ~40 dB. Experiments have been run using three simulated $256 \times 256 \times 256$ ($N = 2^{24}$) HS images: (1) the *Shepp-Logan Phantom* with simulated spectral data[7]; (2) *'Lena'* with simulated spectral data; (3) *Materials*—an HS image with spatially simulated material arrangements of real spectral data obtained from [6]. Single wavelengths of each test image and the spectra in *Materials* is shown in Figure 5.

[7] Each wavelength contains the same image, the idea being simply to simulate the presence of spectral data.

(a) (b) (c)



(d)

Figure 5: Test images. (a) Shepp-Logan Phantom. (b) 'Lena'. (c) Materials, $\lambda = 212$. (d) Spectral signatures of the six different materials in Materials test.

All experiments have been carried out in MATLAB 7.11.0 (64-bit) on a 3.33GHz quad-core desktop computer. Because neither ℓ_1-Magic nor TVAL3 can handle HS data, results for the *Shepp-Logan Phantom* and '*Lena'* tests were obtained by performing recovery on a single wavelength and then extrapolating over 256 wavelengths. However, because *each* wavelength is different for the *Materials* HS test image, no such technique was available and experiments have only been performed using our algorithm.

Table 1 compares the recovery PSNRs and computational times for each algorithm. TVMAβ had the fastest computational times for all tests, with an average speed improvement of 22.4% over TVAL3. PSNR results for our algorithm were within 1 dB of TVAL3 reconstructions for five of the eight comparable tests, and better than the ℓ_1-Magic reconstructions in all cases. Figure 6 shows an example reconstruction of *'Lena'* using TVMAβ.

5 FUTURE WORK

We intend to continue our investigation into TV minimisation and its use in CS-HSI recovery, and improve our understanding of the available TV minimisation algorithms. For TVMAβ, we aim to improve the β stepping, develop more robust stopping criteria, and utilise spectral smoothness.

6 CONCLUSIONS

This paper has outlined the use of CS in HS image acquisition, showing its potential impact in making HSI processing viable to the mainstream. Both acquisition and recovery theory has been explored, during which the

Table 1: Algorithm comparison. Each trial repeated five times and averaged.

Ratio M/N	ℓ_1-Magic [4]			TVAL3 [5]			TVMAβ		
	PSNR (dB)	Time (s)	Time per wavelength (s)	PSNR (dB)	Time (s)	Time per wavelength (s)	PSNR (dB)	Time (s)	Time per wavelength (s)
Shepp-Logan Phantom									
0.15	33.30	342272	1337	**33.85**	1976	7.72	33.41	**1960**	**7.66**
0.20	35.05	311296	1216	**38.29**	1973	7.71	37.01	**1879**	**7.34**
0.30	37.91	361101	1411	**43.94**	1803	7.04	41.37	**1590**	**6.21**
0.50	39.98	370432	1447	**54.15**	1971	7.70	44.23	**1677**	**6.55**
'Lena'									
0.15	26.58	312581	1221	**27.42**	2206	8.62	26.90	**1675**	**6.54**
0.20	27.17	309381	1209	**28.14**	2141	8.36	27.96	**1458**	**5.69**
0.30	28.55	346880	1355	29.33	1957	7.64	**29.69**	**1005**	**3.93**
0.50	30.06	348520	1361	31.19	1894	7.40	**31.96**	**1097**	**4.28**
Materials									
0.15	-	-	-	-	-	-	**26.23**	**1136**	**4.44**
0.20	-	-	-	-	-	-	**29.01**	**1263**	**4.94**
0.30	-	-	-	-	-	-	**30.86**	**1280**	**5.00**
0.50	-	-	-	-	-	-	**31.99**	**1283**	**5.01**

(a) (b) (c) (d)

Figure 6: Recovery example of a single wavelength of 'Lena', with $M/N = 0.3$ (a) Original image (b) Recovered via TVMAβ, $PSNR = 29.69$ dB. (c) Zoom of (a). (d) Zoom of (b).

use of the TV norm in CS-HSI data recovery has been outlined. Our proposed algorithm incorporates a function designed for HS data handling, and when compared with the state-of-the-art TVAL3, boasts improved computational times and mostly comparable, sometimes superior, reconstruction results.

7 REFERENCES

[1] R. J. Bowmaker, R. J. Dunn, K. B. Moynihan, T. J. Roper, and M. Andrews, "Construction of a Practical Hyperspectral Image Acquisition System," in *26th International Conference Image and Vision Computing New Zealand*, 2011.

[2] R. Baraniuk, M. Davenport, M. Duarte, and C. Hegde, *An Introduction To Compressive Sensing*, 1.5 ed. URL: http://cnx.org/content/col11133/1.5/, 2011.

[3] C. Li, "An Efficient Algorithm for Total Variation Regularization with Applications to the Single Pixel Camera and Compressive Sensing," Master of Arts, Rice University, Houston, Texas, 2009.

[4] J. Romberg, "l1-Magic," *Georgia Tech School of Electrical and Computer Engineering*. [Online]. Available: http://users.ece.gatech.edu/~justin/l1magic/. [Accessed: 20-Sep-2012].

[5] "TVAL3 Download." [Online]. Available: http://www.caam.rice.edu/~optimization/L1/TVAL3/welcome-download.html. [Accessed: 24-Sep-2012].

[6] A. M. Baldridge, S. J. Hook, C. I. Grove, and G. Rivera, "The ASTER spectral library version 2.0," *Remote Sensing of Environment*, vol. 113, no. 4, pp. 711–715, 2009.

[7] E. J. Candès and M. B. Wakin, "An Introduction to Compressive Sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 21–30, 2008.

[8] D. L. Donoho, "Compressed Sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, 2006.

[9] E. J. Candès, J. Romberg, and T. Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete. Frequency Information," *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, 2006.

[10] E. J. Candès and T. Tao, "Decoding by Linear Programming," *Information Theory, IEEE Transactions on*, vol. 51, no. 12, pp. 4203–4215, 2005.

[11] J. Romberg, "Imaging via Compressive Sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 14–20, 2008.

[12] E. J. Candès, J. K. Romberg, and T. Tao, "Stable Signal Recovery from Incomplete and Inaccurate Measurements," *Communications on pure and applied mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.

[13] R. Baraniuk and K. Kelly, "Compressive Imaging: A New Single-Pixel Camera," *Rice University DSP*. [Online]. Available: http://dsp.rice.edu/cscamera. [Accessed: 09-Feb-2012].

[14] W. Lee, D. Eason, and M. Andrews, "Compressive Hyperspectral Image Sensing Restoration via Joint TV-L1 Regularisation," *submitted to ENZCon*, 2012.

[15] J. Barzilai and J. M. Borwein, "Two-Point Step Size Gradient Methods," *IMA Journal of Numerical Analysis*, vol. 8, no. 1, pp. 141–148, 1988.

[16] L. Gan, "Block Compressed Sensing of Natural Images," in *Digital Signal Processing, 2007 15th International Conference on*, 2007, pp. 403–406.

Removing Floor Reflections to Improve Robot Navigation with a 3D-TOF Camera

Sean G. J. Kelly, Adrian A. Dorrington

Dept. Engineering, University of Waikato,
Private Bag 3105, Hamilton, New Zealand
Email: sk165@students.waikato.ac.nz

Abstract: Currently time of flight range imaging cameras suffer from distortions in their range measurements due to multipath effects. For a robot to autonomously navigate an unknown environment it requires accurate 3D information about what is in front of it. Multipath effects cause a range imaging camera to see holes in highly reflective floor surfaces, such as vinyl. This paper seeks to remove floor reflections for range imaging cameras. Reflections cause measured points to be further away than the floor really is. To remove the reflections these points are removed. This approach was able to remove all reflections from a video recording in a way which mimicked real time conversion. The method outlined in this paper works well in level environments and where the camera maintains a fixed distance from the floor. However the major pitfall of this technique is that if there are real holes in the floor such as downward stairs they will be removed.

Keywords: Time of flight, Robotic Navigation, Machine Vision, Floor Segmentation, Range Imaging,

1 INTRODUCTION

Highly reflective floor surfaces cause time of flight (TOF) range imaging cameras to see non-existent holes in the floor. These holes make autonomous navigation for a robot system very difficult. An autonomous system will make poor decisions if it has inaccurate data. This paper focuses on removing floor reflections.

Robot navigation is a complex task which can involve multiple sensors and require large amounts of computing power. One of the most difficult tasks for a robot is to know where it is located within 3D space. Many studies [1, 2, 3] have been done dealing with the complex task of mapping an unknown 3D environment and determining the current location of the robot. This is also known as simultaneous localisation and mapping (SLAM) [2].

Time of flight range imaging has advantages over stereoscopic vision. Stereoscopic vision requires image matching to measure distance, a computationally complex process, whereas TOF cameras require little post processing. In principle TOF cameras decrease the overall computational complexity for a robotic navigation system.

There are two main approaches to time of flight range imaging, the *pulse gated* approach that directly measuring the speed of light (using optical shutter technology [4]), and the *Amplitude Modulated Continuous Wave* (AMCW) method the measures the phase change due to propagation delay of the amplitude modulated light. The phase measurement is generated by comparing the received signal with a reference signal [5]. The AMCW method is most common, and therefore will be the focus of this paper.

A well-known problem in range imaging is due to its requirement to actively illuminate a scene. Ideally, light will travel directly from the camera to the scene and straight back again. However, some light will reflect off multiple objects before returning to the camera. This effect, known as multipath interference, introduces distortions in the range data. It has a tendency to round sharp angles such as the corners of a room. It also means that if there is a highly reflective object in the scene, such as exposed metal, it will look vastly different to what it really is.

In this paper, we are particularly interested in the multipath effects that occur due to highly reflective floor surfaces. If an object is on a highly reflective floor surface, to a range imaging camera, it could look twice as big as it should due to its reflection in the floor. If a robot which has a range imaging camera attached had the ability to pick up objects it might attempt to grab a piece of the floor because it appears to be part of the object.

Section 2 covers the theory behind how a range camera works and how the floor reflections were removed. Section 3 covers the physical hardware and setup. Section 4 gives the results and discussion. It is concluded in section 5.

2 THEORY

The camera used for this paper is the SR-4000 developed by MESA Imaging. It uses an amplitude modulated continuous wave (AM-CW) system of measurement. 24 infrared LEDs modulated at 30 MHz illuminate the cameras field of view. This light is scattered by objects in the scene, and a portion returns to the cameras. The modulation envelope of light returning from distance objects is phase shifted due to the propagation delay, and the camera measures this phase shift to determine distance. At 30 MHz the phase change over 360° corresponds to 1 m. Because the camera measures the out and back return path the camera is limited to measuring objects at a maximum range or non-ambiguity distance of 5 m.

To determine the modulation envelope phase, and hence distance, the camera uses a specialised image sensor that is gain modulated, also at 30 MHz. This effectively down-converts the detected illumination signal in a similar way to a lock-in amplifier. Quadrature detection is used, requiring four separate raw exposures with relative phase differences between the illumination and sensor modulation signals of 0°, 90°, 180° and 270° are acquired, producing images C_1, C_2, C_3 and C_4 respectively. The amplitude, A, and phase, θ, for each pixel in the cameras sensor array can be calculated from the first bin of a Fourier transform, simplified to equations 1 and 2 [6, 7]. Note that pixels are treated individually, so pixel notation is omitted for clarity.

$$A = \frac{\sqrt{(C_4 - C_2)^2 + (C_1 - C_3)^2}}{2} \quad (1)$$

$$\theta = \tan^{-1}\left(\frac{C_4 - C_2}{C_1 - C_3}\right) \quad (2)$$

The distance to each pixel can then be calculated from the phase difference and the non-ambiguity range $D_{max} = c/2f_m$. Where c is the speed of light and f_m is the modulated light frequency.

$$D = \frac{D_{max}\theta}{2\pi} \quad (3)$$

Multipath interference is a form of the mixed pixel problem associated with TOF cameras. In an ideal case, light would travel from the camera to the scene and straight back to the camera. In practice this is not the case, light reflects off multiple objects before returning to the camera. These multiple reflections cause distortions in the depth data. The higher the reflectivity of an object is, the more multipath interference it has. Highly reflective floors pose a problem to a robot using a TOF camera because the camera will measure the distance to the refection rather then the distance to the floor, as shown in figure 1. Therefore, these reflections below the floor line appear as a hole in the floor.



Figure 1 - Comparison of a scene with and without multipath interference.

Reflections in the floor cause the TOF camera to "see" points which are further away than they really are. If the camera is located at a fixed angle to the floor, the floor's expected location can be calculated. To correct an image containing reflections, all points further away than where the floor is expected to be can be removed. This can be done by comparing an image to a plane equation that describes the position of the floor. To create the plane equation a calibration image is acquired and a flat plane is fitted to known correct floor points. This calibration image should be made on a floor type that has as little multipath interference as possible.

Another way of solving the multipath interference problem uses a Mixed Pixel Restoration algorithm [8]. It works by comparing two images with different modulation frequencies. This works well for static images but changing between modulation frequencies causes a delay. The delay means it is too slow to work in real time.

3 METHOD

The robot used is the Pioneer 2 designed by Mobile Robots. The case was modified by attaching an optical rail across it as in figure 2. It has an on-board computer and is capable of wireless communication. It is running Gentoo Linux 2.4.22 v3. It is uses a software package called ARIA. Controlling the robot was done by logging in wirelessly and using an ARIA program called "teleop".

The TOF range imaging camera used is an SR4000 developed by Mesa Imaging. It has an image resolution of 144 by 176 pixels and a depth resolution of less than one centimetre. It was mounted low down on the robot using an optical rail and an articulated mount as in figure 2. The camera was connected to a Macbook Pro "core 2 duo" which was running Windows XP service pack 3. Controlling this laptop was done wirelessly using the Windows tool virtual desktop. The software used to retrieve data from the

camera is custom software written using MATLAB by Robert Larkins.

Several camera locations ranging from 0.3 m to 1.0 m off the ground were tested. The final location of 0.3 m was selected because it allowed the camera to be tilted at only a slight angle. This meant higher objects were visible from further away, while minimising the blind spot directly in front of the robot. This improved an operator's experience, who was controlling the robot remotely.

The end of a hallway with a highly reflective vinyl floor was used as the test scene. A watering can painted with non-glossy white paint was placed 20 cm from the end. One hundred and fifty images where taken with a static camera. Those images were then averaged in MATLAB to produce a single image.

Calibration images were taken of a carpet floor that has very little multipath interference. The plane created from the calibration was then used to remove reflections from the test image.

Video data was also recorded of the robot moving toward the watering can and away from it. MATLAB was used to remove reflections from the video in a way that mimicked a real time conversion.



Figure 2 - The Pioneer 2 with SR400 camera and Macbook Pro

4 RESULTS AND DISCUSSION

Figure 4 is a depth image of a watering can at the end of a hallway with a highly reflective vinyl floor. Red indicates distances further from the camera and blue closer to the camera. The end wall is approximately 4 m from the robot. It is a clear example of multipath interference. The watering can is reflected in the floor and the wall behind the watering can is also reflected in the floor up to the line through the middle of the photo. That line is seen as a large drop off to the range camera.

One of the big advantages of TOF camera technology is the ability to view a captured scene from multiple angles. Figure 3 is a side on view of the hallway shown in Figure 4, with a plane drawn through the middle. Below the plane is a large series of points that correspond to reflections in the floor. The points above the plane relate to the sides and end of the hallway.



Figure 3 - Side on view of 3D data, showing a plane where the floor should be in blue, walls in green, and unwanted reflections in red.



Figure 4 - Original range image of the end of the hallway. The blue is correct data for the floor closer to the camera, but note the unwanted reflection of the watering can in the floor further from the camera.

Although the place of the floor can be found quite precisely (because the plane is fit to many points in the calibration data) the precision of the real-time 3D data is quite variable. Typically the SR-4000 has a measurement precision of around 5 to 10 cm, dependent on the amount of light returned to the camera, and hence the objects colour. Dark objects (in the infra-red) have a much worse precision than light objects. In addition small amounts of multi-path interference in areas where reflections are not apparent can also cause the floor to warp upwards, making the floor appear slightly warped or higher than is should be.

To avoid these random variations in the real-time data appearing as obstacles in front of the robot, a threshold is used to alter the height of the plane used

to remove reflections. Threshold levels of 0 cm, 3 cm, and 6 cm were tested and the results compared, producing figures 5, 6 and 7 respectively. The light part of the floor in these figures is where part of the floor is higher than the level of the calibration plane. The dark part is where the reflections have been removed.

With a 6 cm threshold, all of the floor and the reflections have been removed. This improves both manual and autonomous robot control because anything above the floor plane can be considered as an obstacle to be avoided. Whit less threshold, parts of the floor show through, potentially confusing object detection. Visually it is easier for a user to interpret an image without reflections or distortions in the floor. This should improve the operators ability to control the robot. Also, in an autonomous navigation situation the robot will be able to make correct decisions based on its environment. Without removing reflections in the case of figure 4 an autonomous robot will react incorrectly to the environment because it will see a hole in the floor.

A pitfall with having the threshold at 6 cm is that the bottom of objects is removed. This can be seen by comparing figures 5 and 7. Objects smaller than 6 cm could be high enough to disrupt the movement of the robot, but would not be detected.

Frames from the video data have not been included in this report because they show little difference to the static test images. The reflections were able to be removed in real time without a noticeable slowdown in the time taken to render each frame.

With the improvements there are several drawbacks. The largest is that if there really is a hole in the floor the camera will not see it. An example of this is a set of downward stairs as they are totally invisible to this method of reflection removal. This method also relies on the camera maintaining a constant distance from the floor. If the camera changes angle while attached to the robot it will cause the calibration to fail.

To overcome these limitations, more advanced threshold determination techniques may be useful. In addition, it may be possible to combine this method and the Mixed Pixel Restoration algorithm discussed at the end of section 2 in a two-step process. Although the Mixed Pixel Restoration algorithm is to slow to be used in real time, the robot could first stand still for a short period to obtain a measurement of the scene free from multi-path interference, which will enable it to detect the correct floor location and any actual holes in the floor (such as stairs). Once the robot has planned it's path to avoid stationary objects or holes, it could switch to the presented method for real-time navigation, allowing a measurement rate fast enough to detect and avoid moving objects such as people walking in front of it.

5 CONCLUSIONS

A method for removing multipath interference from Time-of-Flight 3D measurements of highly reflective floor surfaces is outlined in this paper. It improves the vision capability for robot navigation applications buy removing incorrect data from the scene, and allows easier identification of obstacles in the robot's path. This technique works well in level environments, however the major drawback is that it cannot identify large drops such as downward stairs.



Figure 5 - Reflection removal with a 0 cm threshold



Figure 6 - Reflection removal with a 3 cm threshold



Figure 7 - Reflection removal with a 6 cm threshold

ACKNOWLEDGEMENTS

The authors would like to thank Lee Streeter, Robert Larkins and the other members of the Chronoptics for their valuable support and assistance.

REFERENCES

[1] Guðmundsson, S. (2006). Robot Vision Applications using the CSEM SwissRanger Camera. Masters Thesis. Technical University of Denmark, Kongens Lyngby, Denmark.

[2] Almansa-Valverde, S. Carlos Castillo, J. Fernández-Caballero, A. (2012). Mobile Robot Map Building from Time-of-Flight Camera. Expert Systems with Applications, Vol 39, Issue 10, DOI:10.1016/j.eswa.2012.02.006.

[3] May, S. Droeschel, D. Holz, D. Fuchs, S. Malis, E. Nüchter, A. Hertzberg, J. (2009). Three-Dimensional Mapping with Time-of-Flight Cameras. *Journal of Field Robotics (JFR)*, Special Issue on Three-dimensional Mapping, Part 2. Volume 26, issues 11-12, pages 934-965. DOI: 10.1002/rob.20321

[4] Gvili, R.; Kaplan, A.; Ofek E.; Yahav, G. (2003). Depth keying. *Proc. SPIE* **2003**, *5249*, 534-545.

[5] Lange, R. (1999) Time-of-Flight range imaging with a custom solid-state image sensor. *Proc. SPIE* **1999**, *3823*, 180-191.

[6] T. Oggier et al., (2003). An all-solid-state optical range camera for 3D real-time imaging with sub-centimeter depth resolution (SwissRangerTM), SPIE Proc Vol. 5249-65, St. Etienne.

[7] McClure, S. (2009). Remote Control of a Robot Using Time of Flight Range Data. BE Design Report. Hamilton, New Zealand.

[8] Dorrington, A.A. Godbaz, J.P. Cree, M.J. Payne, A.D. Streeter, L.V. (2011) Separating true range measurements from multi-path and scattering interference in commercial range cameras, Proc. SPIE 7864, 786404 (2011).

A GPS Signal Synthesiser for Systems Integration Testing of the Transient Array Radio Telescope

C.F. Shaw, T. C. A. Molteno*

*Department of Physics, University of Otago
Box 56, Dunedin, NEW ZEALAND.
Email: tim@physics.otago.ac.nz

Abstract: The Transient Array Radio Telescope is a wide field synoptic synthesis array operating in the GPS L1 band. This paper focuses on the systems integration testing of the data collection and processing unit for this instrument. The required test signals are produced by GPS signal generator, implemented in a Xilinx Spartan-3 Xc3s1000 FPGA. Test results indicate the expected functionality.

Keywords: Systems Integration Testing, Imaging Pipeline, Radio Astronomy, Interferometry, Transient Event, Sensitivity and Sampling, GPS PRN C/A Codes, Correlation.

1 Introduction

The Transient Array Radio Telescope (TART) is a wide field synoptic synthesis array that operates in the GPS L1 band. Its primary purpose is the radio detection of transient astronomical events such as supernovae.

Radio telescopes generally average a signal over a long time period to maximise sensitivity, but this restricts their ability to observe transient events. Because of it's short averaging time, the TART is able to detect events of short duration, at the expense of sensitivity.

The sensitivity of a radio telescope is defined by the signal to noise ratio. From equation 1 it can be shown that this is directly related to the number of samples.

$$\mathcal{R}_{sn} = C\frac{T_A}{T_S}\sqrt{\Delta\nu\tau} \qquad [2] \qquad (1)$$

Where C is a constant, T_A is the antenna temperature, T_S is the system temperature, $\Delta\nu$ is the bandwidth, and τ is the signal averaging time. If a signal is averaged for time τ, it contains approximately $2\Delta\nu\tau$ independent samples, so increased sample size improves sensitivity.

1.1 Global Positioning System L1 Band

A primary consideration when designing a radio telescope is the selection of an operating band. For this instrument, there are compelling reasons to select the GPS L1 band, centered on 1.5742 GHz.

Large astronomical events such as supernovae generate a strong radio signal across the entire spectrum. This includes the almost silent GPS L1 navigation band, which is protected by legislation. GPS signals themselves are very faint, as low as -166 dBw at ground level [1].

GPS satellites may used as guide stars for calibration purposes, as positions are very accurately known. The embedded GPS data time and position stamps each observation precisely.

In order to detect GPS signals, every GPS Radio Front End incorporates a very high performance Low Noise Amplifier (LNA). Use of GPS LNAs allows optimisation of $\mathcal{R}_{sn}$ by reducing T_S in Equation 1.

1.2 Overview of Telescope Sub-Systems and Testing

The TART includes the following subsystems

1. Two to eight antennas.
2. One radio front end / ADC per Antenna.
3. One clock conditioning transceiver per antenna.
4. A clock generation and data receive unit.
5. A data buffering and serial transmission unit.
6. A computer for control, data storage and correlation function calculation.

This paper focuses on the development and testing of sub-systems 5 and 6. These will be tested as a unit using a simulated GPS signal.

2 Pseudo Random Noise Generation

Signals from cosmic sources almost universally have the form of Gaussian random noise [2]. Their autocorrelation yields the Dirac delta function, and this property can be used to determine the position of a source.

Pseudo Random Noise (PRN) is a family of deterministic sequences that share many of the properties of true random noise, but repeat at regular intervals. A PRN sequence of length $2^n - 1$ can be generated using an n-bit Linear Feedback Shift Register (LFSR).

Combining signals from two LFSRs allows the generation of a family of distinct PRN sequences with low cross-correlation. These Gold codes are used in GPS to produce the Coarse Acquisition (C/A) codes [1]. The GPS C/A codes are used as synthetic signals for testing the TART.

2.1 Linear Feedback Shift Register (LFSR)

The functioning of this LFSR is illustrated in Figure 1, and the state diagram in Table 1. The bit positions 3 and 4 provide linear feedback affecting the next state, these are referred to as taps. Note the repeat every 15 iterations.

1. At reset, a seed value is loaded into the shift register. $\{1, 1, 1, 1\}$ is used in this case, but any other non-zero value is acceptable. A seed value of $\{0, 0, 0, 0\}$ would prevent a change of state, and must be avoided.
2. Bits 3 and 4 of the shift register are XORed, and the result stored until the next clock cycle.
3. On the positive clock edge the contents of the shift register are shifted one bit to the right, leaving a space at bit 1. The right most bit is discarded.
4. The result calculated in step 2 is placed at bit 1.
5. Steps 2 to 5 are repeated indefinitely until a reset signal initiates a return to step 1.



Figure 1: A simple 4 bit LFSR. MSB on right. Taps at position 3 and 4

2.1.1 Implementation of Flexible Tapping

Recall that two different LFSRs are required to generate a Gold code sequence. Instead of coding these individually, they are synthesised in a single Verilog module using parameters for register size and tap positions.

Figure 2 illustrates the implementation of flexible tapping. Note that in comparison to Figure 1, each bit of the shift register is tapped, and that a tap register and array of AND gates has been added.

Table 1: State Diagram of LFSR in Figure 1. MSB at right.

Iteration	Shift Register Status	sr(3) XOR sr(4)
1	1 , 1 , 1 , 1	0
2	0 , 1 , 1 , 1	0
3	0 , 0 , 1 , 1	0
4	0 , 0 , 0 , 1	1
5	1 , 0 , 0 , 0	0
6	0 , 1 , 0 , 0	0
7	0 , 0 , 1 , 0	1
8	1 , 0 , 0 , 1	1
9	1 , 1 , 0 , 0	0
10	0 , 1 , 1 , 0	1
11	1 , 0 , 1 , 1	0
12	0 , 1 , 0 , 1	1
13	1 , 0 , 1 , 0	1
14	1 , 1 , 0 , 1	1
15	1 , 1 , 1 , 0	0
Cycle Repeats	1 , 1 , 1 , 1	0

The tap and shift registers are connected to the inputs of an AND gate at each bit position. The AND output connects to the first input of an XOR gate. When the tap register bit is high, the content of the shift register will be applied to the first XOR input, and the circuit will behave as if a tap is present. When the tap register bit is low, the XOR gate will link the output of the previous XOR gate to the input of the next.

The Verilog code used to generate a flexibly tapped LFSR is shown in Listing 1.



Figure 2: One of two 10 Bit Linear Feedback Shift Registers (G_2) required to produce a C/A code. Shift register (top) in initial state. Taps are set at $\{10, 9, 8, 6, 3, 2\}$ by the tap register (bottom).

2.2 Generation of C/A Code

The C/A code generator used to produce the various GPS signals is outlined in Figure [3] and in Listing [2].

Both G_1 and G_2 are 10 bit LFSRs, internally tapped at positions $\{10, 3\}$ and $\{10, 9, 8, 6, 3, 2\}$ respectively. They generate sequences of 1023 chips.

The code phase selector determines which PRN sequence is generated . In this case it is set to add bits 3 and 8 of G_1. This result is then added to the output of G_1 to produce PRN code #31.

The code phase assignments required to generate all 37 PRN codes are shown in Table 3. Of these, the first 32 are

Listing 1: Verilog code for flexibly tapped LFSR

```
module fibonacci_lfsr
        #(parameter
        bits = 4d10,
        taps = 10b1110100110,
        init_state = 10b1111111111)
        (input clock,input reset,output wire [bits−1:0]signal_out);
        reg [bits−1:0]state = init_state;
        always @(posedge clock, posedge reset)
                begin
                if (reset)
                                state <= init_state;
                        else
                                state <= {state[bits−2:0],^(state & taps)};
                end
        assign signal_out = state[bits−1:0];
endmodule
```



Figure 3: C/A code generator. [1]

used by the GPS satellites, with the remainder allocated to other uses, including ground stations.

A variety of PRN outputs were simulated. A comparison with published literature [1] indicates correct operation. Results for PRN #1 are shown in Table 2.

A reset signal returns both LFSRs to all ones. The all ones condition is detected by the Decode unit, allowing a pulse to be generated once every 1023 chips, at the completion of a sequence. This pulse is used to generate a 1kHz clock, and further divided by 20. The resulting 50Hz clock is used to generate ephemeris data that contains the orbital and clock parameters for a satellite.

Table 2: Test output of Verilog C/A code generator module showing first 10 chips of PRN sequence #1

Chip	G_1 Status	G_2 Status	Output
1	1111111111	1111111111	1
2	1111111110	1111111110	1
3	1111111100	1111111100	0
4	1111111000	1111111001	0
5	1111110001	1111110010	1
6	1111100011	1111100101	0
7	1111000111	1111001011	0
8	1110001110	1110010110	0
9	1100011100	1100101101	0
10	1000111000	1001011010	0

Listing 2: Verilog Code for C/A Code Generator

```
module c_a_code_generator
        #(parameter
        G1_bits = 4'd10,
        G1_taps = 10'b1000000100, // Shift register G1
        G1_init = 10'b1111111111,
        G2_bits = 4'd10,
        G2_taps = 10'b1110100110, // Shift register G2
        G2_init = 10'b1111111111,
        phase_sel_taps = 10'b0000100010) // Default PRN #1
        (input clock, input reset, output wire kHz_out,
        output wire c_a_code_out);
        wire [G1_bits−1:0] G1_out;
        wire [G2_bits−1:0] G2_out;
        // Pulse provided to mark the start of the C/A code ------------------------
        assign kHz_out = & G1_out; // output only if all bits 1
        // Instantiate shift register G1 ----------------------------------------------
        fibonacci_lfsr #(.bits(G1_bits), .taps(G1_taps),.init_state(G1_init))
        G1(.clock(clock),.reset(reset),.signal_out(G1_out));
        // Instantiate shift register G2 ----------------------------------------------
        fibonacci_lfsr #(.bits(G2_bits), .taps(G2_taps),.init_state(G2_init))
        G2(.clock(clock),.reset(reset), .signal_out(G2_out));
        // Combine signals from G1 & G2 to produce c/a code -------------------
        assign c_a_code_out = (^(G2_out & phase_sel_taps))^G1_out[G1_bits−1];
endmodule
```

Table 3: Code phase selection required to generate GPS PRN signals. Only codes 1 to 32 are used for satellites. C/A codes 34 and 37 are common. [1]

PRN #	Code Phase	PRN #	Code Phase	PRN #	Code Phase
1	$2 \oplus 6$	**14**	$7 \oplus 8$	**26**	$6 \oplus 8$
2	$3 \oplus 7$	**15**	$8 \oplus 9$	**27**	$7 \oplus 9$
3	$4 \oplus 8$	**16**	$9 \oplus 10$	**28**	$8 \oplus 10$
4	$5 \oplus 9$	**17**	$1 \oplus 4$	**29**	$1 \oplus 6$
5	$1 \oplus 9$	**18**	$2 \oplus 5$	**30**	$2 \oplus 7$
6	$2 \oplus 10$	**19**	$3 \oplus 6$	**31**	$3 \oplus 8$
7	$1 \oplus 8$	**20**	$4 \oplus 7$	**32**	$4 \oplus 9$
8	$2 \oplus 9$	**21**	$5 \oplus 8$	**33**	$5 \oplus 10$
9	$3 \oplus 10$	**22**	$6 \oplus 9$	**34**	$4 \oplus 10$
10	$2 \oplus 3$	**23**	$1 \oplus 3$	**35**	$1 \oplus 7$
11	$3 \oplus 4$	**24**	$4 \oplus 6$	**36**	$2 \oplus 8$
12	$5 \oplus 6$	**25**	$5 \oplus 7$	**37**	$4 \oplus 10$
13	$6 \oplus 7$				

3 FPGA System Integration

3.1 Overview

The Data Buffering and Serial Transmission (DBST) unit is to be tested in isolation from the rest of the TART system. The sampling clock and sampled signals required for systems integration and testing are generated internally. Figure [4] shows the unit in its test configuration.

3.2 Generation of Clock Signals

The GPS system requires the entire C/A code to repeat every millisecond. There are 1023 chips in the code sequence, so the C/A Code clock must be $1\text{ kHz} \times 1023 = 1.023\text{ MHz}$. GPS receivers typically sample at 16.368 MHz, 16 times this rate.

Unfortunately, the Xc3s1000 FPGA system used to produce the test signal cannot generate these frequencies internally, so 1 MHz and 16 MHz clocks are used instead.

A 50 MHz clock is increased to 160 MHz using the chips hardware fractional frequency multiplier, and divided by 10 and 16 to give 16 MHz and 10 MHz clocks. The 10 MHz clock is further divided down to 1 MHz.

3.3 The Control Buffer Finite State Machine

The control buffer records samples of fixed length from two to eight sources. The sampled data is then sent to the UART for transmission a to a PC via an RS232 link. The states of this Finite State Machine (FSM), and the transitions between them are described as follows;

Idle: This is the initial state. It is also entered upon system reset, or after all the contents of the memory have been read. Read and Write Pointers are reset to zero. The *SAMP BEG* signal selects the write state.

Write: In this state, the write pointer is advanced on the positive edge of sample clock. When the memory is full, the FSM changes to the read state.

Read: In this state, the contents of the control buffer are written to the UART a block at a time. The read pointer advances on the positive edge of the sampling clock until a *TX FULL* signal is received from the UART indicating that it's FIFO is full. At this point the FSM changes to the read wait state. When the control buffer memory has been fully emptied, the FSM is returned to the idle state

Read Wait: This state is entered from the read state when the UART FIFO is full. When the *TX EMPTY* signal indicates that the UART FIFO has emptied, the FSM returns to the read state.

Memory Utilisation and Sample Sizes

The control buffer utilises up to 54 kB of block RAM available on the Xc3s1000 chip. The width of the memory is determined by the number of signals to be simultaneously recorded - between two and eight. In this implementation, two channels are required, so the memory is configured as 2×2^{17} bits. Eight channels would be configured as 8×2^{15} bits. In each case 32 kB of block RAM are used. The use of three or seven channels would be most efficient, and allow the utilisation of 48 kB.

3.4 The Universal Asynchronous Receiver / Transmitter

The UART module provides an RS232 interface for communication with a computer. The interface provides adequate speed, is readily available, is straightforward to implement. The required drivers and DB9 connector are available on the Spartan-3 development board used.

The baud rate, number of data bits and stop bits, and the with of the FIFO are defined as parameters. These are set as 115200 kBit/s, 8 bits, 1 bit, and 16 bytes respectively. Existing code for the UART [3] has been modified to suit this application.

The UART module incorporates a baud rate generator, a receive unit, a transmit unit, and two FIFOs. While the facility for full duplex communication is provided, only the transmit functionality is used a this time.

Each word transmitted contains a one bit sample from up to eight input channels. The use of eight channels is dictated by the maximum word size in the RS232 standard. Modification of the Verilog code would make it possible to use more than eight channels. This would require multiple bytes of data to be sent per sample. Because only two channels are used, data is confined to bits 0 and 1. The other bits are set to 0,0,1,1,0,0. This results in a byte that corresponds to the ASCII character 1, 2, or 3. A 1 indicates that the first channel is high, a 2 that the second channel is high, and a 3 indicates that both channels are high.

3.4.1 Asynchronous Clocks

There are two asynchronous clock signals provided to the UART. The first is the 50 MHz clock used by the baud rate generator, and the second is the sampling clock at 16.368 MHz (or 16.0 MHz for testing).

The transmit unit operates off a 50MHz clock, and indicates transmission of a byte of data by a one clock cycle wide tick. The FIFO uses the sampling clock , and provides the next byte of data upon receiving the *TX DONE TICK*. The width of the tick provided by the transmit unit is shorter than the clock cycle of the FIFO. Unless this tick happens to coincide with the rising edge of the sampling clock, it will be missed, and the same byte will be sent multiple times.

This issue is addressed by the provision of a modified RS Flip-Flop to make the *TX DONE TICK* into a wider *TX DONE TOCK* so that it can be seen.

The FIFO provides a *TX START* signal to the transmit unit whenever it contains data. This signal must be held low until the *TX DONE TOCK* is complete.

3.5 Computer Software

A Python script performs the functions detailed below;

1. RS232 port is opened, and data stored as a text file.
2. The string is converted first to an integer array (containing 1 to 3), then to binary.
3. Each column of the array (an input channel) is converted to one bit sign magnitude format (-1 or 1).
4. The cross-correlation function of the two signals is calculated and displayed. Results written to a file.



Figure 4: Block diagram of the top level FPGA implementation.

4 Testing

4.1 Observation of Auto-correlation Peaks

For the first test, two PRN #1 C/A code sequences were generated in phase at 1 MHz, sampled 16 MHz, and correlated in accordance with Equation 2 [4].

The correlation function θ of two binary sequences a and b is defined as

$$\theta(\tau) = \sum_{i=0}^{2^N-1} \chi a(i) \chi b(i+\tau) \tag{2}$$

Where χ is the operator mapping $[0, 1]$ onto $[-1, 1]$.

Figure [5] shows the region near the main correlation peak at zero lag. Note also the presence of minor peaks to the left and right. These are expected for all non-maximal sequences including Gold codes.

A key advantage of Gold codes is that these minor correlation peaks lie within bounds defined by Gold [4] as

$$|\theta| \leq \begin{cases} 2^{(n+1)/2} + 1 \text{ for n odd} \\ 2^{(n+2)/2} + 1 \text{ for n even} \end{cases} \tag{3}$$

Note that $n = 10$ (even) for all Gold codes used in GPS.

A comparison of the measured and theoretical values is worthwhile. The maximum correlation peak has a value



Figure 5: Auto correlation of 2^{14} bits from PRN sequence #1 sampled at 16MHz. An enlargement of the central section shows the major correlation peak at zero lag, and the minor peaks to either side.

of 16384, while the absolute value of largest minor correlation peak is 1184. Dividing these values by 16 (samples/chip) gives 1024 and 74 respectively.

The expected value of the main correlation peak [5] [4] is shown in Equation 4

$$\theta_{peak} = 2^n - 1 = 1023 \tag{4}$$

The remaining values are expected to satisfy the inequality

$$\begin{aligned} \theta &\leq 2^{((10)+2)/2} + 1 \\ &\leq 65 \end{aligned} \tag{5}$$

The theoretical (1023, 65) and measured values (1024, 74) are close, and discrepancies will be investigated.

4.2 Observation of Periodicity

Any periodicity present in a signal can be established by examining it's auto-correlation function. Figure [6] shows the auto-correlation of a pair of PRN #1 codes. The code has a chip rate of 1 MBit/s, and repeats every 1023 cycles. Three major correlation peaks are clearly visible with a separation of $1023\mu s$, as expected.

4.3 Observation of a Phase Shift

The correlation function was calculated for a pair of PRN #1 signals with a 20 chip phase offset. Figure [7] shows the individual samples near the main correlation peak at $-20\,\mu s$. This position is as expected for a chip period of $1\,\mu s$. Within one chip of the peak, the correlation increases linearly, giving the expected triangular function [6].



Figure 6: Auto-correlation of 2^{16} bits from PRN sequence #1 sampled at 16MHz. Note periodicity of major peaks.



Figure 7: The correlation two PRN #1 sequences with phase delays of 20 and 40 chips respectively. The correlation peak has moved from the centre position indicating a 20 μs delay. Note also the linear increase in the correlation function within 1 chip of the major peak.

4.4 Observation of Cross-correlation

A key advantage of Gold codes is that they allow generation of a family of codes for which the cross-correlation functions are both small and well defined. This allows multiple signals to share the same band, and is essential in all code division multiple access systems including GPS. The cross-correlation function of C/A PRN sequences #1 and PRN #13 is shown in Figure 8. The cross-correlation also obeys the inequality in Equation 3.

5 Conclusions and Further Work

Testing indicates the data collection and processing units function as expected. Opportunities for future improvement are also indicated.



Figure 8: Cross-correlation between different C/A codes

The sample size and sensitivity is limited by the quantity the FPGA's internal block RAM. The use of external DDR RAM would allow increased sample size and sensitivity.

The buffer type and serial connection speed prevent real time sampling. The implementation of an asynchronous FIFO and a faster communication protocol would allow a larger sample size and increased sensitivity.

The correlation function from Python's NumPy library is too slow for real time computation, a greater bottleneck even than the RS-232 interface. Furthermore, the correlation values calculated were only accurate at close to 0 μs offset. Speed and accuracy could be improved by using an FFT based algorithm. By using the FPGA's on chip DSP, these computations could be performed in real time.

References

[1] A. Annex, "Global positioning system standard positioning service signal specification," 1995.

[2] A. Thompson, J. Moran, and G. Swenson, "Interferometry and synthesis in radio astronomy," *New York, Wiley-Interscience, 1986, 554 p.*, vol. 1, 1986.

[3] P. Chu, *FPGA prototyping by Verilog examples: Xilinx Spartan-3 version*. Wiley-Interscience, 2011.

[4] R. Gold, "Optimal binary sequences for spread spectrum multiplexing (corresp.)," *Information Theory, IEEE Transactions on*, vol. 13, no. 4, pp. 619–621, 1967.

[5] K. Borre, D. Akos, N. Bertelsen, P. Rinder, and S. Jensen, *A software-defined GPS and Galileo receiver*. Birkhhauser, 2007.

[6] R. Dixon, *Spread spectrum systems: second edition*. John Wiley & Sons, Inc., 1984.

TreeTwist: a spiral grain scanner for increment cores

Dr M. P. Hayes* and N. Pink†

*Department of Electrical and Computer Engineering,
University of Canterbury
Christchurch, NEW ZEALAND.
Email: michael.hayes@canterbury.ac.nz

†School of Forestry,
University of Canterbury
Christchurch, NEW ZEALAND.

Abstract: Foresters are interested in the non-destructive determination of the mechanical properties of timber that can be obtained from standing trees. One property of interest is the radial variation of the spiral grain angle. This paper describes a device that estimates the spiral grain angle using ultrasonic time of flight measurements from increment cores.

Keywords: Increment cores, non-destructive testing timber, ultrasonics

1 Introduction

Foresters are interested in growing high quality trees that are suitable for producing strong and stable timber [1]. Increment cores are a low cost and effective method of non-destructively sampling a standing tree. These are typically 10–16 mm in diameter and are removed with a coring tool, ideally through the centre of the tree. One parameter of interest is the spiral grain angle. Wood fibres form a spiral within a tree and the angle of this spiral with respect to the longitudinal axis of the tree is called the spiral grain angle [2]. Variations in the spiral grain angle lead to twist in timber as it dries. They also lead to variations in the longitudinal stiffness of the timber.

The spiral grain angle can be determined by microscopic examination as the core is sliced [1]. However, this is time consuming and destructive of the core. Recently, a Spiralometer device was developed by Industrial Research Limited, New Zealand, in conjunction with Scion [3]. This instrument provides a measure of relative grain angle at 1 mm intervals along the length of a 10 mm increment core [4]. It uses the relative intensity of laser transmission through the core; maximum transmission occurs when the tracheids are coincident with the laser beam.

In this paper, a different device for measuring the spiral grain angle is presented. This also uses increment cores but measures the ultrasonic time of flight of stress waves through the cores. In addition to estimating the spiral grain angle, the radial variation in stiffness parameters can be estimated if the radial density distribution is measured.

We start by presenting a short background to ultrasonic stiffness measurement of timber in Section 2. The device we call TreeTwist is described in Section 3. Preliminary results obtained with this device are shown in Section 4 and discussed in Section 5. Finally, we conclude in Section 6.

2 Background

The spiral grain angle can be determined by measuring the time for a stress-wave to propagate across an increment core. The speed of propagation depends on the mechanical stiffness and the density of the wood. The relation is not straightforward, since the propagation of stress waves in timber is complicated by the anisotropic and inhomogeneous nature of wood [5].

The stiffness of wood depends on the structure and orientation of its cell walls. Thus, in theory, the spiral grain angle can be found by rotating the increment core to find the angle where the fastest propagation speed is measured, since this corresponds to the angle of the fibres. However, this is not entirely accurate since the fastest speed occurs along the direction of the cellulose microfibrils in the S2 layer of the cell walls and these are rotated with respect to the wood fibres by the microfibril angle (MFA) [6].

A number of techniques have been proposed to solve the inverse problem of estimating the stiffness parameters from group speed measurements over a range of angles [7–9]. Unfortunately, there is no closed-form solution for determining stiffness parameters from group speed measurements (even by approximating wood with orthotropic symmetry to reduce the number stiffness parameters to 9

from the 21 required to model full anisotropy). A common approach to estimate the stiffness parameters is to minimise the squared error between the measured and estimated group speeds. An advantage of measuring the group speed over a number of azimuth angles is that it makes the estimates more robust to local inhomogeneities that can distort the stress wave propagation times.

The forward problem of estimating the group speed given a stiffness model is not straightforward either. Given an angle of propagation, the phase speed can be computed by solving the Kelvin-Christoffel equations. Unfortunately, there is no direct solution for finding the group speed and a search is required; the details are beyond the scope of this paper.

3 TreeTwist

A block diagram of the TreeTwist system is shown in Figure 1. It consists of two main sub-systems; the actuator controller sub-system and the ultrasonic measurement sub-system. These are interfaced to a PC using USB.

Positioning of the transducers is performed by three SMC actuators: a slider table, a rotary table, and a gripper. The ultrasonic transducers are attached to the gripper and are positioned along the increment core by the slider table. The increment core is attached to the rotary table.

Rather than using the expensive proprietary SMC controllers, the actuators were interfaced to a pair of low-cost Arcus PMX-2ED-SA two axis combined motor controller and driver modules. The Arcus modules have inputs for the quadrature encoders and limit switches.

The ultrasonic measurement system is based around the Treetap 7c system interfaced to a pair of Fakopp 90 kHz piezoelectric transducers[1]. The Treetap 7c system is a custom made device designed at the University of Canterbury for stress-wave measurements in wood. It has been used in a number of projects, including the Treedisk machine [10] used for scanning tree disks. It consists of a high-voltage pulse generator for generating stress waves using one of the transducers. It also has two channels with programmable gain control for sampling the stress waves measured by receiver transducers. Power is supplied via the USB; this is also used for configuration and control.

Software to control Treedisk was written in Python to run under Linux using a number of classes to control each sub-system. The graphical user interface was generated using Glade[2]. Numerical calculations were performed using SciPy[3] and NumPy[4] and symbolic calculations were performed with SymPy[5] computer algebra system.

[1] `www.fakopp.com`
[2] `glade.gnome.org`
[3] `www.scipy.org`
[4] `numpy.scipy.org`
[5] `sympy.org`

A photo of the prototype TreeTwist system with an increment core is shown in Figure 2. An increment core sample is loaded into the device by hand, by placing it into the long recess and moving the sliding end of the device up to the sample. Two sharp spikes hold the sample in a chuck mounted on the rotary actuator, aligned by sliding phosphor-bronze collars at each end.

An initial calibration phase homes the actuators and then the increment core length is measured by stepping the slider actuator until a limit switch is tripped. A scan then consists of rotating the increment core to the next angular position and stepping the ultrasonic measurement transducers along the length of the increment core. The sample spacing is limited by the size of the transducer tips. Currently, the limitation is 1 mm steps every 6° in angle without the transducer tips homing into the previous indentation.

For each scan position, the transducers are inserted into the increment core using the gripper actuator. A constant force is applied (limited by the stepper motor current) and the resulting transducer separation is measured by counting the quadrature encoder pulses. A number of ultrasonic pulses are generated by the Treetap unit and the resulting stress waves are recorded and averaged to improve the signal to noise ratio.

There are three modes of operation. The fastest mode simply applies a constant force to the ultrasonic measurement transducers and measures the resulting transducer separation. The penetration distance then has to be estimated assuming a constant diameter of the increment core. However, as the increment core dries it shrinks non-uniformly and the diameter is not constant. There are also variations in the thickness produced by the increment corer. For more accurate estimation of the penetration, the transducers can be moved in steps while measuring the ultrasonic response to determine the point of initial contact. A constant force is then applied and the penetration is calculated from the difference in the transducer separation. This mode is considerably slower. A third mode of operation is to estimate the thickness as before but not to penetrate the increment core. With this mode the speed estimates are poorer since the signals are weaker but holes are not made in the increment core.

4 Results

To determine the validity of the TreeTwist system, the stiffness parameters for *Pinus radiata* were calculated using Cave's model [6, 11], with a microfibril dispersion standard deviation of 10 degrees. This model associates most of the stiffness to the S2 layer but underestimates the transverse stiffness provided by the S1 and S3 layers [12].

The group speeds as a function of rotation angle were then calculated from the stiffness parameters and are plotted in Figure 3. This shows the group speed of the longi-



Figure 1: TreeTwist block diagram. Tx and Rx denote the transmitter and receiver piezoelectric transducers.



Figure 2: Photo of the TreeTwist system with an increment core. A collar supports the increment cores close to the ultrasonic measurement transducers. The rotary actuator is at the centre right and the linear actuator is behind it. The black box is Treetap 7c; this contains the high voltage pulse generator and waveform recording electronics. The Arcus stepper motor controllers and power supply are mounted under the slider table.



Figure 3: Simulated *Pinus radiata* group speed variation with azimuth angle (the grain angle is at 0°) and microfibril angle (MFA). The estimated results for an empirical parametric model are superimposed.

tudinal stress wave estimated at 6 degree increments of azimuth angle. Superimposed are the estimated group speeds using a simple empirical model based on four parameters found using an optimisation procedure. For a small MFA the fit is excellent and degrades when the MFA is 45 degrees. As expected, the along-grain speeds reduce with increasing MFA and the fastest speeds are not along the grain, due to the opposite orientation of the microfibrils in neighbouring cell walls.

Data measured using TreeTwist on a 12 mm increment core are shown in Figure 4. The measured delay, transducer tip separation, and estimated group speed are shown as functions of the scan position and rotation angle. There is a clear variation of the speed with rotation angle as expected. The speed also varies along the scan and is slowest at the pith, estimated to be at position 80 mm.

The empirical model was then fitted to the measured speed data and the results are shown in Figure 5. The estimated spiral grain angle is relative to the orientation of the increment core in the scanner. To obtain an absolute value would require accurate marking of the core when it was extracted from the tree. However, the important measurement is the how the spiral grain angle varies radially across the tree. There are two obvious artefacts at around scan position 160 mm. This is corroborated by the model fitting error shown in Figure 5(b). The error results from a failure in the estimation of the group speed from the measured waveforms. This could be caused by a local inhomogeneity resulting in a much weaker signal.

Figure 5(c) shows two of the estimated stiffnesses; the longitudinal stiffness C_{33} and the tangential stiffness C_{22}. These have been estimated on the false premise that the density is unchanging along the increment core. In practice, the density varies radially across the tree and between



(a)



(b)



(c)

Figure 4: Raw TreeTwist scan data (sample T-11). Step size 1 mm, angle increment 6°. (a) measured time delay, (b) transducer separation, (c) estimated speed of propagation.



(a)



(b)



(c)

Figure 5: Derived results from TreeTwist scan (sample T-11). Step size 1 mm, angle increment 6°, $\rho = 450\,\mathrm{kg/m^3}$. (a) estimated spiral grain angle (SGA), (b) model fitting error, (c) estimated longitudinal and tangential stiffnesses assuming a constant density.

early-wood and late-wood of each growing season.

Although the absolute stiffnesses are unknown, the relative stiffness estimates are meaningful. For example, the pith can be seen to occur at around position 80 mm where the anisotropy is smallest. The asymmetry of the estimated stiffness with position is probably due to a lean of the tree when it was young. It appears that there is compression wood on the left hand side; not only is the stiffness lower but the transducer separation distance is larger, indicative of a higher density.

5 Discussion

The increment scanner can accurately determine the variation in spiral grain angle without knowledge of the wood density. The precision depends on the signal to noise ratio and the number of samples. Fitting a model to the measured data helps to make the spiral grain estimate more robust to local inhomogeneities that can distort the stress wave propagation times.

The relative longitudinal and tangential wood stiffnesses can be estimated from the TreeTwist measurements. However, due to the scanning geometry, not all the wood stiffness parameters can be estimated. Furthermore, estimation of the absolute wood stiffness parameters requires knowledge of the wood density. While this can be measured from X-ray absorption, ideally it would be estimated by the core scanner to ensure accurate registration. Initial empirical observations indicate there is some correlation of density with probe penetration but this needs to be quantified. For accurate estimation of the penetration depth it is desirable for the increment cores to have a consistent thickness. However, as they dry they shrink non-uniformly and warp. The thickness can be estimated by observing the ultrasonic waveforms as the transducers are lowered toward the increment core. However, this slows the scanning and currently the TreeTwist machine is being modified to measure the core thickness mechanically.

6 Conclusions

A device for estimating the spiral grain angle and wood stiffness from tree increment cores has been presented. Initial results show promise for the characterisation of wood properties. Further work is required to estimate density, possibly from penetration measurements.

7 Acknowledgements

The authors would like to thank Prof. John Walker of the School of Forestry at the University of Canterbury, for his support and encouragement, and students Benjamin-Washington Yule and Karl Barrett for their contributions to the hardware and software.

References

[1] C. Wessels, F. Malan, and T. Rypstra, "A review of measurement methods used on standing trees for the prediction of some mechanical properties of timber," *European Journal of Forest Research*, vol. 130, no. 6, pp. 881–893, 2011.

[2] J. Harris, *Spiral grain and wave phenomena in wood formation.* Springer-Verlag, 1989.

[3] M. Riddell, D. Cown, J. Harrington, J. Lee, and J. Moore, "Assessing spiral grain angle by light transmission–proof of concept," *IAWA Journal-International Association of Wood Anatomists*, vol. 33, no. 1, p. 1, 2012.

[4] B. Brashaw, V. Bucur, F. Divos, R. Gonçalves, J. Lu, R. Meder, R. Pellerin, S. Potter, R. Ross, X. Wang, *et al.*, "A worldwide research update," *Forest Products Journal*, vol. 59, no. 3, p. 7, 2009.

[5] V. Bucur, *Acoustics of wood.* CRC Press, 1995.

[6] I. D. Cave, "The anisotropic elasticity of the plant cell wall," *Wood Sci. Technol.*, vol. 2, pp. 268–278, 1968.

[7] B. Audoin and C. Bescond, "Measurement by LASER-generated ultrasound of four stiffness coefficients of an anisotropic material at elevated temperatures," *J. Nondestructive Evaluation*, vol. 16, no. 2, pp. 91–100, 1997.

[8] D. Lévesque, N. Legros, and A. Ajji, "Ultrasonic determination of mechanical modulii of oriented semicrystalline polymers," *Polymer Engineering and Science*, vol. 37, pp. 1833–1844, Nov. 1997.

[9] A. G. Every and W. Sachse, "Determination of the elastic constants of anisotropic solids from acoustic-wave group-velocity measurements," *Physical Review B*, vol. 42, pp. 8196–8205, Nov. 1990.

[10] S. Knopp and M. Hayes, "Tree disk machine," in *ENZCON'10, Electronics Conf. New Zealand*, (Hamilton, New Zealand), 2010.

[11] I. D. Cave, "The longitudinal Young's modulus of *Pinus Radiata*," *Wood Sci. Technol.*, vol. 3, pp. 40–48, 1969.

[12] H. Qing and J. Leon Mishnaevsky, "3D hierarchical computational model of wood as a cellular material with fibril reinforced, homogeneous multiple layers," *Mechanics of Materials*, vol. 41, pp. 1034–1049, 2009.

Compressive Hyperspectral Image Sensing Restoration via Joint TV-L1 Regularisation

William Y.L. Lee*, Duncan Eason‡, Mark Andrews†

Department of Electrical and Computer Engineering
University of Auckland, Auckland, New Zealand
Email: *william.lee@auckland.ac.nz, ‡duncan.eason@auckland.ac.nz, †m.andrews@auckland.ac.nz

Abstract: Compressive sensing is a recently introduced signal acquisition technique that can significantly reduce the number of measurements necessary for reconstructing a signal to a high degree of accuracy. This has considerable practical benefits to hyperspectral imaging where the current cost and effort required to collect the vast amount of data hinders its applicability to numerous field of sciences. Current state-of-the-art algorithms for compressive sensing are designed for solving problems with single regularisers (such as the sparse-inducing L_1-norm and total variation). This paper presents a method for solving compressive hyperspectral image sensing problem using joint L_1-norm and total variation regularisation as an on-going effort in improving compressive hyperspectral image restoration. Preliminary results suggest that such regularisation gives better results than using the individual regularisers alone. Current work in progress includes incorporating spatial-spectral regularisation for hyperspectral image recovery and analysing new compressive sampling schemes to exploit such regularisation.

Keywords: Compressive sensing, hyperspectral imaging, image restoration, total variation, regularisation.

I. INTRODUCTION

Hyperspectral imagery is an increasingly popular imaging technique that captures the spatial information of a scene across a large number of wavelength bands, spanning across the visible and infrared spectra. By exploiting the intrinsic spectral signature of the objects within these images, hyperspectral imagery has shown great promise in image enhancement, object recognition and target detection [1] [2]. Because of its non-contact and non-destructive nature, hyperspectral imaging has great potential in numerous fields such as food safety, medical sciences, forensics etc. where contaminations from external sources could potentially have a significant impact to what follows [3] [4]. However, due to its current complexity, cost and the effort required to collect the vast amount of data, most existing hyperspectral imaging applications are limited to mapping terrestrial resources via satellites equipped with hyperspectral camera systems where the economic benefits are more significant.

Compressive sensing is a recently emerged signal acquisition technique that enables a signal or image to be accurately recovered, in its compressed form, from relatively few random linear projections of the signal. Provided that the signal is *sparse* (contains only a few non-zero coefficients) in a certain transform basis, the number of measurements required can be far fewer than that imposed by the Shannon-Nyquist sampling theorem. This opens up the possibility of simpler, cheaper and more efficient alternatives than existing solutions. Consequently, compressive sensing has received increasing attention in applications with very large, high dimensional datasets such as that of hyperspectral imagery where the current cost and time necessary to capture the data outweighs its practical use.

Since the number of collected measurements is well below the signal dimension itself, the process of recovering the signal from its compressive measurements becomes an ill-posed linear inverse problem. Consequently, most approaches tackle the linear inverse problem that arises from compressive sensing by finding a solution $\hat{x}$ that minimizes the regularised objective:

$$\hat{x} = \arg\min_x \frac{1}{2}\|y - \Phi x\|_2^2 + \tau\|\Psi^{-1}x\|_1 \qquad (1)$$

where x is the $N \times 1$ signal to be recovered, y is the $M \times 1$ measurement vector, Φ is a measurement matrix of dimension $M \times N$, $\|\cdot\|_1$ (L1-norm) is the *sparse inducing regulariser*, Ψ^{-1} is the basis transformation in which the signal is sparse in and τ is the *regularisation parameter*. In many compressive sensing image restoration problems, the total-variation (TV) regulariser has also proven to be successful:

$$\hat{x} = \arg\min_x \frac{1}{2}\|y - \Phi x\|_2^2 + \tau\, TV(x) \qquad (2)$$

where $TV(x)$ is the *piece-wise smoothness inducing regulariser*. Current state-of-the-art performances in solving compressive sensing problems have been obtained with algorithms that solve (1) or (2) efficiently [5-11].

Inspired by the success of (1) and (2), we consider incorporating both regularisers into a single joint TV-L1 objective for compressive sensing hyperspectral image restoration:

$$\hat{x} = \arg\min_x \frac{1}{2}\|y - \Phi x\|_2^2 + \tau\|\Psi^{-1}x\|_1 + \lambda\, TV(x) \qquad (3)$$

The underlying rationale for such regularisation is that natural images are known to be sparse under a variety of transforms (the well-known example is the Wavelet transform) and, simultaneously, piece-wise smooth in the image domain. These two characteristics cannot be favoured by either (1) or (2) alone and hence the desirability to incorporate them jointly.

The proposed approach is an instance of the so-called *alternating direction method of multipliers (ADMM)* [12], which allows the decomposition of the objective in (3) into smaller sub-problems in which the convergence has been proved. Furthermore, our particular formulation allows the sub-problems to be solved in parallel; this is highly desirable due to the complexity and size of hyper-spectral datasets.

It is worth mentioning that the problem (3) above is not completely new and has been suggested in the past. In [13] [14], the authors have expressed the desirability of combining L1 and TV regularizers in the context of magnetic resonance image reconstruction and compressive sensing image restoration, however, no algorithm is proposed to solve it. In [15], an algorithm to solve linear inverse problems with 2 arbitrary regularisers has been proposed in the context of image de-blurring, but not in the framework of ADMM and the convergence of the algorithm is not guaranteed. More recently, [16] proposed solving (3) in the context of fluorescence molecular tomography; however, the algorithm therein solves (3) indirectly via a modified version of the objective and the method described does not allow parallelisation.

This paper is divided as follows: §II briefly introduces the concepts of compressive sensing necessary for the remaining of the paper and §III overviews the standard ADMM algorithm. Our formulation of (3) for compressive hyperspectral image sensing restoration in the framework of ADMM is then presented in §IV along with some preliminary results. Finally, current work in progress is described in §V and the conclusions are drawn.

II. COMPRESSIVE SENSING

The conventional approach to reconstructing any signal from its measured data requires the sampling rate to be twice the bandwidth present in the signal, or the so-called *Nyquist rate*. However, it is becoming apparent that more often than not, most of the signals can be stored and represented accurately by only a small fraction of the data size implied by the Nyquist criterion. This is achieved via *transform coding*, a signal compression technique that exploits the property of sparsity. Compressive sensing (CS) was introduced as an extension to the idea of transform coding in that instead of exploiting sparsity as a post-processing technique to retain the sparse coefficients of the signal (thus reducing the data size), CS exploits sparsity to achieve direct acquisition of these sparse coefficients at sub-Nyquist rates [17].

Consider an image with N pixels represented by the vector x, and suppose that x can be represented by just $K \ll N$ coefficients under a certain linear transform $\alpha = \Psi^{-1}x$. CS aims to make just M measurements $y(= \Phi x)$, such that $K \leq M \ll N$ and recover the transformed version of the image, α, via recovery algorithms that solves, for example, (1) or (2). It has been shown that provided the measurement matrix Φ, satisfies the so-called *Restricted Isometry Property* (RIP) [18] then it is a sufficient condition for the recovery of α (and hence x). Matrices that satisfy the RIP include those that generate



Figure 1. Comparison between the traditional sampling scheme and compressive sensing. An image is sampled at the Nyquist rate, after which only the largest coefficients in the wavelet domain are kept. Compressive sensing directly estimates these largest coefficients at sub-Nyquist rate.

each element of Φ from an independent and identically distributed Gaussian random variable, or that take random rows from an orthogonal transform matrix (such as the Discrete Fourier Transform matrix) as the rows of Φ (see [19] for a comprehensive overview of CS). An illustration of the relationship between transform coding and CS is shown in Fig. 1.

III. ALTERNATING DIRECTION METHOD OF MULTIPLIERS (ADMM)

The alternating direction method of multipliers is an algorithm that solves problems in the form:

$$\min f(x) + g(z) \quad subject\ to\ \ Ax + Bz = c \tag{4}$$

with variables $x \in R^n$ and $z \in R^m$, where $A \in R^{p\times n}, B \in R^{p\times m}, c \in R^p$. The algorithm works by first forming the *augmented Lagrangian* of (4):

$$L_\rho(x,z,w) \triangleq f(x) + g(z) + w^T(Ax + Bz - c) + \frac{\rho}{2}\|Ax + Bz - c\|_2^2 \tag{5}$$

where w denotes the associated Lagrange multipliers and $\rho > 0$ is a constant. ADMM then proceeds as shown in Table 1.

Standard ADMM Algorithm
1. Set $k = 0$, and choose $\rho > 0$, z_0 and w_0
2. While stopping criterion does not hold
3. $x_{k+1} = \arg\min_x L_\rho(x, z_k, w_k)$
4. $z_{k+1} = \arg\min_z L_\rho(x_{k+1}, z, w_k)$
5. $w_{k+1} = w_k + \rho(Ax_{k+1} + Bz_{k+1} - c)$

Table 1. The standard ADMM algorithm.

Conveniently, the last two terms in (5) can be written as a single quadratic (plus a constant independent of x and z) via completing the squares:

$$L_\rho(x,w,u) \triangleq f(x) + g(z) + \frac{\rho}{2}\|Ax + Bz - c + u\|_2^2 - \frac{\rho}{2}\|u\|_2^2 \tag{6}$$

where $u = \frac{1}{\rho} w$. This gives rise to a more compact form of the ADMM algorithm, shown in Table 2, which this paper adopts.

ADMM (Scaled dual form) Algorithm
1. Set $k = 0$, and choose $\rho > 0$, z_0 and u_0
2. While stopping criteria does not hold
3. $x_{k+1} = \arg\min_x f(x) + \frac{\rho}{2}\|Ax + Bz_k - c + u_k\|_2^2$
4. $z_{k+1} = \arg\min_z g(z) + \frac{\rho}{2}\|Ax_{k+1} + Bz - c + u_k\|_2^2$
5. $u_{k+1} = u_k + (Ax_{k+1} + Bz_{k+1} - c)$

Table 2. The scaled dual form of the standard ADMM algorithm.

It is also recommended to update the quadratic penalty parameter ρ after each iteration in order to improve convergence [12][20]. A heuristic for such an update that works well in general is given in [12]. Finally, the convergence of the ADMM algorithm above has been shown in [12][21].

IV. PROPOSED APPROACH

Assuming each band of the hyperspectral image is sampled under the same measurement matrix Φ, the resulting measurements can be written as:

$$Y = [y_1, \dots, y_p] = \Phi[x_1, \dots, x_p] = \Phi X \tag{7}$$

where p is the total number of bands in the hyperspectral imagery.

A. Algorithm Formulation

By penalising each individual band separately, (3) can be re-written as:

$$\hat{X} = \arg\min_X \frac{1}{2}\|Y - \Phi X\|_F^2 + \tau \sum_{i=1}^{p} \|\Psi^{-1} x_i\|_1 + \lambda \sum_{i=1}^{p} \|D x_i\|_1 \tag{8}$$

where $\|\cdot\|_F$ is the Frobenius norm and $TV(x)$ is taken as the anisotropic total variation defined by $TV(x) = \left\|\begin{bmatrix} D_v \\ D_h \end{bmatrix} x\right\|_1$ where D_v and D_h are the first-order horizontal and vertical difference matrices respectively. Inspired by [22], the minimization problem (8) can be translated to (4) via the mapping:

$$f(x) = 0, Z = \begin{bmatrix} Z_1 \\ Z_2 \\ Z_3 \end{bmatrix} = \begin{bmatrix} \Phi \\ \Psi^{-1} \\ D \end{bmatrix} X = AX \Rightarrow AX - Z = 0$$

and $g(z)$ given by:

$$g(z) = \frac{1}{2}\|Y - Z_1\|_F^2 + \tau \sum_{i=1}^{p} \|(Z_2)_i\|_1 + \lambda \sum_{i=1}^{p} \|(Z_3)_i\|_1$$

where $(Z_j)_i$ is the ith column of the matrix Z_j. With these definitions, assuming Ψ^{-1} is an orthogonal transform i.e., $\Psi\Psi^{-1} = \Psi\Psi^T = I$ and A has full column rank, Step 3 of the ADMM algorithm reduces to solving the equation:

$$(\Phi^T\Phi + D^T D + I)X_{k+1} = \left(\Phi^T(Z_{1_k} - U_{1_k}) + \Psi(Z_{2_k} - U_{2_k}) + D^T(Z_{3_k} - U_{3_k})\right) \tag{9}$$

where U_1, U_2, U_3 are matrices containing the associated scaled Lagrange multipliers. Step 4 of the algorithm can be decoupled into 3 independent steps that can be solved simultaneously:

$$Z_{1_{k+1}} = \arg\min_z \frac{1}{2}\|Y - Z\|_F^2 + \frac{\rho}{2}\left\|\Phi X_{k+1} - Z + U_{1_k}\right\|_F^2 \tag{10}$$

$$Z_{2_{k+1}} = \arg\min_z \tau \sum_{i=1}^{p} \|(Z)_i\|_1 + \frac{\rho}{2}\left\|\Psi^{-1} X_{k+1} - Z + U_{2_k}\right\|_F^2 \tag{11}$$

$$Z_{3_{k+1}} = \arg\min_z \lambda \sum_{i=1}^{p} \|(Z)_i\|_1 + \frac{\rho}{2}\left\|D X_{k+1} - Z + U_{3_k}\right\|_F^2 \tag{12}$$

Furthermore, $Z_{1_{k+1}}$, $Z_{2_{k+1}}$ and $Z_{3_{k+1}}$ has the closed form solution of:

$$Z_{1_{k+1}} = \frac{1}{1+\rho}\left(Y + \rho\left(\Phi X_{k+1} + U_{1_k}\right)\right) \tag{13}$$

$$Z_{2_{k+1}} = \mathrm{SoftThresh}\left(\left(\Psi^{-1} X_{k+1} + U_{2_k}\right), \frac{\tau}{\rho}\right) \tag{14}$$

$$Z_{3_{k+1}} = \mathrm{SoftThresh}\left(\left(D X_{k+1} + U_{3_k}\right), \frac{\lambda}{\rho}\right) \tag{15}$$

where $\mathrm{SoftThresh}(x, t) = \mathrm{sgn}(x)\max\left((|x| - t), 0\right)$ is the well-known *soft-thresholding operator* [23] applied element-wise to the input matrix x. The resulting algorithm is shown in Table 3. In this form, it is clear that the resulting algorithm follows a *decomposition-coordination* procedure: the original problem is decomposed into smaller sub-problems (Step 4) in which the solutions can be found in parallel and are coordinated (via Step 3) to find the solution to the original problem.

CIRTV-L1R(Compressive Image Restoration via TV-L1 Regularization)
1. Set $k = 0$, and choose $\rho, \tau, \lambda > 0$, $Z_0 = 0, U_0 = 0$
2. While stopping criterion does not hold
3. $X_{k+1} = (\Phi^T\Phi + D^T D + I)^{-1}\left(A^T(Z_k - U_k)\right)$
4. a. $Z_{1_{k+1}} = \frac{1}{1+\rho}\left(Y + \rho\left(\Phi X_{k+1} + U_{1_k}\right)\right)$
b. $Z_{2_{k+1}} = \mathrm{SoftThresh}\left(\left((\Psi^{-1} X_{k+1} + U_{2_k}\right), \frac{\tau}{\rho}\right)$
c. $Z_{3_{k+1}} = \mathrm{SoftThresh}\left(\left(D X_{k+1} + U_{3_k}\right), \frac{\lambda}{\rho}\right)$
5. $U_{k+1} = U_k + (AX_{k+1} - Z_{k+1})$

Table 3. The proposed algorithm



Figure 2. Visual comparison between the different recovery methods. Each row shows a particular band of the hyperspectral image recovered from CIRTV-L1R. (a) Ground truth; (b) Minimum energy reconstruction (i.e., no regularizer) ; (c) Reconstruction from L1 regularization ; (d) Reconstruction from TV regularization; (e) Reconstruction from joint TV-L1 regularization. Note that the restored images are inverted to provide a clearer visual comparison of the reconstruction quality.

B. Compressive Hyperspectral Image Sensing Restoration

To demonstrate the benefit of joint TV-L1 regularization, several 256×256 test images consisting of a group-sparse structure (composed of randomly placed white blocks in a black background) were randomly created to simulate the different bands of a hyperspectral image, Fig. 2a. shows two particular bands. By doing so, the sparcifying transform for the sparsity inducing regulariser is simply the identity, i.e., $\Psi^{-1} = I$. The compressive sensing measurement process is simulated by taking Φ as a random partial Fourier matrix (see §II) i.e., $\Phi = PF$ where P is a random partial permutation matrix and F is the discrete Fourier matrix. Since $(\Phi^T\Phi + D^TD + I)$ is symmetric and positive definite, (9) (i.e., Step 3 of the algorithm) is solved using the conjugate gradient (CG) method. It is also worth mentioning that (9) need not be solved exactly, the convergence still holds provided the error sequences between the true and the approximate solution of (9) are absolutely summable [21]. In our experiments, solving (9) took no more than 6 CG iterations. Finally, the algorithm is terminated when the relative change in the objective value falls below a certain tolerance.

A visual comparison of the reconstruction results, from noisy ($\sigma = 0.1$) compressive measurements with an under-sampling ratio of $\frac{M}{N} \simeq 0.09$, between no regulariser, single L1/TV regulariser and joint TV-L1 regularisers is shown in Fig. 2b-2e. The regularisation parameters for each of the case are chosen experimentally to give the minimal mean square error. Fig. 3 shows the mean square error plot against the relative deviation from the optimal regularisation parameter for each of the regularised case above. From the plot, it is apparent that the joint TV-L1 regularisation achieved better results, quantitatively, than a single regularizer alone. Lastly, it is worth noting that the average reconstruction time per band took no longer than 7 seconds when performed under MATLAB v7.13 (R2011b) running on an Desktop PC with an Intel Core2 Quad CPU at 2.6 GHz and 4 GB of RAM. This is comparable with most state-of-the-art algorithms that solve (1) and (2). See [11] for a recent comparison of the state-of-the-art algorithms for solving (1) and (2).



Figure 3. Mean square error plot against the relative deviation from the optimal regularization parameter.

V. WORK IN PROGRESS

The proposed approach restores individual bands of the hyperspectral image simultaneously. However, restorations occur independently as the regularisers are applied spatially across each band without any consideration for the structural redundancies along the spectral dimension. It is a well-known fact that the spectral signature of most real-life objects exhibit a smooth response, therefore it is

hoped that exploiting such structure could potentially lead to improved restoration. This could be as simple as incorporating yet another regulariser into (8):

$$\hat{X} = \arg\min_{X} \frac{1}{2}\|Y - \Phi X\|_F^2 + \tau \sum_{i=1}^{p} \|\Psi^{-1} x_i\|_1 + \lambda \sum_{i=1}^{p} TV(x_i) + \gamma \sum_{i=1}^{N} \left\|\Psi_\lambda^{-1}(X^T)_i\right\|_1 \quad (16)$$

where Ψ_λ^{-1} is a transform which sparsifies the hyperspectral image along the spectral dimension, $(X^T)_i$ is the ith column of X^T, i.e., the spectral signature of the ith pixel and γ is the associated regularisation parameter. Furthermore, incorporating this *spectrally sparse inducing* regulariser, may entail a further reduction in the number of measurements necessary, for example, instead of sampling each and every band of the hyperspectral image, only a random subset of the bands are sampled, that is:

$$Y = \Phi X P_\lambda^T \quad (17)$$

where P_λ is another random partial permutation matrix. Preliminary results suggest that it is still possible to recover the hyperspectral cube even when only 40% of the bands are sampled. Current work in progress focuses on incorporating (16) and (17) into the framework of ADMM and further analysis of the sampling scheme (17) and other alternatives.

VI. CONCLUSIONS

A method for joint TV-L1 regularisation for compressive hyperspectral image sensing restoration using the alternating direction method of multiplier has been presented. Preliminary results suggest that the joint TV-L1 regularisation gives better reconstruction from compressive measurements, both visually and quantitatively, than using individual regularisers alone. Current and future work includes improving the reconstruction quality through joint spatial-spectral regularisation and analysis of new sampling schemes to exploit such regularisation.

VII. REFERENCES

[1] H. Liang, "Advances in Multispectral and Hyperspectral Imaging for Archaeology and Art Conservation," *Applied Physics A*, vol. 106, no. 2, pp. 309-323, 2012.

[2] D. Manolakis, D. Marden and G. A. Shaw, "Hyperspectral Image Processing for Automatic Target Detection Applications," *MIT Lincoln Laboratory Journal*, vol. 14, pp. 79-116, 2003.

[3] O. Carrasco, R. B. Gomez, A. Chainani and W. E. Roper, "Hyperspectral Imaging Applied to Medical Diagnoses and Food Safety," in *SPIE*, 2003.

[4] S. Dowler, "Applications of Hyperspectral Imaging Techniques to Forensic Image Analysis," Ph.D. Dissertation, The University of Auckland, 2010.

[5] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183-202, 2009.

[6] J. Bioucas-Dias and M. Figueiredo, "A New TwIST: Two-Step Iterative Shrinkage/Thresholding Algorithms for Image Restoration," *IEEE Transactions on Image Processing*, vol. 16, no. 12, pp. 2992-3004, 2007.

[7] S. Wright, R. Nowak and M. Figueiredo, "Sparse Reconstruction by Separable Approximation," *IEEE Transactions on Signal Processing*, vol. 57, no. 7, pp. 2479-2493, 2009.

[8] M. Figueiredo, R. Nowak and S. Wright, "Gradient Projection for Sparse Reconstruction: Application to Compressed Sensing and Other Inverse Problems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 4, pp. 586-597, 2007.

[9] M. Afonso, J. Bioucas-Dias and M. Figueiredo, "Fast Image Recovery Using Variable Splitting and Constrained Optimization," *IEEE Transactions on Image Processing*, vol. 19, no. 9, pp. 2345-2356, 2010.

[10] C. Li, "An Efficient Algorithm For Total Variation Regularization with Applications to the Single Pixel Camera and Compressive Sensing," M.S. thesis, Rice Univ., Houston, TX, 2009.

[11] D. Eason, W. Lee and M. Andrews, "Compressively Sensed Hyperspectral Image Recovery using Total Variation Minimization by Approximation," *submitted to Electronics New Zealand Conference (ENZCon)*, 2012.

[12] S. Boyd, N. Parikh, E. Chu, B. Peleato and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1-122, 2011.

[13] M. Lustig, D. Donoho and J. Pauly, "Sparse MRI: The Application of Compressed Sensing for Rapid MR Imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182-1195, 2007.

[14] E. Candes and J. Romberg, "Practical Signal Recovery from Random Projections," in *Proc. of the SPIE International Symposium on Electronic Imaging: Computational Imaging III*, 2005.

[15] J. Bioucas-Dias and M. Figueiredo, "An Iterative Algorithm for Linear Inverse Problems with Compound Regularizers," in *Proc. of the 15th IEEE International Conference on Image Processing*, 2008.

[16] J. Dutta, S. Ahn, C. Li, S. Cherry and R. Leahy, "Joint L1 and Total Variation Regularization for Fluorescence Molecular Tomography," *Physics in Medicine and Biology*, vol. 57, no. 6, pp. 1459-1476, 2012.

[17] D. Donoho, "Compressed Sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289-1306, 2006.

[18] E. Candes and T. Tao, "Decoding by Linear Programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203-4215, 2005.

[19] M. Davenport, M. Duarte, Y. Eldar and G. Kutyniok, "Introduction to Compressed Sensing," in *Compressed Sensing: Theory and Applications*, New York, Cambridge University Press, 2011, pp. 1-65.

[20] J. Nocedal and S. Wright, Numerical Optimization, 2nd Edition, Springer, 2006.

[21] J. Eckstein and D. Bertsekas, "On the Douglas-Rachford Splitting Method and the Proximal Point Algorithm for Maximal Monotone Operators," *Mathematical Programming*, vol. 55, pp. 293-318, 1992.

[22] M. Afonso, J. Bioucas-Dias and M. Figueiredo, "An Augmented Lagrangian Approach to the Constrained Optimization Formulation of Imaging Inverse Problems," *IEEE Transactions on Image Processing*, vol. 20, no. 3, pp. 681-695, 2011.

[23] D. Donoho, "De-noising by Soft-Thresholding," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 613-627, 1995.

Cascomp BJT Amplifier vs. Traditional Configurations

Harrisson Jull, Toby Balsom, and Jonathan Scott

University of Waikato
School of Science and Engineering
harrisson.j@hotmail.co.nz

Abstract: All transistor circuits introduce distortion. In Radio Frequency (RF) circuits, the third-order distortion components are the most important. The quest for more linear circuits has become more important with complex-modulation as used in modern cellular phone systems. Quinn's Cascomp Amplifier, first reported in the 1970s, promises ideal linearity and can deliver close to that promise. We review the theory and address the question of why the Cascomp has not replaced other configurations in amplifiers where low distortion is important. Calculations are supported by measurements. A new, alternative variant of the Cascomp topology is introduced and compared with the existing configuration. We assert that the improved linearity comes at such a price in gain that it does not make sense to use the configuration in broadband RF circuits.

Keywords: Cascomp, Compensation, Error-correction, Feed-forward, Intermodulation distortion, Linearity, Transconductance amplifiers.

1 INTRODUCTION

The collector current of a Bipolar Junction Transistor (BJT) obeys the well-known nonlinear equation

$$I_C = I_S\, e^{\frac{V_{BE}}{V_T}}. \tag{1}$$

When a signal is imposed on the base-emitter junction, a nonlinear function of that signal appears in the collector. This results in distortion of the signal. For example, when the input signal is a sinewave, the collector current contains both the original sinewave and its harmonics, and we say that the amplified signal contains "harmonic distortion" [1]. Of more interest to communications engineers is the distortion that arises when the input signal consists of multiple sinewaves. In the case of two sinewave inputs, some algebra shows that the output signal will contain signals at various frequency components as depicted in Figure 1 [1]. The tones whose frequencies lie nearby to the two input frequencies, f_1 and f_2, are the ones considered most objectionable because they cannot easily be filtered out. It is common to measure the extent of the distortion introduced by an amplifier with a measure such as Third-order Intermodulation Distortion ("IM3") or Third-order Intercept (TOI) [1].



Figure 1. Components involved in harmonic and intermodulation distortion (IMD) from two input tones.

There has been considerable effort spent in recent decades to find a configuration of transistors that produces less distortion. The Differential Pair (DP), shown in Figure 2(b), has a tanh transfer function not an exponential one. This is claimed to be more linear than the traditional Common-Emitter (CE) amplifier, shown in Figure 2(a). Tanh is symmetrical, so it yields no second-order components, and thus produces much lower harmonic distortion figures. However, it offers less advantage if one is interested only in the IM3 figure.



Figure 2. (a) Common-emitter, (b) Differential pair and (c) Cascomp configurations discussed in this paper.

2 COMPARISON OF DP AND CE CONFIGURATIONS

To compare the DP to the CE configuration, we consider a two-tone test applied to each amplifier. In each case, we will consider the input peak voltage at the base of the transistors and the output peak current taken at the collector. The measurements were taken this way because, at small signal, the resistance seen by the collector and the voltage supply have no effect on the output current, and thus our comparison is independent of these.

Collector bias currents of the CE and DP amplifiers are set at 5 mA. The Values of R_1 and R_2 are set low enough to mitigate the effects of β variations on the quiescent collector current. The value of R_C was set low enough that compression did not occur until higher voltages. This was done because the TOI is determined when the output is a linear function of the input. Increasing the value of R_C means that compression occurs earlier and the TOI may need to be acquired at lower input voltages.

Figure 7 and the upper half of Table 2 present a comparison between the CE and DP circuits with respect to gain (g_m) and TOI. Simple theory predicts that the gain of the CE will be 192 mS and the (single-ended) gain of the DP half that compared to the CE or 96 mS. The theory does not allow for parasitic resistance in the BJT nor Early effect or β variations, and so yields slightly optimistic numbers compared with the simulations and measurements. Intersil's CA3083 transistor-array chip was used after carefully determining the SPICE parameters using an Agilent E5270B. Simulation and measurement agree well. The SPICE parameters are shown in Table 1.

Table 1. Extracted SPICE parameters for the transistors in intersil's CA3083.

IS (Saturation current)	9.6E-15 A
NF (Forward ideality factor)	1.004
NR (Reverse ideality factor)	1.012
VAF (Forward Early voltage)	54 V
VAR (Reverse Early voltage)	6 V
BF (Forward β)	224
BR (Reverse β)	25
RE (Terminal emitter resistance)	1.2 Ω

With equivalent collector quiescent currents, the DP has ~6 dB less gain than the CE, and 1.4-1.7 dB worse IIP3. The DP also has lower fifth-order and seventh-order intermodulation products than the CE. If the user is interested in THD there is an improvement, but if TOI is the measure, the DP performs worse in both ways.

To summarise, the DP reduces even order nonlinearities and THD in the output, but has slightly worse TOI with half the g_m. For this reason more elegant circuits were designed using feed-forward and feed-back techniques [3] to cancel nonlinearities. One circuit of particular interest is the Cascomp feed-forward amplifier [4]-[6].

3 QUINN'S CASCOMP AMPLIFIER

The Cascomp feed-forward amplifier is shown in theoretical form in Figure 3. The Cascomp configuration used in this paper is shown in Figure 2(c). The source impedance (R_S in Figure 3) of the input voltage source is not shown in the circuits in Figure 2, but is implicit to the voltage inputs applied to these circuits. R_B (Figure 3) is used to compensate for the beta dependant gain caused by R_S. It is not included in Figure 2(c) because it adds complexity to the circuit and is not essential to our analysis. The π configuration of the circuitry below the emitters of the main differential pair in Figure 3 performs exactly as the T configuration at the same nodes in Figure 2(c).

The Cascomp amplifier uses the cascoding stage of the main amplifier to replicate the nonlinearities produced by $\Delta V_{BE} = V_{BE1} - V_{BE2}$ not being equal to zero. The error amplifier, ideally represented in Figure 3 but appearing as the inner DP in Figure 2(c), then senses the replicated voltage $\Delta V'_{BE} = V_{BE3} - V_{BE4}$ and adds a copy of this to the output of the cascaded DP. This ideally perfectly cancels the distortion in signal at the output. For a more thorough explanation of this profound process the reader is referred to the references cited!



Figure 3. Patrick Quinn's Cascomp feed-forward error correction amplifier.

Nonlinear cancellation occurs in the Cascomp when the transconductance of the error amplifier, $g_{ME} = 1/R_M$, where R_M is the emitter degeneration resistance of the main DP. In Quinn's implementation, the value of I_E is suggested to be half that of I_M. This is because if I_E is too low the DP error amplifier causes spurious nonlinearity through its own tanh characteristic. When I_E is too large, the base currents drawn by the inputs of the error correction amplifier become so large that the current in the transistors of the main DP are no longer the same as the currents in the cascading devices [5]. In Figure 3, the collector currents of Q1 and Q3 would no longer be the same.

3.1 Cascomp Considerations

When designing a Cascomp amplifier many values need to be determined to ensure perfect cancellation. Taking into consideration the conditions for compensation, and after some algebra, the transconductance of a single-ended output of the Cascomp becomes

$$g_M = \frac{1}{2R_M}. \tag{2}$$

It is most common to use a degenerated differential pair as the error amplifier. Assuming $\beta \gg 1$ and satisfying the condition $g_{ME} = 1/R_M$,

$$R_E = R_M - \frac{2V_T}{I_E} \tag{3}$$

where R_E is the degeneration resistor in the error amplifier. It is impossible to obtain more than a certain gain for a given value of tail current I_M, because the required value of R_E would be less than zero. In practice, Quinn's theory requires R_M to have a value of at least $1/g_{ME}$, severely limiting the achievable overall g_M. Refer to Figure 4 for a visual representation of this limit.



Figure 4. This plot shows the value of R_E required to achieve a given transconductance in the overall amplifier for various ratios of the two tail currents when I_M=20 mA (■), I_M=10 mA (●), I_M=5 mA (♦). Cascomp transconductance (—).

Using this relationship, suitable resistor values were chosen. Simulation and measurement showed that the actual value needed for the degeneration resistors in the error amplifier was slightly lower than the theoretical value. This is attributed to the transistor's terminal resistance in the emitter leads, as well as added resistances from connections in both tail circuits. Finite beta and early effects also contribute to the deviation from theory.

Also note that the Cascomp requires sufficient degeneration to produce significant compensation (cancellation of distortion products). This null occurs when $R_M = 30\ \Omega$ for the currents used in this paper. A value of 33 Ω was used to produce a g_M close to the maximum g_M attainable, while ensuring a high level of compensation. This also meant that the null would be present with shifts due to component tolerances.

4 COMPARISON OF THE CASCOMP AGAINST CE AND DP

To compare the Cascomp with the CE and DP we consider the same two-tone test in section 2. Table 2 and Figure 7 show the comparison between the Cascomp (single-ended), CE and DP circuits with respect to gain (g_m) and TOI. Theory (Equation 2) predicts that the gain of the Cascomp will be $g_m = 15.2\ mS$. Simulation and measurement agree well with theory.

With equivalent collector quiescent currents, the Cascomp has 14 dB less gain than the DP, and 20 dB less than the CE. There is vast improvement in TOI over both the CE and DP.

The TOI of the Cascomp was taken at an output current of -131 dBA. This was the lowest point measurable owing to the limits of dynamic range in the spectrum analyser. Whilst this is not a perfect indication of the true TOI, the values show dramatic improvement over the CE and DP. There is an improvement in IIP3 of over 38 dB compared to the CE and DP. Figure 7 shows the slope of the third-order product as 110 dB/dec instead of the typical 60 dB/dec. This is due to the fifth-order intermodulation product's contribution to the third's coefficient.

To summarise, the Cascomp reduces Intermodulation Distortion (IMD) and increases the TOI. The downfall of the Cascomp is that the degeneration needed for compensation reduces the gain (g_m) to a tenth of the gain possessed by the CE and a fifth the gain of the DP.

5 ALTERBATIVE CASCOMP CONFIGURATIONS

According to Equation 3 there are multiple points where compensation (distortion nulling) occurs for different tail currents in the error amplifier. This theory holds true only when there is sufficient degeneration in the error amplifier that it approximates the ideal of Quinn's theory. When there is low degeneration, the null formed by compensation does not exist, because the error amplifier contributes its own distortion. This circumstance is not taken into account by the original theory.

Simulation and optimisation have shown that there are a range of currents that have compensation at two resistance values. One of these values is the value predicted by Quinn's theory (Equation 3). The second value is the point where the distortion contribution of the error amplifier itself is taken into account, and it corresponds to a higher tail current value. This increased current also means an increase in the Cascomp's gain. It is possible to think of this second cancellation point as the place where the distortion of the error amplifier is additionally pitted against the nonlinearity of the main amplifier.

To compare these two points simulation was done to find the point at which compensation occurs when there is no degeneration in the error amplifier. This point would bring added gain to the Cascomp which it severely needs. The main amplifier used in the previous sections was kept constant with only the error amplifier being altered. The input voltage was kept at a constant 0.02 V (-34 dBV). Figure 5 and Figure 6 display the third-order intermodulation distortion and fundamental gain when the error amplifier is so modified. The inverted spikes in Figure 5 are due to large simulation steps.



Figure 5. Output level of the Cascomp's third-order intermodulation product as a function of the error amplifiers tail current and degeneration resistance.



Figure 6. Output level of the Cascomp's carrier as a function of the error amplifiers tail current and degeneration resistance.

The theoretical point from the previous sections occurred when $I_E = 5\,mA$ with a 17.6 Ω degeneration resistor. The point where no degeneration was needed occurred when $I_E = 16.4\,mA$. This produced a ~6 dB increase in gain and a ~3 dB decrease in TOI compared with the traditional Cascomp. Figure 7 and Table 2 compare the new design with the previously mentioned amplifiers.

In section 3 it was explained why the current should be half that of the main amplifier. Because of the increase in current the output current will be influenced by the base currents from the error amplifier. This effect can be cancelled by a combination of various techniques including of scaling transistors [7].

Table 2. Third-order intercept and transconductance. Values given are rounded to three significant figures.

Common-emitter	**IIP3 (dBV)**	**Gm (mS)**
Simulated	-15.4	147
Measured	-14.8	146
Differential pair		
Simulated	-16.8	76.5
Measured	-16.1	74.3
Cascomp		
Simulated	>22.7	15.3
Measured	>23.3	15.2
Alternative Cascomp		
Simulated	>19.3	30.7



Figure 7. Output level as a function of input level for fundamental and third-order products for CE, DP and Cascomp amplifiers with similar device operating points. Simulation (—), symbols are measured data. CE fundamental (●), CE third-order product (■), DP fundamental (♦), DP third-order product (▲), Cascomp fundamental (+) and Cascomp third-order product (x).

6 REALISATION IN HBT

Although the Cascomp amplifier allows for significant improvement in TOI compared to other topologies, the gain-bandwidth product of the amplifier is much smaller. The compensation trades off gain for linearity at such a rate that the gain is smaller than the CE and DP. In terms of bandwidth, the first pole of the Cascomp is inherently diminished, compared with the other topologies. If all circuits were moved to a RF Heterojunction Bipolar

Transistor (HBT) technology, we would expect a higher overall gain-bandwidth product, but comparably the Cascomp would still be less effective compared to other topologies. If the Cascomp was to be realised to compete in the RF amplifier market, where linearity and gain-bandwidth are important, we cannot say that we expect it to perform better.

7 CONCLUSIONS

The CE outperforms the DP in gain and TOI but not in THD, Signal to Noise and Distortion (SINAD) ratio and Spurious-Free Dynamic Range (SFDR). If the output signal is filtered, leaving only the carriers with the intermodulation products closest to them, then the CE is superior to the DP.

The Cascomp dominated the CE and DP in TOI, THD, SINAD and SFDR, even when the calculated TOI was extremely underestimated.

With the revelation of a secondary compensation point with no degeneration in the error amplifier, an alternative topology was created for the Cascomp. The new topology offers an increase in gain for a small decrease in TOI. This is a new result that has not appeared before in the open literature.

The gain-bandwidth product of the Cascomp is smaller than the CE and DP. The bandwidth of the Cascomp is inherently diminished due to a larger time constant through its signal path. If all circuits were moved to an RF HBT technology, we would expect a higher overall gain-bandwidth product, but the Cascomp will still be less effective than the other topologies, and is likely to have insufficient gain to allow broadbanding.

8 REFERENCES

[1] Sansen, W, "Distortion in elementary transistor circuits", *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process*, vol. 46, no. 3, pp. 315-325, Mar, 1999.

[2] D. A. Neaman, "Basic BJT amplifiers," in *Microelectronic Circuit Analysis and Design*, 3rd ed., New York, NY: McGraw-Hill Companies, Inc., 2007, ch. 6, sec. 6.2.1, pp. 375-376.

[3] J. Vanderkooy and S. P. Lipshitz, "Feedforward error correction in power amplifiers," in *J. of the Audio Engineering Society*., vol. 28, no. 1/2 pp. 1-16, Jan/Feb, 1980.

[4] P. A. Quinn, "Feed-forward amplifier," U.S. Patent 4 146 844, Mar, 27, 1979.

[5] P. A. Quinn, "A cascode amplifier nonlinearity correction technique," in *IEEE Int. Solid-State Circuits Conference*., New York, NY, 1981, pp. 188-272.

[6] S. Simpkins and W. Gross, "Cascomp feedforward error correction in high speed amplifier design," in *IEEE J. Solid-State Circuits Conference*., vol. sc-18, no. 6 pp. 762-764, Dec, 1983.

[7] K. G. Schlotzhauer *et al.*, "Cascode feed-forward amplifier," U.S. Patent 4 322 688, Mar, 30, 1982.

Mobile SPOT: A Wireless Sensor and Actuator Node for Applications based on Internet of Things

Udayanto Dwi Atmojo, Zoran Salcic, Kevin I-Kai Wang
Embedded Systems Research Group, Department of Electrical and Computer Engineering
The University of Auckland
38 Princes Street, Auckland, New Zealand, 1010
udwi469@aucklanduni.ac.nz, z.salcic@auckland.ac.nz, kevin.wang@auckland.ac.nz

Abstract— **We describe a new mobile wireless sensor and actuator node, called Mobile SPOT that has the ability to sense its immediate physical environment for two purposes, (1) obstacle detection and autonomous navigation and, (2) acquisition of information as a part of a ubiquitous wireless sensor network and communication with the Internet. A Mobile SPOT first integrates elements of two existing technologies, the SunSPOT wireless sensor node and the LEGO Mindstorms robot, and then extends the integrated platform with new sensing and actuating capabilities, as well as very powerful software programming paradigm. This paper focuses on hardware features of the Mobile SPOT with a brief introduction of its software.**

Keywords: SunSPOT, LEGO Mindstorms, Robot, Internet of Things, Wireless Sensor and Actuator Node.

I. INTRODUCTION

Wireless sensor and actuator networks (WSAN) have emerged as the core technology of the future Internet of Things (IoT) [1]. They comprise large numbers of wireless nodes which are capable of sensing of and actuating to their immediate environment. The nodes are typically stationary at fixed locations (fixed nodes), but emerging applications have shown the trend of integrating mobile sensing and actuating nodes in forming more versatile and robust systems. Together, fixed and mobile nodes represent an ideal base for many new applications in the context of future Internet, where individual nodes will be IP addressable and accessible via Internet enabled mechanisms, such as TCP/IP sockets and web services, suited for small and low power computing platforms. In this paper we present a new type of node, called Mobile SPOT, which integrates the Internet enabled SunSPOT [7] wireless node with necessary mechanical movement and manipulation capabilities based on the physical configuration of LEGO Mindstorms robots [3]. This integration provides a powerful, yet low cost and flexible platform for fast prototyping of mobile sensing and actuating nodes.

The first step of the integration was to replace the brain of the LEGO Mindstorms robot, the NXT brick, which is an ARM-7 based embedded platform. The NXT brick has limited capabilities in computing power and interfaces with external sensors/actuators and does not satisfy many requirements of a mobile node. However, LEGO Mindstorms mechanical parts enable easy creation of different physical configurations of mobile nodes and hence it is partially retained in the new node to achieve fast prototyping. On the other hand, SunSPOT [7] represents one of the most advanced wireless sensor nodes, based on a high performance ARM-9 microprocessor. SunSPOT comes with a set of built-in sensors, including a 3-axis accelerometer, a 3-colour light sensor and a temperature sensor, which provide basic sensing capability [8]. SunSPOT can also be easily extended and customised by connecting external sensors and actuators through a number of industrial standard serial buses, such as I^2C bus and UART, and generic digital/analogue IO pins. In addition to the hardware platform, SunSPOT runs Squawk VM [6], which is a small Java Virtual Machine that executes on top of the ARM-9 microprocessor, and as such enables programming in Java directly without having any operating systems. Squawk VM provides a comprehensive APIs to interface with and control of built-in and external sensors and actuators. Further, Squawk incorporates the state of the art 6LoWPAN protocol stack [5] through the built-in 2.4GHz CC2420 radio chip [2], which allows each SunSPOT node to be IP addressable and Internet accessible. Designers can develop application programs easily in Java, or a higher level concurrent programming language, SystemJ [4] , which is particularly suitable for WSAN and the Internet of Things (IoT) applications that involve concurrency, reactivity, sensor fusion, and communications with other parties in a bigger system. Figure 1 describes the layers containing essential supporting elements of the customised Mobile SPOT.



Figure 1. Layers of the customised Mobile SPOT

In addition to all capabilities of the original LEGO Mindstorms with the NXT brick, the Mobile SPOT introduced in this paper is capable of navigating in a physical environment autonomously, sensing its immediate environment for navigation and data acquisition and communicating with other fixed wireless sensor nodes, using the 6LoWPAN protocol. Such nodes can be characterised and referred to as Mobile Wireless Sensor and Actuator Network (MWSAN) nodes. In a Mobile SPOT, external ultrasonic/infrared rangefinders are connected to SunSPOT through the I^2C bus and generic analogue inputs, to sense obstacles in the physical environment.

Together with two Mindstorms motors controlled through digital IO pins, Mobile SPOT can autonomously navigate in a physical environment. Built-in light and temperature sensors can be used to sense the ambient light intensity and temperature. All the sensing and actuating capabilities can be used from standard Java or wrapped within a runtime support for higher level concurrent programming language, SystemJ. In this paper, we describe the details of system design approach and hardware implementation of the Mobile SPOT with its sensing and actuating capabilities, as well as very brief elements of software support. Simple application scenarios are provided to demonstrate how Mobile SPOT navigates in the physical environment and communicates with other fixed and mobile nodes in a WSAN. As such, the Mobile SPOT can be considered as the basic infrastructure for the IoT-based applications.

The rest of the paper is organised as follows. Section II details the system design of the Mobile SPOT and its capabilities. Section III explains the hardware implementation of the Mobile SPOT. Section IV concludes the paper and provides planned future works based on the current progress.

II. System Design Approach

In this section, the design approach of how to integrate various sensing and actuating components into a Mobile SPOT is presented. The SunSPOT platform has a layered architecture and consists of three physical layers, namely the battery, the eSPOT board and the eDEMO board, stacked on top of each other. The eSPOT board is the main board of SunSPOT, which contains the ARM-9 microprocessor and the 2.4GHz radio transceiver. The eDEMO board consists of an auxiliary AVR-based microprocessor, which controls the on-board sensors, and generic digital and analogue IOs. The main ARM-9 microprocessor communicates with the auxiliary AVR through the high speed SPI bus.

The overall system architecture of the Mobile SPOT is shown in Fig. 2. The target applications of Mobile SPOT is to act as a mobile node which navigates within the physical environment, senses ambient conditions and communicates with other fixed/mobile nodes within the same WSAN. The built-in accelerometer and light sensors are used to measure linear acceleration of the Mobile SPOT and the light intensity in its surroundings. In order to achieve autonomous navigation, external sensors and actuators must be added to the original SunSPOT. In this project, a number of ultrasonic and infrared rangefinders are added to detect any obstacles in a Mobile SPOT's immediate environment. The sensing inputs are used to enable autonomous navigation by using a typical wall-following mechanism.



Figure 2. Mobile SPOT system diagram

Referring to Fig. 2, the ultrasonic rangefinders are digital sensors, which are linked to SunSPOT through the I^2C bus, whereas the infrared sensor is attached to one of the analog inputs on the eDEMO board. In this project, Devantech SRF08 ultrasonic sensor is chosen for its robust and accurate measurement, along with its outstanding measurement range from 6 cm up to 11 meters. Multiple ultrasonic rangefinders can be integrated to a single SunSPOT node via the serial I^2C bus. In addition to the ultrasonic sensors, a Mindstorms compatible sensor, SHARPGP2Y infrared rangefinder, is also adopted. This sensor has a shorter detection range of between 20 cm to 1.5 meters. Similar to the infrared rangefinder, a passive infrared (PIR) sensor is integrated through another analogue input of SunSPOT. This sensor is used for detecting human presence in front of the Mobile SPOT, instead of detecting rigid obstacles for Mobile SPOT navigation. Fig. 3 illustrates the detection range and the Field of View (FOV) for various rangefinders and PIR sensor. It is worth mentioning that the ultrasonic rangefinder has a cone shape FOV, whereas the infrared range finder detects an obstacle that is within its straight line of sight. Unlike the rangefinders, which detect rigid, fixed obstacles, the PIR sensor detects the presence of a moving human within the range of 10 meters and 160° FOV.



Figure 3. The sensors FOV

Based on the different sensing capabilities, ultrasonic and infrared rangefinders measurement results are used in different ways. The front ultrasonic sensor is used to detect any obstacles in front of the Mobile SPOT and also to keep it "on track" while performing "wall following" navigation. The right ultrasonic and the infrared rangefinders are used mainly to maintain the distance between the Mobile SPOT and the wall within the desired range while it is navigating autonomously. Based on the sensor inputs, a typical wall following algorithm is implemented in Java to compute the necessary motor control outputs (referring to section III). Four digital outputs (H0-H3 in Fig. 2 and Fig. 5) in the SunSPOT are used to generate PWM control signals to an external H-bridge, which in turn drives two LEGO Mindstorms motors. All the onboard sensors, external sensors, and actuators can be accessed and controlled through the standard APIs provided by the Squawk VM.

In order to communicate with other nodes within the same system, wireless communication is achieved through the onboard 2.4GHz CC2420 radio transceiver. An IP-enabled, 802.15.4-based, 6LoWPAN protocol stack is also provided in the standard APIs of the Squawk VM. TCP or UDP sockets can be easily established between two SPOT nodes for communication purposes. The ability of the Mobile SPOT to establish wireless communications provides many desired features. The Mobile SPOT is able to collect information and/or commands from a remote node or to exchange information with other Mobile SPOTs. Collected sensor data can also be uploaded to a remote PC equipped with a SunSPOT base station which allows remote monitoring of the behaviour and the immediate environment of the Mobile SPOT. Moreover, this allows teleoperation of the Mobile SPOT by issuing commands via a remote PC.

III. IMPLEMENTATION

This section describes the actual hardware and software implementation of the Mobile SPOT, targeting autonomous navigation based on the wall-following mechanism, remote data upload and teleoperation of the Mobile SPOT. The constructed Mobile SPOT is shown in Figure 4. The original LEGO NXT brick (controller) is removed and replaced with a SunSPOT wireless node, which is put on top of an electronic circuit board attached to the remaining LEGO Mindstorms mechanical build.



Figure 4. A snapshot of the constructed Mobile SPOT

The overall hardware implementation is illustrated in the schematic diagram shown in Fig. 5. Both ultrasonic sensors are connected to the SunSPOT through I^2C. The infrared rangefinder and the PIR sensor are interfaced to the analogue inputs (A0 and A1, respectively, as can be found in [8]).

The sensors are powered through a L7805 5V voltage regulator and also from the SunSPOT power output (a 3.3V power output pin to power the PIR sensor). The two SRF08 ultrasonic sensors are installed and facing the front and the right side of the Mobile SPOT, respectively, as mentioned in the previous section. The SHARPGP2Y infrared rangefinder is located at the right front corner, pointing at 45-degree angle. Figure 6 depicts the installed ultrasonic and infrared sensors on the Mobile SPOT.



Figure 5. Complete Hardware Schematics



Figure 6. The rangefinders positions on the Mobile SPOT

The PIR sensor itself is installed on an elevated position, facing the front direction of the Mobile SPOT, in order to detect the presence of a moving human, as shown in Fig. 7.



Figure 7. PIR sensor on an elevated position

The Mobile SPOT navigates in the physical environment using a "right wall following" mechanism. This scenario is based on a simple concept of maintaining a robot at a given distance from the "wall" on its right side while moving. If the robot drifts away from the specified distance (either greater or less than the desired distance), the robot has to alter its movement to bring itself back to the "track" to be within the desired distance to the wall. The Mobile SPOT "wall following" mechanism implements a PD (Proportional and Derivative) control algorithm. The algorithm implemented in Java will produce a pair of PWM (Pulse Width Modulation) signals with variable duty cycles to the SunSPOT digital output pins (H0, H1 and H2, H3), which drive an external H-bridge IC (SN754410 as in Fig. 5). Each pair of PWM signals act as differential inputs to the H-bridge to generate a single PWM control signal that controls a DC motor. Four digital outputs are required to produce two pairs of differential inputs, which drive the left and right motor of the Mobile SPOT. Two motors are sufficient for providing necessary movements through a physical environment on a plane surface.

To demonstrate the implemented functionalities, where the Mobile SPOT can detect nearby objects in its vicinity and perform right wall-following strategy, a simple testbed as depicted in Fig. 8 is built. Sensor fusion process is performed on the measurement results of all rangefinders. With the designed sensors configuration, the robot is able to navigate autonomously.



Figure 8. The robot while navigating through a shaped maze

A PD control algorithm is implemented on the Mobile SPOT to facilitate smooth robot movement. Equation (1) shows the general formula of a Proportional and Derivative algorithm.

$$y(t) = Kp \cdot e(t) + Kd \cdot \frac{d}{dt} e(t) \qquad (1)$$

In Equation (1), *y(t)* represents the output of the system, *Kp* corresponds to the proportional constant, *Kd* denotes the derivative constant, while *e(t)* defines the difference between the actual measured value (process value) and the desired value (setpoint). The algorithm takes the difference between the desired "range to the right wall" and the current measurement from the rangefinders and outputs PWM signals with variable duty cycles through SunSPOT digital output interface (H0-H3) [8] which will drive the SN754410 H-bridge circuit.

Through this mechanism, the algorithm will keep the Mobile SPOT moving along the right wall while maintaining the specified distance to the wall, and also to avoid any obstacles that might interfere with its movement. Figure 9 shows the flow diagram of the implemented PD control.



Figure 9. PD Control diagram

RF(t) denotes the current measurement results from all rangefinders, RF(t-1) denotes the previous results from all rangefinders, RFdesired is the desired distance between the Mobile SPOT and obstacles/wall, E(t) represents the difference between the measured range and the desired range, E(t-1) represents the previous difference between the measured range and the desired one, while PWMLeft and PWMRight correspond to the PWM signals for the left motor and the right motor, respectively. The PD control block takes inputs E(t) and E(t-1) and processes them to produce variable PWM signals for both motors. Figure 10 shows the refinement of the PD control block in Figure 9.



Figure 10. Refinement of the PD Control block

Eusf(t) and Eusf(t-1) correspond to the current and the previous (delayed) range values returned by the front ultrasonic sensor, while Eusr(t) and Eusr(t-1) represent the current and the previous range measurements from the right ultrasonic sensor, and EIR (t) and EIR(t-1) denote the current and the previous ranging results obtained from the infrared rangefinder, respectively. Each of the current range measurements is multiplied by specified proportional constants (Kp1, Kp2, and Kp3 respectively), while each of the previous sensor results is multiplied by particular derivative constants (represented as Kd_1, Kd_2, and Kd_3 respectively).

```
if(frontsonar>=rangelimit && IRranger>=rangelimit
&& rightsonar<rangelimit)
{
  leftPWM = fullspeed -
            (Kp*(rangelimit-rightsonar)-
             Kd*((rangelimit-rightsonar)-
            (prevrightsonar-rangelimit));

  rightPWM = fullspeed;

if (leftPWM<speedthreshold)
{
  leftPWM=speedthreshold;
}
}
… //other possible sensors outcomes with similar
calculations and thresholdings

robot.ActuateMotor(leftPWM,rightPWM); //invoke a
Java method to actuate both motors
```

Listing 1. PD control with decision and thresholding algorithm

The multiplication results from the same sensor are summed together, thus producing three PWM duty cycle representations. Next, these PWM results are inserted to the Decision and Thresholding block that saves several usable formulas to calculate the resulting PWM outputs. This block takes into account all possible sensor ranging outcomes (which correspond to the Mobile SPOT's immediate environment) and decides which formula to be used to generate the resulting PWM duty cycles for both motors based on the current rangefinders outcomes. Listing 1 provides a snapshot of a Java code that implements the PD algorithm on the Mobile SPOT.

In addition to navigating freely in the environment, the Mobile SPOT also collects information of its local environment, such as ambient light intensity and human presence, and upload the collected data to a remote PC. During its operation, the Mobile SPOT is constantly transmitting information (namely presence detection from the PIR sensor, accelerometer readings, light intensity, and reading from the built-in temperature sensor in the SunSPOT) wirelessly to a remote PC host via a fixed SunSPOT base station using a built-in datagram-based communication protocol. The SunSPOT base station is connected to a permanent power supply. Therefore, base station is usually used as a gateway interface for linking ordinary PC with 6LoWPAN based wireless network. Figure 11 illustrates a snapshot of how the received information is presented to the user through the PC Graphical User Interface.



Figure 11. A Graphical Interface that provides information to the user

```
//acquire access to accelerometer device
(happens in //other device)
IAccelerometer3D accel = (IAccelerometer3D)
Resources.lookup(IAccelerometer3D.class);

//open UDP connection to the target IP and port
number
url = "udp://" + IPaddress + ":" + port;
UDPConnection sender = (UDPConnection)
Connector.open(url);

//get tilting angle in x axis
tilt_x = accel.getTiltX();
//get tilting angle in y axis
tilt_y = accel.getTiltY();
//get tilting angle in z axis
tilt_z = accel.getTiltZ();
//wrap data into single instance for sending
tiltres = wrapReading (tilt_x,tilt_y,tilt_x);

//send reading through UDP protocol
Datagram dg =
sender.newDatagram(client.getMaximumLength());
dg.writeUTF(tiltres);
sender.send(dg); //send the datagram to the
destination
```

Listing 2. Java code to acquire accelerometer reading

Besides autonomous navigation and real-time remote monitoring, a user is also able to manipulate the motion of the Mobile SPOT by using another SunSPOT node as a remote controller. As each SunSPOT node has an onboard accelerometer, the tilting or panning information can be collected through accelerometer readings and transmitted to the remote Mobile SPOT. Based on the received tilting information, the Mobile SPOT can adjust its direction and speed of motion. Listing 2 provides a part of Java code on how the magnitude of tilting and panning is acquired by making use of the onboard accelerometer and the reading is sent to the remote Mobile SPOT through a UDP socket.

Once the reading has been acquired, the remote controller SPOT transmits the reading to the Mobile SPOT, in which the Mobile SPOT will adjust its own movement (direction and speed) based on the reading it received. Listing 3 shows a part of the Java code to receive data and process the data to generate motion commands on the Mobile SPOT.

```
UDPConnection receiver = (UDPConnection)
Connector.open("udp://:"+portnumber);

Datagram dg = (UDPDatagram)
receiver.newDatagram(receiver.getMaximumLength()
);
//receive reading through UDP protocol
Datagram dg =
receiver.newDatagram(client.getMaximumLength());

While (true)
{
  //receive tilting magnitude from the sender
  receiver.receive(dg);
  tiltres = receiver.readUTF();

  //extract reading for each axis
  tilt_x = tiltres.getElementAt(0);
  tilt_y = tiltres.getElementAt(1);
  tilt_z  = tiltres.getElementAt(2);

  //calculate PWM duty cycle for both motors
  leftPWM = robot.CalculateLeftPWM
            (tilt_x,tilt_y,tilt_z);
  rightPWM = robot.CalculateRightPWM
            (tilt_x,tilt_y,tilt_z);

  robot.ActuateMotor(leftPWM,rightPWM);
}
```

Listing 3. Java code on the robot side to receive data from other SunSPOT

The datagram received contains the reading from the remote controller SPOT. The reading can be easily extracted (through readUTF() method) , and then unwrapped through three consecutive methods (getElementAt()). The values returned by these methods will be subjected to PWM calculation methods for both motors. The resulting PWM calculations for both motors are used by the ActuateMotor() method, which will drive the Mobile SPOT with a particular direction and speed based on the received tilting magnitude.

IV. Conclusions and Future Works

Mobile SPOT, an autonomous wireless sensor and actuator node for the future Internet of Things applications, is described. The node builds on integrating technologies, of an advanced wireless node, SunSPOT, and mechanical parts of the LEGO Mindstorms robot, and extends the capability and applicability of original LEGO Mindstorms robot in terms of sensing, actuating, communication and software development tools. The Mobile SPOT has enough processing power to enable smooth movement based on its own sensing and motor control capabilities, by avoiding simple obstacles and sensing its immediate environment and communicating wirelessly via the Internet enabled 6LoWPAN protocol. Therefore, it can be used in creating versatile wireless networks and application systems in the context of the Internet of Things and collaborate with the same type of mobile nodes or fixed WSAN nodes placed in the environment.

Based on the current progress, the next step is to incorporate multiple Mobile SPOTs, targeting more complex operations in a collaborative manner. Such distributed and concurrent applications pose difficulties for system designers to ensure system functionalities at design time. Therefore, a high level concurrent programming language, SystemJ, designed specifically to handle highly concurrent and distributed applications, will be used for this type of scenario. Together with multiple fixed and Mobile SPOTs, the ability of using SystemJ to decompose complex system behaviours and implement system functionalities will be demonstrated.

References

[1] Atzori, L., Iera, A., and Morabito, G., "The Internet of Things: A survey," Computer Networks, vol. 54, pp. 2787-2805, 2010.

[2] Hong, J. and Hu, B., "IEEE 802.15.4 compliant RF module design," in Microwave, Antenna, Propagation and EMC Technologies for Wireless Communications, 2005. MAPE 2005. IEEE International Symposium on, 2005, pp. 910-913 Vol. 2.

[3] LEJOS. (16 October). NXJ Technology. Available: http://lejos.sourceforge.net/nxj.php

[4] Malik, A., Salcic, Z., Roop, P. S., and Girault, A., "SystemJ: A GALS language for system level design," Computer Languages, Systems, & Structures, vol. 36, pp. 317-344, 2010.

[5] Mulligan, G., "The 6LoWPAN architecture," presented at the Proceedings of the 4th workshop on Embedded networked sensors, Cork, Ireland, 2007.

[6] Simon, D., Cifuentes, C., Cleal, D., Daniels, J., and White, D., "Java on the bare metal of wireless sensor devices: the squawk Java virtual machine," presented at the Proceedings of the 2nd international conference on Virtual execution environments, Ottawa, Ontario, Canada, 2006.

[7] Smith, R. B., "SPOTWorld and the Sun SPOT," presented at the Proceedings of the 6th international conference on Information processing in sensor networks, Cambridge, Massachusetts, USA, 2007.

[8] Sun-Labs, "Sun™ SPOT eDEMO Technical Datasheet Rev 8.0," Sun Oracle America Inc, Menlo Park, California.2010.

Vision-based Positioning of Robots

Alain A. Evina and Charles Tsui

School of Electrical Engineering and Trades

Centre of Engineering Excellence

Manukau Institute of Technology

Auckland

andj1@manukaumail.com, charles.tsui@manukau.ac.nz

Abstract: This paper discusses applicability of computer vision and wireless communication in robot control. The experimental set up consists of two robots working together to find a ball in a playing field and then to drop the ball at a specified position. The objective was to eliminate the use of remote controller which emphasises a direct human interaction with the control of the robots. A camera was used to acquire a snapshot of the playing field. During the operation, the snapshots were segmented and used to navigate the robots to the position of the ball. Two improvements to the approach have been indentified and possible solutions were recommended.

Keywords: Image Segmentation, Embedded Control, Robot positioning, Junior Soccer Robot, Wireless Communication.

1 INTRODUCTION

Since the first public announcement of the Robot World Cup Initiative (RoboCup) in 1993 [1], it has been widely used by researcher to promote artificial intelligence (AI) robotics development. The RoboCup Soccer games have been developed into five categories, Simulation, Small-Size, Middle-Size, Standard Platform and Humanoid leagues. All competitors are using vision based techniques, one way or the other to facilitate AI control on the robots to achieve the goal.

RoboCup Junior [2] is an offspring of RoboCup that targets younger students (technically up through age 19) as extra-curricular activities. Its Soccer Challenge game allows autonomously running robots to detect an infrared emitting ball. The playing field has special pattern to help robots finding their orientation. In [3] an attempt of introducing the camera was made. However the field pattern was still the key factor for robot orientation. In [4], it is clearly explained how to give robot vision but expensive hardware is needed. A special miniaturised camera was attached to the robot. The price of such camera was around 350USD, too expensive for the project explained in this paper.

This paper explains an approach that prototype the Small-Size League of the RoboCup Soccer game [5]. A low cost camera was used to provide a global view of the field to help identifying the robots' positions in the field. The overall objective was to eliminate the use of remote control which emphasizes a direct human interaction with the control of the robot. An image segmentation technique was designed to distinguish the robot, the ball and the target point. This technique was combined with a robot navigation technique to successfully navigate two robots in a playing field. The approach does not require the playing field with a particular pattern for the robots to follow. However, the field and destination location colours must be differentiable and each robot must have unique colour for identification.

2 HARDWARE

The experimental set up (Fig. 1) consists of two identical robots working together to find a ball in the playing field, one webcam, an overhead light, a computer, and three wireless modules. The closest robot to the ball finds the ball and the second robot starts following the first when the ball is at the target point on the playing field. The camera serves as an eye. It maintains the feedback loop that update the position of the robots to the computer in real time. The computer serves as a brain. The wireless modules receive commands from the computer and feed them to the robots by the communication port. Once the command is received, an internal program interprets the command and the robots take actions appropriately. At the completion of the execution, a message is send back to the computer.



Figure 1. Photo of the system set up.

1.1 The Robots

The robots used in the project were two identical Super RCU-based (Robot Control Unit, Fig. 2) robots developed by JoinMax [6]. Note that the robot itself is a combination of a RCU, tires, light-sensors, roof (two-colour) identifier and a wireless module attached to it. The RCU is based on an ATMega16L programmable microcontroller; the program is written on the computer and downloaded to the RCU via the communication port (COM port).

The robots come with programming software RoboExp which combines two programming languages, G language (a graphical language) and C. G programming language is a set of icon that the user needs to combine in an appropriate sequence. When each icon is fully configured and the sequence is well ordered the program is automatically generated. There is low or no need of having text-based programming skill, but do need knowledge on implementing sequential and combinational logic with timing control of events. The RoboExp development software allows exporting to text-based C language. A skilful C programmer can build a very complex program that can make the super RCU accomplish complicated tasks.

The memory of the super RCU has 16k storage capability. It comprises 24 ports; 8 analogue and 12 digital ports. It has also 2 special motor ports, one serial communication port to download the program. A power socket for use with external power supply and one power indicator light are provided as well. The small LCD display can display either numbers or characters and can be used to display messages received or send by the RCU.



(External)



(Internal, showing the ATMega16L microcontroller)

Figure 2. Super RCU.

1.2 The Camera and Overhead Lighting

The Logitech c920 [7] (Fig. 3) with onboard H.264 compression and full 1080p high-definition video recording was used in the project. The build-in compression that compresses video and photos to a desired size allowed for faster image processing. The camera was used to take snapshot of the field every 9 seconds and the images were stored in a file inside the computer for processing.

The camera was fixed on a stand above the operation area and a set of LEDs was placed on each side of the camera to help eliminate shadowing and hence aid the colour segmentation which was crucial for the robot vision. The LED lights reduced the variations in the colour appearance of roof-top on the robots, ball and the playing field (Fig. 1).

Figure 3. C920 Logitech Camera.

1.3 The Playing Field

The field was made from wood and cardboards. The wood frame was 1.2 m long by 1 m wide. The floor of the field was covered with a white paper border by large black area along the edge [8]. The black border on the field was used to identify the boundary of the field. The destination target area is coloured with light purple.

1.4 The Wireless Modules

To avoid the use of cords three Xbee modules [9] were used to facilitate wireless communication between the robots and the command computer. Using ZigBee standard, the Xbee modules are configured as one coordinator connected to the computer and two end-nodes, each end-node connected to one of the robots. The coordinator receives commands from the computer and feed into the serial communication port of each robot. The Xbee modules have a communication range of 100m at a supply voltage of 3.4V was more than sufficient in our case.

3 SOFTWARE

Three separate programs were used to run the project: image acquisition, image segmentation and robot guiding and Robot command interpretation program (written in C and run by the robots). First two programs were written in Java and are run concurrently on the computer.

At the start, the Image acquisition program initiates the camera to continually take pictures of the field and its contents. During playing, the image segmentation and robot guiding program grabs the saved picture, breaks the picture into pixels of RGB components and performs image segmentation. This helps to recognise everything of interest that is in the playing field, i.e. robots, ball and destination.

Once the picture is taken and processed, the most advanced robot, the nearest to the ball, is chosen and the program selects the address of that robot to send a command. The chosen robot picks and drops the ball to the destination point. The second robot follows the first after the task is completed (Fig. 4).



Figure 4. Flowchart of image segmentation program.

1.5 Guiding the robot

The computer is also responsible for guiding the robots (Fig. 5). The action command is determined by an algorithm based on the relative position of the target point compared to the current position of the robots. Each command is send as an ASCII code. To turn left the computer send L to the robot and the internal program inside the robot interpret letter L as a left turn. Same procedure occurs when it is an R (right) turn. The decision on the guiding computer is based on the centroid of each object of interest on the photo after colour segmentation. If the position Y of the robot is greater than the position Y of the ball and the position X of the robot less than the position X of the ball the computer send R and it is interpreted by the robot as right. This calculation is performed as long as the robots are running.

In the robot, a block function was written for each of the listed commands. Once the command is received that block executes for a specific amount of time. After that, the robot goes to idle position waiting for the next com-

mand. Table 1 shows the list of commands received by the robots.

Table 1: List of robot movement commands.

F (Forward)
L (Left)
R (Right)
P (Pick)
D (Drop)
T (Turn)
S (Stop)

4 RESULTS AND DISCUSSION

The robots were successfully tested. They were able to pick the ball and go to the destination following commands received. Picking up the ball and moving it around was sometimes very difficult. The ball can easily stick onto the field and the robots tend to climb on top of the ball. This may be fixed by introducing an appropriate part that could either kick or pick up the ball.



Figure 5. Flowchart of guiding robot algorithm.

This project demonstrated the design with a working prototype. It has successfully direct the Super RCU-based robots to move the ball to a defined target area without the use of additional sensors (such as infrared) to find the ball. Colour segmentation with appropriate algorithm on the guiding computer was able to find the target, calculate the relative locations and send the necessary command to the robots. However, robots are a bit primitive. They take quite an amount of time to accomplish the tasks. When motors were programmed at higher speed the robot overshoot. When they run at lower speed, much longer time was taken to process and execute each command making the overall respond very slow. Using advanced microprocessors would improve the speed.

It was also found that robot identification was influenced by variations of the light intensity of the environment. External sun light tends to affect colour segmentation. A set of LED lights was used to counter add on this and provided improvements. Therefore, if the project runs in a controlled environment where there is a steady light intensity, this problem can be avoided.

The main difference between RoboCup Small-Size League [5] and this project is that the RoboCup used a field having no special feature, but just a single green coloured mat or carpet. For this project, the target area was painted with a special colour to distinguish it from the floor. The similarity is both having a special colour pattern on top of each robot for indentifying them.

5 CONCLUSIONS AND FUTURE WORK

This paper described the successful prototyping and testing of a low cost vision-based robot positioning system. However, the overall processing was somehow slow and influenced by the quality of the lighting. A faster processing robot together with a more uniform lighting would improve the image processing task, and therefore enhance the overall performance of the system. Due to such slow response time of the system the robot soccer game algorithm was not fully implemented.

Moreover, further work is necessary to investigate the benefits of performing colour segmentation using more advanced software such as MATLAB to take advantage of its image processing tool. There is room to improve on the robot's dynamic mechanical behaviour and the algorithm to execute movement commands.

6 REFERENCES

[1] The Robocup Federation . (2012). *A Brief History of RoboCup* [Online]. Accessed on 12th October, 2012: http://www.robocup.org/about-robocup/a-brief-history-of-robocup/

[2] RoboCup Junior. (2012). *Soccer Challenge* [Online]. Access on 10th November 2012: http://rcj.robocup.org/soccer.html

[3] A. Junghan. (2001). *Collaborative Robotics with Lego Mindstorms,* Master Thesis, Karlsruhe University of Applied Science, Germany [Online]. Accessed on 10th November 2012: http://junghans-schneider.de/download/masterthesis-junghans.pdf

[4] S. Meloan (July 2003). *Futurama: Using Java Technology to Build Robots That Can See, Hear Speak, and Move* [Online]. Accessed on 10th November 2012: http://www.oracle.com/technetwork/articles/javase/robotics-138733.html

[5] The Robocup Federation . (2012). *Laws of the RoboCup Small Size League 2012* [Online]. Accessed on 10th November 2012: http://small-size.informatik.uni-bremen.de/_media/rules:ssl-rules-2012.pdf

[6] JoinMax Digital (2012). *Super RCU* [Online]. Accessed on 10th November 2012: http://www.roboexp.com/products/list.asp?id=8

[7] Logitech (2012). *Logitech HD Pro Webcam C920* [Online]. Accessed on 10th November 2012: http://www.logitech.com/en-nz/webcam-communications/webcams/hd-pro-webcam-c920

[8] H. H. Lund and L.Pagliarini. (2000). *Robocup Jr. with LEGO Mindstorms*, In Proc. of IEEE International Conference on Robotics and Automation, 2000, vol.1, pp. 813-819.

[9] Digi International Inc. (2012). *RF modules utilizing the ZigBee PRO Feature Set* [Online]. Accessed on 10th November 2012: http://www.digi.com/products/wireless-wired-embedded-solutions/zigbee-rf-modules/zigbee-mesh-module/xbee-zb-module#overview

Efficient Feature Based Visual Servoing for Onboard Computer Vision

S. Duncan, M. Hayes, A. Bainbridge-Smith

*Department of Electrical and Computer Engineering,
University of Canterbury,
Christchurch, New Zealand

Abstract: A robust and efficient algorithm for pose estimation is presented based on restricting the permissible matches between reference and query sets of features. Possible matches are restricted based on proximity to reference features projected using an *a priori* pose estimate with an estimate of its accuracy. Update rate is improved over RANSAC pose estimation, at best by a factor of 3. Robustness to camera deviations, and to reference and query images exhibiting repetitive patterns or many similar features (e.g. brick walls, text). Tests run on a Gumstix single board computer using 640×480 images, with varying poses. While performance is improved, realtime performance is still unfeasible.

Keywords: Visual Servoing, Feature Detection, Embedded Computer Vision, Unmanned Vehicles

1 Introduction

Traditionally, the computationally expensive nature of computer vision (CV) algorithms and the large data volumes associated with visual data streams has precluded real time performance in embedded computing environments. In this paper an efficient and robust feature based visual servoing algorithm using FAST [1] feature detection and BRIEF [2] feature descriptors is presented. The ePnP [3] algorithm is used for efficient pose estimation and performance is improved by restricting permissible feature matches based on the pose estimate and an estimate of its accuracy. Similar work has been done by Voigt et al in [4] and by Dorini and Goldstein [5] with the aim of improving robustness. The material in [4] is similar but uses binocular vision and relies in part on inertial cues. In [5] the focus is on robust feature tracking in the general case. The work presented in this paper is aimed at achieving real time performance in computationally constrained environments.

The theory of estimating the pose accuracy is presented in section 2, with a brief overview of reduced complexity of feature matching. In section 3 a constrained matching and pose estimation algorithm is developed, based on the results from section 2. The algorithm is then tested on a Gumstix[1] single board computer (SBC) as described in section 4. The results of testing are presented in section 5. This paper concludes with a discussion of the results in the context of using a low powered SBC for the control of a small quadrotor.

2 Visual Servoing

The basis for feature based visual servoing is pose estimation. Pose estimation uses a set of 2D-3D point correspondences in image space and world space to estimate the pose of the camera. In visual servoing the pose is compared to some set point, and control signals are generated to minimize the error between the two, so as to position the camera relative to some real object or target [6]. Point correspondences are generated by matching a set of reference features with known world locations (the training set) to a set of features detected in a frame from the camera (the query set). The matching process is computationally expensive, with a complexity of

$$O(N_T N_R H), \tag{1}$$

where H is the complexity of matching one feature [7], and N_T and N_Q are the number of training and query features. However, Hajjdiab and Laganiere [7] have shown

[1] http://gumstix.com



Figure 1: An illustration of the constrained matching process

that restricting the permissible matches based on location can reduce the complexity to

$$O\left(\frac{N_T N_R H}{\rho}\right), \tag{2}$$

where ρ is a factor defined by the spatial density of the query features and the size of the permissible match region.

The size of the region can then be used to trade performance for robustness. The rest of this section provides an outline of how an approximate measurement of the accuracy of the pose estimate is used to determine the size of the permissible match region.

2.1 *Pose Accuracy*

Due to quantization noise and other disturbances, the location of each training and query feature cannot be measured precisely. As a result, the coordinates of each feature can be represented by a vector of Gaussian random variables [4] with a mean position $\bar{\mathbf{p}}$ and covariance matrix $\mathbf{R}$. The variance of the feature locations results in a similar uncertainty in the pose estimate. Typically it is impractical to measure the covariance of the points or the pose directly, and instead the reprojection error is used as a more readily accessible correlate,

$$\sigma^2 = \frac{1}{N}\sum_{i=1}^{N} ||\mathcal{P}(\mathbf{p}_i) - \hat{\mathbf{p}}_i||^2, \tag{3}$$

where $\mathcal{P}(\mathbf{p}_i)$ is the projection of the world point matched to the detected point $\hat{\mathbf{p}}_i$, and N is the number of point correspondences.

σ^2 is then used to determine the radius of the permissible match region,

$$r = \alpha + \sigma^2\delta, \tag{4}$$

where α is the base radius of the region, δ is a scaling factor, and r is the radius of a circular region of permissible matches as shown in Figure 1. Circular regions are used for simplicity in this paper.

3 Constrained Feature Matching

Based on equations (3) and (4) the constrained matching process is implemented as illustrated by Figure 1. The algorithm uses a RANSAC based pose estimation algorithm to find the initial estimate of the pose, and its accuracy. Once initialized, the following steps are performed for each camera frame:

1. Project training features into frame n using the *a priori* pose estimate from frame $(n-1)$.
2. Use (4) to calculate permissible region size.
3. Construct a match mask for OpenCVs match function.
4. Match features using the mask.
5. Update the pose, and the reprojection error using (3).
6. Threshold σ^2 to detect 'marker lost' state.

   $lost \rightarrow$ re-initialize.

   $\overline{lost} \rightarrow$ run from step 1.

The above process is valid if motion between frames is slow relative to the frame rate, a valid assumption for visual servoing.



(a) (b)

(c) (d)

Figure 2: Point correspondences found by matching training features to query features in a live camera feed and in the wall dataset[2]. (a), (c) unconstrained matching, which requires a RANSAC outlier rejection scheme to produce good pose estimates. (b), (d) constrained matching, where outliers are eliminated by restricting the feature search and match regions.

3.1 Initialization

The initial pose estimate is found using a RANSAC based scheme to find an accurate pose estimate in the presence of false matches. Once a sufficiently small σ^2 is found, the corresponding pose estimate is used as the *a priori* estimate in step 1 above.

3.2 Robustness

Robustness refers to the mechanism for recovering the target if it is temporarily occluded or otherwise 'lost' resulting in inaccurate pose estimates. For this paper it is assumed that loosing the marker is a result of occlusion or distortion (e.g. motion blur), and that the marker is always in the field of view (FOV) of the camera.

A threshold is applied to σ^2 to detect a lost marker situation. Once detected, the program returns to an initialization state until a sufficiently small σ^2 is measured. Acceptable σ^2 values are dependant on the accuracy of the training and query feature sets, and vary between images. In this paper $\sigma^2 = 2$ was used for the live camera data, and $\sigma^2 = 3.5$ for the wall dataset.

[2] `http://www.robots.ox.ac.uk/~vgg/data/data-aff.html`

4 Testing

Tests were run in both desktop and embedded environments and for both constrained matching and a RANSAC based alternative for comparison. The RANSAC approach is widely used as a robust way to estimate pose from a set of point correspondences containing false matches. An unconstrained algorithm without RANSAC was also tested as a reference for the update rate performance, which is dependent on the number of reference features. The number of query features is maintained at approximately 500 by adjusting the FAST threshold based on the number of detected features.

4.1 Performance Metrics

The performance metrics of interest are the maximum frame rate and the accuracy and variance of the pose estimate. Frame rate is measured by timing a large number (10^4) of frames and using the average processing time per frame to calculate the frame rate.

Accuracy is compared via the mean and variance of the reprojection errors of each method. Variations in the reprojection typically correlate with variations in the pose estimate. Because the tests are run with a stationary camera, the variance of the reprojection error as well as the mean is recorded.

4.2 Equipment

The embedded environment used to measure the performance of the approach is a Gumstix Overo Earth SBC. The ARM based SBC has a small form factor which lends its self to unmanned aerial vehicle (UAV) control. Processing power is limited, with the CPU clock limited to $600\,MHz$ and an advertised performance of 12000 Dhrystone MIPS[3].

A Point Grey Chameleon CCD camera is used for image acquisition. Prior to testing the camera was calibrated; the intrinsic camera matrix, and the lens distortion coefficients are known.

4.3 Procedure

Tests were run with the camera fixed at a known ground-truth pose. A single test was run for each algorithm and dataset, and the pose, reprojection error, and processing time per frame were all logged. The timer measures processing time with an accuracy to μs, and only records the time required to process a frame and update the pose. The data was then processed and is presented in Figure 3.

5 Results

The test procedure was run as described in section 4, and the results presented in Figure 2 and Figure 3. Figure 2 illustrates the operation of the constrained matching approach, showing the false matches when matching is not constrainted. Figure 3 shows the performance metrics calculated for both algorithms. The frame rate of an additional unconstrained algorithm with no RANSAC outlier rejection is shown as a reference frame rate in Figure 3c. This unconstrained algorithm produces erroneous pose estimates, and is only used as a performance reference for the frame rate metric.

5.1 Frame Rate

In Figure 2 the point correspondences are shown as coloured lines between features in the reference (left) and query (right) images. By inspection it can be seen that the correspondences in Figure 2b are all correct, whereas in Figure 2a a significant proportion of the matches are incorrect, or outliers. The correspondences in Figure 2a can be used to estimate the pose, subject to a computationally expensive RANSAC outlier rejection scheme with a non-deterministic run time. Under constrained matching there are no outliers (Figure 2b). Constrained matching therefore does not require the expensive outlier rejection stage, and therefore more consistent and higher frame rates are possible. Figure 3c shows that the constrained approach is slightly faster than RANSAC outlier rejection, and the frame rate is also less variable. The difference between the two methods, however, is small. Even in the best case, constrained matching still does not afford real time performance.

5.2 Pose Variability

Figure 3 shows that the constrained method allows for consistently faster frame rates than the RANSAC approach, with comparable error performance. In Figure 3d unconstrained matching with no outlier rejection is shown for reference.

5.3 Tracking

As well as being faster and less variable than the RANSAC approach, constrained matching is also more robust camera pose deviations. Figure 2 illustrates this robustness; The camera pose of the query image differs significantly[4]. from the camera pose used to capture the train image, however the constrained approach still produces only correct matches. In addition to correct matches, there is less variability difference as the pose deviates from the training pose, in constrained matching compared with RANSAC outlier rejection.
The constrained approach is also more robust (as in Figure 2) when the marker exhibits repetitive patterns. This is also illustrated in Figure 2, where the text pattern contains many similar features.

6 Discussion and Future Work

In order to achieve stable control of a quadrotor, minimum update rates of around 10 Hz have been suggested[5]. The results shown in this paper fall short of this requirement, remaining limited to 2 Hz.

This limited performance is due in part to the test platform. The Gumstix was chosen for its form factor which is ideally suited to processing onboard small aerial vehicles. While the next generation of Gumstix SBCs are much more capable, and include multiple core architectures, faster clock speeds, ARM extensions, and DSP or GPU coprocessors, the most significant limitation of ARM SBCs for computer vision is OpenCV hardware optimization. Because OpenCV was originally an Intel research project, hardware optimizations are mostly targeted at Intel architectures, and more recently support for CUDA. ARM cores support many data processing extensions, and

[3] https://www.gumstix.com/store/product_info.php?products_id=211

[4] This excludes deviations with rotation about the z axis. For rotation invariance in the image plane, an alternate feature descriptor such as ORB [8] can be used.

[5] www.pxihawk.eth.chz



Figure 3: Performance of each method measured over 10^4 frames. (a), (c) Average frame rate of each method for the book live camera target and wall dataset respectively. (b), (d) Mean squared reprojection error for the constrained matching algirthm and for the RANSAC approach for the book live camera target and the wall dataset respectively.

include coprocessors which have the potential to significantly accelerate embedded performance of OpenCV. As of this writing, however, ARM architecture optimization in OpenCV is limited. Intel-based SBCs remain the most capable platforms for onboard computer vision.

6.1 Limitations

Pose estimation, and calculating its accuracy, are non-linear; A large σ^2 indicates an uncertain pose, however small values of σ^2 are also possible with some uncertain poses. σ^2 is only a valid estimate of the pose accuracy for small pose errors. It was also observed that for erroneous estimates with very large z translations (several orders of magnitude larger than the correct pose) produced very small values for σ^2. This is because reprojected points converge to a single point as the z translation relative to the camera tends toward infinity.

The hysteresis between frames limits the maximum speed of the target relative to the camera proportionally with the frame rate. Because the *a priori* pose estimate comes from the previous frame, motion which is fast compared to the frame rate requires much larger permissible match regions. To remedy this, equations (3) and (4) can be adapted to account for the accuracy and the direction of motion between previous frames as well as pose accuracy.

An additional improvement which would reduce the cost of feature detection is to also reduce the feature search area. Currently the FAST algorithm searches the whole FOV, however, the target may only occupy a small portion of this. Constraining the search region as well as the permissible matches can further improve performance.

7 Conclusion

Constrained feature matching was found to improve efficiency and robustness over RANSAC based outlier rejection. Update rates are improved by a factor of 3 in some cases, however even the best results are still insufficient for realtime performance. In the best case, the frame rate is limited to a maximum of about $2\,Hz$, significantly less than the $10\,Hz$ requirement for realtime quadrotor control. Robustness to poses differing from the training pose, and to markers containing many similar features is improved over RANSAC pose estimation in some cases. The results suggest, however, that the robustness is limited to pose variability rather than the pose estimate its self.

References

[1] E. Rosten and T. Drummond, "Machine learning for high-speed corner detection," in *European Conference on Computer Vision*, vol. 1, pp. 430–443, May 2006.

[2] M. Calonder, V. Lepetit, M. Ozuysal, T. Trzcinski, C. Strecha, and P. Fua, "BRIEF: Computing a Local Binary Descriptor Very Fast," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 7, pp. 1281–1298, 2012.

[3] F. Moreno-Noguer, V. Lepetit, and P. Fau, "Accurate non-iterative O(n) solution to the PnP problem," in *IEEE International Conference on Computer Vision, Rio de Janeiro, Brazil*, 2009.

[4] R. Voigt, J. Nikolic, C. Hürzeler, S. Weiss, L. Kneip, and R. Siegwart, "Robust embedded egomotion estimation," in *Proc. of The IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, (San Francisco, USA), September 2011.

[5] L. Dorini and S. Goldenstein, "Unscented feature tracking," *Computer Vision and Image Understanding*, 2010.

[6] D. Kragic and H. I. Christensen, "Robust visual servoing," *The International Journal or Robotics Research*, pp. 923–939, 2003.

[7] H. Hajjdiab and R. Laganiere, "Complexity analysis of feature-based image matching," 2009.

[8] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to sift or surf," in *International Conference on Computer Vision*, (Barcelona), 11/2011 2011.

Bayesian Inference of Position from Short Samples of GPS Signals

Tom Holmes, T.C.A Molteno*

*Department of Physics, University of Otago, Dunedin, New Zealand

tim@physics.otago.ac.nz

Abstract: We describe a Bayesian approach to GPS positioning where post processing of ultra-short sequences of captured GPS signal data is used to infer receiver location. This paper extends a least-squares optimization scheme by analyzing the binary signal data using Bayesian inference with Markov chain Monte Carlo sampling. This allows not only postition and time estimates but also their relative uncertainty to be determined. The outlined approach reaches final solutions, giving position estimates typically within 250 meters of true receiver location.

Keywords: Global Positioning Systems (GPS), Animal Tagging, GSM telemetry.

A typical GPS device requires at least 30 seconds to calculate a position from cold start. In this paper we present a system that dramatically reduces the time-to-first-fix by simply storing a small (1ms) sample of GPS signal data. The short recording time (1ms) leads to a factor of 30,000 reduction in power required for a fix. Consequently much lighter batteries, and or solar power systems can be used, and the tag-weight is correspondingly lower.

We have previously used a least-squares algorithm to do this processing [1]. Bayesian inference with Markov chain Monte Carlo (MCMC) sampling can extract valuable information from limited data but often at the cost of high compute cost [2]. This paper outlines how one may apply such techniques to estimate receiver position and GPS time and also provide a quantitative measure of uncertainty in each estimated parameter.

1 Bayesian Inference

The GPS receiver stores N data samples from the GPS L1 band at 1.57542 GHz. This data is typically digitized at one-bit resolution and at a sampling frequency of 8 MHz. The forward map $F_k : \vec{x} \mapsto y_k$, maps the model parameters $\vec{x}$ (latitude, longitude, etc.) to the N data samples y_k, $k = 1, 2, ..., N$ [3]. As our data is subject to random variation (the amplitudes of the signals of interest are well below the thermal noise floor), this falls into the category of statistical inference [4]. We employ Bayes' Theorem,

$$f(\vec{x}|y_k) = \frac{1}{f(y_k)} f(y_k|\vec{x}) f(\vec{x}) \tag{1}$$

where $f(\vec{x})$ is the prior pdf (probability density function), $f(y_k|\vec{x})$ is the likelihood pdf and $f(\vec{x}|y_k)$ the posterior pdf. The $\frac{1}{f(y_k)}$ term can be thought of as a normalization factor and is often left out, giving instead [5]

$$f(\vec{x}|y_k) \propto f(y_k|\vec{x}) f(\vec{x}). \tag{2}$$

The prior probability $f(\vec{x})$ represents everything we know about our model parameters $\vec{x}$ before the data y_k has been collected. For example, if the GPS receiver was on a vehicle, we can be fairly certain that the receiver's altitude will be non-negative and also less than a few thousand meters depending on our prior knowledge of the receiver environment. The accuracy of a real time clock on board the receiver can also be expressed in the form of a probability density function for the time that the sample is taken.

Evaluation of the forward map $F_k(\vec{x})$ for a given state of the parameters $\vec{x}$ gives us the noise free data $\bar{y_k}$ we would expect to see at the receiver output, provided the forward map correctly modelled the system. Assuming a basic noise model of additive Gaussian noise with standard deviation σ, ie. $y_k = \bar{y_k} + \mathcal{N}(0, \sigma)$, we find the following expression for the likelihood [3],

$$f(y_k|\vec{x}) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2} \sum_{k=1}^{N} \left[y_k - F_k(\vec{x})\right]^2\right). \tag{3}$$

Multiplying the likelihood by the prior gives (up to a normalization constant) the posterior distribution $f(\vec{x}|y_k)$ – the probability distribution for the true values of $\vec{x}$. The statistics of the posterior distribution provide us with both estimates of our parameters $\vec{x}$, as well as quantitative information about our certainty of these estimates.

1.1 The Forward Map

Our forward map has four parameters, the receiver latitude, longitude altitude and time. Combining these with the satel-

lite (SV) ephemerides data, we can calculate each SV's expected C/A code phase, carrier Doppler shift and signal amplitude [6]. The first step is to convert the receiver's position from WGS84 coordinates (latitude, longitude and altitude on an ellipsoid) into the ECEF coordinate system (earth centered, earth fixed) as used in the ephemerides to describe the GPS satellite postions and velocities [7]. The L1 signal emitted from each satellite begins as a sine wave carrier multiplied by the satellite's C/A code, $C(t)$, giving

$$aC(t_t)\sin(\omega_c t_t) \tag{4}$$

where ω_c is the frequency of the L1 carrier signal and a is the signal amplitude. Note that the SV's data code has been left out as it only has possible bit transitions every 20 milliseconds. This introduces the possibility for some C/A codes to appear inverted and is dealt with in another part of the final code. The time of transmission, t_t, is equal to the time of reception, t, minus the time of flight, τ, giving

$$aC(t-\tau)\sin(\omega_c t - \omega_c \tau). \tag{5}$$

A signal's time of flight τ, can be broken down into,

$$\tau = (n+\delta_{C/A})T_{C/A} = (m+\delta_{L1})\frac{2\pi}{\omega_c} \tag{6}$$

where $T_{C/A}$ is the C/A code period (1ms), $n, m \in \mathbb{N}$ and $\delta_{C/A}, \delta_{L1} \in [0,1)$.

It is convenient to split the absolute time parameter t into two parts. The first, which we call the C/A code time offset, is the difference between the transmission time and the closest C/A code epoch beginning. The C/A code time offset will be expressed as $\delta_t T_{C/A}$ where $\delta_t \in [0,1)$, as it represents a fraction of one C/A code epoch. All SV clocks are synchronized (with atomic clock precision after application of clock errors from the ephemerides [8]) such that each will begin transmission of every C/A code epoch at the same instant, thus $\delta_t T_{C/A}$ is common to all SVs. The second part of t, which we will label t_z, is the reception time relative to the closest C/A code epoch beginning. Similarly, inside the sine function we split t into the same t_z and a carrier signal phase offset ϕ.

As the Doppler shifts for each SV signal are different and ϕ is relative to the Doppler-shifted carrier frequency, this introduces 32 new phase-offset parameters into the forward map (one carrier phase offset for each SV). The introduction of new parameters may seem counter-intuitive as it increases model complexity, but as will be shown later, it serves to 'smooth out' the likelihood space by removing carrier-phase perturbations. The absolute time still plays a role in the forward map, as it determines the positions of the satellites and thus their expected C/A code phases and Doppler shifts. Combining 5 and 6 we get

$$\begin{aligned} aC(t_z+&[\delta_t - n + \delta_{C/A}]T_{C/A}) \\ &\sin(\omega_c t_z + \phi - \omega_c(m+\delta_{L1})\frac{2\pi}{\omega_c}. \end{aligned} \tag{7}$$

Since both $C(t)$ and $sin(t)$ are periodic, with periods $T_{C/A}$ and $\frac{2\pi}{\omega_c}$ respectively, the integers n and m may be removed. Everything in the sine function with no dependence on t_z can be absorbed into the phase constant, ϕ, simplifying 7 down to

$$aC(t_z + [\delta_t - \delta_{C/A}]T_{C/A})\sin(\omega_c t_z + \phi). \tag{8}$$

For simplicity we ignore the effects of the Earth's ionosphere and troposphere on the propagating signal (including a time delay and frequency shift [9]) and assume a basic model of signal amplitude $A(\theta)$ (absorbing a from above) which depends only on the satellite's angle of elevation, θ. Adding in the signal's Doppler shift ω_{ds}, we get at the receiver's antenna,

$$A(\theta)C(t_z + [\delta_t - \delta_{C/A}]T_{C/A})\sin(t_z[\omega_c + \omega_{ds}] + \phi). \tag{9}$$

Multiplication of the signal with a local oscillator inside the receiver then downsamples the L1 carrier to the center frequency ω_{cf}. The signal is filtered to retain only the carrier's bandwidth (about 2Mhz [10]), then I and Q sampled at a user configurable sampling frequency ω_s, giving

$$A(\theta)C(t_{z_k} + [\delta_t - \delta_{C/A}]T_{C/A})\exp(jt_{z_k}[\omega_{cf} + \omega_{ds}] + \phi). \tag{10}$$

The sampling process involves digitizing to one bit precision, thus combining the signals from all satellites m, we get

$$\begin{aligned} H\Bigg(\sum_{m=1}^{32}\Big[&A(\theta^m)C(t_{z_k} + [\delta_t - \delta^m_{C/A}]T_{C/A}) \\ &\exp(jt_{z_k}[\omega_{cf} + \omega^m_{ds}] + \phi^m)\Big]\Bigg), \end{aligned} \tag{11}$$

where $H(x)$ is the Heaviside step function,

$$H(x) = \begin{cases} 0, x < 0 \\ 1, x \geq 0 \end{cases}$$

and the sampling frequency ω_s determines the spacing between samples, ie,

$$t_{z_k} - t_{z_{k-1}} = \frac{2\pi}{\omega_s}.$$

The forward map, as specified, requires more than just the three postion parameters and time. The additional, 'nuisance' parameters include the receiver center frequency ω_{cf}, sampling frequency ω_s, C/A code phase offset δ_t and the 32 signal carrier phases ϕ^m, giving a total of 39 dimensions in our state space.

1.2 Characterising the Posterior

We use Markov Chain Monte Carlo (MCMC) sampling [5] and the Metropolis Hastings algorithm to characterize the posterior. The random walk MCMC, as the name suggests,

takes steps of random size and direction in each dimension of the target distubution. Each proposed step is either accepted or rejected, ie. the chain changes or stays the same, with a probability of acceptance proportional to the extent that the step leads to a point of higher probability density, ie. the chain is more likely to move 'uphill'. In this way, our Markov chain spends more time in regions of higher probability density than it does lower. Plotting the chain's parameter states at each step in histograms gives an approximation of each parameter's marginal distribution. Figure 4 shows the result of running an MCMC on the posterior distribution of a typical data set.

The random walk MCMC method used is most simply applied in situations where the target distribution is uni-modal and decays smoothly away from the mode. This allows the chain to follow a positive gradient to find the mode. If, instead, the distribution decays rapidly from the mode, ie. a sharp peak, then has a near zero gradient for the remainder of the search space, the chain has difficulty being steered toward the mode and may never reach it. Searching the surface of the Earth for a peak with a width of around 30 meters posed a problem for sampling our posterior distribution, with the chain, more often than not, failing to finding the correct mode even after millions of steps. To tackle such a problem we use more informative priors, that reduce the search space or perform some reduction on the data[11]. By using less data we effectively widen the peak for which we are searching. In this paper, the inverse problem was broken down into three steps, each refining our knowledge of the parameters. The first is the Doppler-fix, followed by a more accurate phase-fix and finally a raw-data fix step.

2 Searching for a Solution

2.1 The Doppler Fix

The Doppler-fix step is a data reduction that provides a rough estimate of receiver position to seed the MCMC in subsequent steps. This is derived from the step of the same name in the Fastfix scheme [1]. The raw signal data obtained from the GPS receiver is cross-correlated with each of the 32 SV C/A codes, with each correlation repeated for a range of possible Doppler shifts, covering the receiver's center frequency $\pm$5kHz [6]. Figure 1 shows the Doppler shifts for the first eight satellites using data recorded at 7.22pm on 19th May 2010.

In essence we have added another step to our forward map and should rewrite it accordingly. However, correlating C/A codes for 32 SVs for 50 possible Doppler shifts each, takes significant CPU time ($\sim$ 20s), making thousands of forward map evaluations impractical. To simplify matters, we can create a new approximate (Doppler) forward map F_d which takes a parameter vector $\vec{x}_d$, consisting of the receiver's position, GPS time and center frequency. Rather than producing a binary signal vector as noise free data output, the Doppler forward map uses only the calculated Doppler shifts and



Figure 1: Doppler shifts for visible satellites. The dominant yellow peak indicates the signal from SV number 7 is the strongest visible and has been Doppler shifted by around -1000Hz.



Figure 2: Simulated Doppler shifts produced by evaluating the Doppler forward map at the correct receiver position.

satellite elevation angles (to determine strength) to simulate data in the form of Figure 1. We will refer to the reduced (real) data as $\vec{y}_d$.

The result of evaluating the Doppler forward map at the correct receiver location and time (corresponding to the data used to produce Figure 1) is shown in Figure 2. The error associated with the Doppler shift data is not simply additive Gaussian noise, however, treating it as such to obtain a Doppler likelihood function L_d, still holds merit. Our reduced data Doppler likelihood function is expressed as

$$L_d(\vec{y}_d|\vec{x}_d) \propto \exp\left(-\frac{[\vec{y}_d - F_d(\vec{x}_d)]^t[\vec{y}_d - F_d(\vec{x}_d)]}{2\sigma_d^2}\right). \quad (12)$$

Maximizing $L_d(\vec{y}_d|\vec{x}_d)$ is equivalent to minimizing the sum of squared differences between the real reduced data $\vec{y}_d$ and the simulated reduced data $F_d(\vec{x}_d)$ [3]. Figure 3 shows the structure of this Doppler likelihood in the vicinity of the receiver's true position, conditional on the parameters not plotted having been set to their estimated true values.



Figure 3: Doppler likelihood distribution of latitude and longitude parameters (conditional on set values of all other forward map parameters) with (roughly) Gaussian shaped peak centered around the true receiver location. Maximum value is 23km away from the true position.

The Doppler forward map is a function of five parameters (latitude, longitude, altitude, time and center frequency), we use random walk MCMC to characterize the PDF of these parameters. Figure 4 shows the result of running the MCMC on the Doppler likelihood. Our prior-knowledge of GPS system time is given by the real time clock onboard the receiver. Although this time often has an accuracy of only one or two seconds, it is more reliable than that provided by the Doppler-fix ($\sim$ 50s standard deviation), thus the histogram of chain states for the time parameter provides no new information and is omitted from the figure. Similarly, our prior knowledge of receiver altitude ($0 < altitude < 1000$m for example) is more accurate than that provided by a Doppler-fix ($\sim$ 50km standard deviation). This particular Doppler-fix gives a mean latitude and longitude of (-45.4948, 169.6732) which, ignoring altitude error, is 77.1km away from the true position (-45.8639, 170.5133). The standard deviation of the latitude states in the chain is 0.6716 degrees ($=$ 74.70km) and the standard deviation of the longitude is 1.7238 degrees which at -45.4948 degrees latitude is equivalent to 134.6km. These are typical Doppler-fix values, as outlined in Section 3.

The means and standard deviations calculated from the Doppler-fix Markov-chain are used to approximate Gaussian distributions for each parameter. These are passed to the next step of the fix process, the phase-fix, as prior distributions. If the standard deviations are over a certain limit, the Doppler-fix can be flagged as a poor fix and further processing can be aborted.



Figure 4: Running a Markov chain of length 20000 (excluding burn-in) on the Doppler likelihood (of Figure 3), we can plot histograms of the states occupied by each parameter. These histograms approximate the parameter's marginal distributions. The histograms for the latitude and longitude states in the example can be closely approximated with Gaussian distributions, the means of which are seen to lie very close to the true parameter values (indicated by the red lines).

2.2 *The Phase Fix*

The phase-fix step is conceptually equivalent to the Doppler-fix, a new (Phase) forward map has been created to reflect a data reduction process. First, we have some prior knowledge of parameters (the Gaussian distributions produced from the Doppler-fix). The likelihood function associated with the Phase forward map is then explored using the random walk MCMC, the resulting chains are used to approximate Gaussian distributions for each parameter and finally these are passed as prior information to the next and final step, the raw-data fix. The reduced data used in the phase-fix step, denoted $\vec{y}_p$, is derived from the same C/A code correlations used in the Doppler-fix. In this case, however, for the Doppler shift value which produces the highest correlation magnitude, we take the entire correlation vector and use the C/A code phase information it contains.

Just as in the Doppler-fix, the new Phase forward map F_p is created to simulate the reduced data for a given state of the input vector $\vec{x}_p$. The parameters of $\vec{x}_p$ are the receiver's latitude, longitude, altitude and time plus the global C/A code

phase offset parameter δ_t, which was introduced in the original forward map outlined in section 1.1. The likelihood function L_p for the phase-fix,

$$L_p(\vec{y}_p|\vec{x}_p) \propto \exp\left(-\frac{[\vec{y}_p - F_p(\vec{x}_p)]^t[\vec{y}_p - F_p(\vec{x}_p)]}{2\sigma_p^2}\right), \quad (13)$$

has structure (conditional on the set altitude, GPS time and C/A code phase offset) shown in Figure 5 for the example data file.



Figure 5: Phase likelihood distribution of latitude and longitude parameters (conditional on set values of every other forward map parameter) with (roughly) Gaussian shaped mode. The latitude and longitude axes are relative to the true receiver position. Maximum value is 217m away from the receiver location.

We explore this likelihood using a random walk MCMC. The fix given by the latitude and longitude means (-45.8642, 170.5126) is now 61.9 meters away from the true position. The standard deviations of the latitude and longitude are 262 meters and 359 meters respectively. Again these are typical values for the phase-fix step. We can now estimate the altitude and GPS time more accurately than our prior knowledge. The mean of our GPS time offset tells us (with an associated uncertainty given by the standard deviation) the error in the receiver's onboard real time clock. As the altitude's pdf is not well approximated by a Gaussian, the mean and standard deviation of the altitude parameter in the chain are poor estimators of the true altitude and its uncertainty. Instead, a other statistics of the distribution should be used.

2.3 The Raw-Data Fix

The final stage of the positioning algorithm is the 'raw-data' fix. Here we compare the binary I and Q data, captured from the GPS receiver with the data simulated from the forward map from Section 1.1. Figure 6 shows the structure of the likelihood function around the known receiver position. The raised ridges correspond to wavefronts of signals from each satellite, they should all overlap (increasing the probability density) at the true recevier position.



Figure 6: Raw data likelihood space, centered on known true receiver postition, showing only latitude and longitude distributions conditional on all other parameters set to their known/estimated values. The raised ridges here can be thought of as wavefronts produced from each satellite.

3 Results

From the 12th to the 21st of May 2010, a GPS receiver was left on the roof of the Science III building at the University of Otago. Roughly every hour, the receiver captured and stored 4ms of binary data, sampled at 8.184MHz. In total 220 data files were created.

3.1 Doppler Fix Results

The latitude and longitude means for each Doppler fix are shown in Figure 7, where the majority are seen to lie within a 100km radius of Dunedin city.

Using the known position of the receiver, a latitude and longitude error was calculated for each fix. These are plotted in a histogram in Figure 8 along with a histogram of the latitude and longitude standard deviations. We see the latitude and longitude standard deviations are largely concentrated around 25km and 70km respectively. The absolute position errors (excluding altitude), ie. distances from estimated lat/long to true lat/long for each fix, have a mean of 40.2km and a standard deviation of 65.6km.

Doppler-fixes with large standard deviations ($\sim$ 900km away) lead to a wildly incorrect phase-fix and are excluded from further processing. Based on the 300km seperation between adjacent C/A code phase wavefronts, a reasonable cutoff for longitude standard deviations in the Doppler stage



Figure 7: Doppler-fix positions from 220 fixes. The majority of position estimates are congregated around Dunedin city.



Figure 8: Top two histograms show the error magnitudes for the latitude and longitude position estimates for each of the 220 fix files. Bottom two histograms show the latitude and longitude standard deviations for each fix.

may be 200km. Using this value excluded 8 of the 220 data files from the phase-fix stage.

3.2 *Phase Fix Results*

The phase-fix results from the 212 data files that met the criteria for an acceptable Doppler-fix are overlayed on a map of Dunedin, shown in Figure 9. Here we see they all lie within the city, with the worst fix (upper right corner) being only 7.7km from the known position.

Implementing a standard-deviation threshold of 1000m as an acceptance criterion leaves 205 of the 212 phase-fixes to be passed to the raw-data fix stage. The absolute position errors (again excluding altitude) for these 205 'good' fixes

Figure 9: phase-fix positions from the 212 accepted data files overlayed on a map of Dunedin city. With the exception of one major outlier in the top right corner of the image, most fixes are in the city center (where the University of Otago campus and thus the GPS receiver are located).

have a mean of 131.9m and standard deviation 134.8m, an improvement upon the Doppler fixes by a factor of over 300. Though choosing suitable cut-off criteria at each stage may seem critical, it is really only a matter of convenience for a specific application. Should a particularly bad fix happen to meet the chosen criteria, it will most likely produce a fix in the following stage with a relatively high standard deviation, at which point it can be dealt with accordingly.

3.3 *Raw-Data Fix Results*

The raw-data state is final stage in the fixing process. When evaluating the Doppler and phase-fixes, 4ms of data (sampled at 8.184MHz = 32736 bits) was found to be a suitable amount to provide reliable results.

Cutting out all fixes with a latitude or longitude standard deviation greater than 120m, 185 out of the original 220 remain (84% 'successful') and are shown overlayed on a map of the University of Otago campus in Figure 10. The position estimates for the remaining 185 are all seen to lie within campus, where a red cross marks the true receiver postion. The absolute position errors for these 185 fixes have a mean of 50.8m and standard deviation 29.7m.

4 Conclusions

Using Bayesian inference techniques to get the most information from limited data has proven to be a very successful approach to GPS position estimation. Signal processing methods often provide quick and easy solutions to problems like GPS positioning. In situations where processing power is readily available, Bayesian analysis of data can usually



Figure 10: Final fix positions from the 185 accepted data files. All fixes lie within the University of Otago campus and are clustered around the Science III building (indicated with a red cross), the location of the GPS receiver.

give more in depth results. The future work discussed introduces ample room for improving the results obtained and the methods employed can be easily adapted to suit any number of similar applications.

References

[1] T. Molteno, "FastFix GPS positioning algorithm, unpublished.".

[2] W. Bolstad, *Understanding computational Bayesian statistics*. Wiley, 2010.

[3] C. Fox, S. Tan, and G. K. Nicholls, *Inverse Problems, lecture notes for ELEC404*. University of Otago, 1997.

[4] G. D'Agostini, *Bayesian reasoning in data analysis: A critical introduction*. World Scientific Pub Co Inc, 2003.

[5] D. Gamerman and H. Lopes, *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*. Chapman & Hall/CRC, 2006.

[6] K. Borre, D. Akos, N. Bertelsen, P. Rinder, and S. Jensen, *A software-defined GPS and Galileo receiver: a single-frequency approach*. Birkhauser, 2007.

[7] J. Zhang, K. Zhang, R. Grenfell, and R. Deakin, "Gps satellite velocity and acceleration determination using the broadcast ephemeris," *The Journal of Navigation*, vol. 59, no. 02, pp. 293–305, 2006.

[8] G. Xu, *GPS: theory, algorithms, and applications*. Springer Verlag, 2007.

[9] E. Kaplan and C. Hegarty, *Understanding GPS: principles and applications*. Artech House Publishers, 2006.

[10] A. R. Corporation, *Navstar GPS Space Segment/Navigation User Interfaces, Interface Specification IS-GPS-200D*. US Air Force, Apr. 2000.

[11] B. Carlin and T. Louis, *Bayesian methods for data analysis*, vol. 78. Chapman & Hall/CRC, 2009.

Low-power Telemetry System for Wildlife Tracking

Bradley Panckhurst, Phill Brown, T. C. A. Molteno*

*Department of Physics, University of Otago
Box 56, Dunedin, NEW ZEALAND.
Email: tim@physics.otago.ac.nz

Abstract: This paper is about the design, development and testing of a telemetry system for wildlife tracking devices to replace GSM. Various frequency bands were considered and the 434 MHz and 868 MHz SRD bands were chosen using Frequency Shift Keying modulation. A daughter card was designed, developed and tested in both clear line of sight and foliage situations. We measured a range of 400 m in foliage and 3 km clear line of sight.

Keywords: Wildlife tags, Telemetry, GSM, Short Range Devices

1 Introduction

This paper outlines the design, development and testing of a low-power wildlife-tag telemetry system based upon frequency bands allocated to Short Range Devices (SRD). The current Otago wildlife tagging system uses Global Positioning System (GPS) to calculate the animals position, then uses the GSM (Global System for Mobile communications) telemetry system to transmit the collected data back to a central base station. Figure 1 shows a tagged Royal Albatross with a GSM telemetry enabled wildlife tag on its back.

GSM operates worldwide, however the infrastructure is concentrated around high human population densities. The likelihood of remote wildlife populations having GSM coverage is very slim. Figure 2 shows the GPS data collected from the Royal Albatross tagging project. Each colour represents a different albatross, the GPS data shows that the albatross are capable of traveling a long way out to sea over a 1000 km where there is no GSM coverage. This is where the GSM telemetry system for the wildlife tags fails, it requires a significant amount of power when attempting to transmit the collected data back to the central base station when there is no GSM coverage.

Because of these problems alternative telemetry systems are being considered using the SRD frequency bands. SRD frequency bands are covered by a General User Radio Licence, therefore they are free to use and allows for a wildlife telemetry system to be created in tagging locations. The constant development in wireless technology means that there are a number of high performance radio ICs capable of telemetry. These are small, have low-power consumption and can transmit as far as 10 km, making them a good candidate for the replacement of the GSM telemetry system.



Figure 1: The GSM telemetry enabled wildlife tag on the back of a Royal Albatross. Photo courtsey of Keith Payne.



Figure 2: Map of GPS data received from the Royal Albatross tagging project

2 GSM Module Problems

The GSM module part of the wildlife tags has four main problems; Reception, Power Consumption, Weight and Administrative Overheads.

2.0.1 Reception

The current GSM module utilises the existing mobile communication infrastructure to send the data back to a base station which gathers it into a central database. This works well in places where there is coverage. Where the tags are currently deployed; using the GPS/GSM tags does not provide the best use of the tag's ability to send data back to the base station, and diminishes the effectiveness of the wildlife tags as it adds weight and increases battery usage for a telemetry function that is not able to be used.

2.0.2 Power Consumption

The GSM module uses a Telit GE865-QUAD, designed for extremely compact application due to it low profile and small size. It has an extended supply voltage range of 3.22 - 4.5 V DC. Due to a high switched off average current, of up to 62 μA. The whole GSM module is turned off until the tag attempts to send data back to the base station, which happens after the GPS fixs.

The GSM system is made in a way that the RF transmission is not continuous, and the relative peaks can be as high as approximately 2 A. These high peak currents drain the battery quickly, meaning that the ability for data retrieval by the GSM network is restricted.

2.0.3 Weight

In New Zealand, the Department of Conservation recommends an upper limit of 3% of the animal's body weight for a tracking tag. [1] The current weight of the GSM module shown in Figure 3 is 5.74 grams but this does not include the antenna or the SIM card which is required for the module to work. Of the 5.74 grams, 3.2 grams is attributed to the weight of the Telit GE865-QUAD module. This is the heaviest component of the existing tags.

2.0.4 Administrative Overheads

Using the GSM network requires a substantial amount of administration work. Each sim card, associated with each wildlife tag, has to first be registered to the network provider. Then money has to be added to the registered account for the module to be able to send data via SMS to the base station.

Figure 3: Telemetry board from a GSM tag without antenna or SIM card. The board is dominated by the GSM module (large rectangle on the left).

3 System Design

There is a wide range of radio bands for telemetry; for a replacement to the GSM module radio bands covered by a general user licence were considered. The decision to use the SRD radio bands was made as they overlap a number of ISM radio bands with the same output power constraints. At the same time providing a larger range of bands in which to work.

3.1 SRD Frequencies

SRD is intended for a broad range of uses, such as [2]:

- Access control (including door and gate openers)
- Local Area Networks
- Medical implants
- Radio frequency identification
- Smart sensor networks
- Telemetry

In New Zealand, short range devices are covered under the licence name General User Radio Licence for Short Range Devices. A number of general conditions apply to all transmissions under the General User Radio Licence for Short Range Devices, one of which states that the maximum power assigned to the respective frequency must not be exceeded. There are many frequencies available for short range devices: from 9 kHz to 245 GHz [3]. Not all of these frequencies may be used for the wildlife telemetry tag, as either their use is limited to some other function, or certain frequencies would not be desirable due to antenna length / power usage. Table 1 shows the certain frequencies which would be suited to replace the GSM module.

Low (MHz)	High (MHz)	Reference Frequency (MHz)	Maximum Power dBW e.i.r.p.
433.05	434.79	433.92	-16.0
864.00	868.00	866.00	6.0
921.50	928.00	924.75	6.0
2400.00	2483.50	2441.75	6.0

Table 1: Selected Short Range Device frequencies [3]

3.2 Texas Instruments CC1120

Several SRD frequency band radio transceivers were considered. As a base for the new micro-power telemetry (μPT) module, the Texas Instrument CC1120 was chosen as it is a small package chip with ultra-low power and the largest claimed transmission range of the options considered.

Of the many features and functions this chip has, the key ones for the wildlife application are:

- Wide supply voltage range (2.0 V - 3.6 V)
- Low current consumption:
  - RX: 22 mA peak current in high performance mode
  - TX: 45 mA peak at +14 dBm
  - Power down: 0.3 μA
- Enhanced Wake-On-Radio functionality for automatic low-power receive polling
- Support for auto-acknowledge of received packets
- Support for re-transmission

3.3 PCB Design

Based upon a four layer reference design from Texas Instruments, the module's size was dictated by the current generation of wildlife tags. The PCB layout reduced the reference design to two layers and was constrained by the standard module boundary and connector . The largest component on the μPT module is the connector, the next is the CC1120 transceiver chip; the other components (compared to these two) are insignificant in physical size. A 50 Ω antenna was used, so a coplanar waveguide was employed for impedance matching. The tags were machined out of 1.6 mm PCB using a LPKF Protomat S-62 PCB prototyping machine, then with the exception of the CC1120 transceiver chip the components were soldered by hand. The CC1120 chip was reflow soldered, as the bottom of the chip is a ground connection, inaccessible when hand soldering. The completed tag is shown in Figure 4.



Figure 4: Close up on the μPT module prototype, centred on the Texas Instrument CC1120, for size reference the chip is 5mm x 5mm.

4 Testing and Results

The μPT module was compared against the GSM module weaknesses and tested to quantify its effectiveness in the two different scenarios it will face: foliage, and unimpeded line of sight tests.

4.1 Power Usage

Overall the μPT module power usage is a lot lower than the GSM module. This is because the GSM system has relative peaks as high as 2 A, compared to the μPT module where the highest current draw occuring at full transmit output power (+14 dBm) is only 45 mA. As the μPT modules will utilise Wake on Radio they will be consistently on, and checking for a data packet every one second, the average current comsumption is 15 μA. This can be lowered further by checking less frequently.

4.2 Weight

The current GSM module has a weight of 5.74 grams, not including the antenna or the SIM card. The first generation μPT module improves on this with a weight of 3.25 grams, and the second generation μPT module its over 50% lighter than the GSM module at 2.82 grams (excluding the antenna). This weight could be improved further by using the 0.5 mm PCB thickness used for the GSM module, the low-power telemetry module prototypes are currently construction using 1.6 mm PCB. This would be expected to further reduce the weight by 0.5 grams.

4.3 Transmitted Power

The transmit output values were set with the power at +14 dBm and with a center frequency of 868 MHz. The evaluation module spectrum had a peak centered at 868 MHz and has an output power of +14 dBm. The μPT module shows a peak centered at 868 MHz; however the



Figure 5: Results of line of sight distances around Dunedin with path-loss model. The Friis Transmission equation closely models the evaluation module when the antenna is in a vertical orientation. From a received signal strength of approximately -100 dBm the packet error rate starts to increase, and both of the whip antennas are affected by high packet error rates.

transmit output power is +12 dBm. This means that, compared to the evaluation module, along the RF path from the CC1120 chip to the antenna there is a 2 dBm loss.

4.4 Receiver Performance

The evaluation module reported a received signal strength indication of -7 dBm, when connected to a transmit signal of -16 dBm. However the value in the received signal strength indication registers is a raw value. To obtain the correct level of the applied signal a offset has to be added [4]. This means that for the setup used the offset which corrects the registers value is off by 9 dBm.

The result of the μPT module reported a received signal strength indication of -17 dBm, when connected to a transmit signal of -16 dBm. As the register set up for both of these modules was the same, and even though the received signal strength indication was similar to the transmitted signal, the offset is incorrect by 9 dBm. Along the RF path from the antenna to the CC1120 chip there is a loss of 10 dBm compared to the evaluation module.

4.5 Range Testing

The TI CC1120 chip claims a line of sight range of 10 km. As there is some RF path-loss within the μPT module, measuring a maximum range is important as it allows for the deployment of base stations to effectively cover the required area. Due to the geographic layout of Dunedin, the sites were chosen as they provided clear line of sight to the target evaluation module that was located on the top of the Science 3 building at the University of Otago.

Figure 5 shows the results with a fitted path-loss model for the received signal strength indication versus distance and with the packet error rate versus received signal strength indication. The model used is the Friis Transmission Equation [5], commonly used in long-distance radio communication. For the equation, values from the 868 MHz evaluation module were used. Figure 5 shows that the Friis Transmission equation closely models the evaluation module when the antenna is in a vertical orientation. This means that if we know the gain of the antennas (transmit and receive), the wavelength, and the transmitted power, justified estimates based upon this model in free space can be made.

Note that there are number of tests with a 100 % packet error rate, the received signal strength indication for these



Figure 6: Results of transmission through foliage along the 400 m path, origin 20 cm above the ground, with foliage path-loss models. The foliage path-loss models show that for path-loss through foliage only, the FITU-R + Fitted ground reflection model fits the data best. However, the Weissberger model better represents results with the inclusion of the free space over the reservoir as well as foliage cover from the 250 m test site onwards. At around a received signal strength of approximately -110 dBm the packet error rate starts to increase, and both of the whip antennas are affected by high packet error rates.

has been added into Figure 5. However, they do not provide an accurate assessment signal strength of as no packets were received.

4.6 Performance in Foliage

As not all wildlife live in an environment which is "free space", testing was conducted in foliage. The three test paths were repeated twice; one just above the ground and the other in a tree approximately two metres high.

The first foliage path-loss model is Weissberger's modified exponential decay model [6]. It is calculated with the inclusion of free space path-loss, the Friss Transmission equation and is valid for frequencies from 230 MHz to 95 GHz. It is applicable for areas where the path is blocked by dense, dry, in-leaf trees found in temperate climates, similar to the conditions found in New Zealand.

The second foliage path-loss model is the Fitted ITU Recommendation (FITU-R) model combined with a Fitted Ground Reflection model [7]. The Fitted ITU Recommedation model used in the conjunction with the overall model, was optimized using the least squared error fit for several sets of measurement data with in-leaf foliage[7].

The ground foliage test (shown in Figure 6), led to a number of conclusions. The first is that for continuous foliage paths, the FITU-R + Fitted ground reflection model accurately predicts the path-loss. The second is that terrain which blocks the line of sight transmissions will cause essentially 100 % packet loss, where only the evaluation module was able to receive any packets, (and it only received two). The third is that given some area of free space transmission in the RF path, the Weissberger model more accurately predicted the path-loss compared to the FITU-R + Fitted ground reflection model.

The tree foliage test led to two conclusions. The first is that given some area of free space transmission in the RF path the Weissberger model more accurately predicted the path-loss compared to the FITU-R + Fitted ground reflection model. The second observation is that for continuous foliage paths with the target, neither model accurately predicted the path-loss.

5 Conclusions and further work

This project has shown that the μPT module is capable of replacing the current GSM module. The module is small, cheap and as required has a lower power consumption, but still provides a useful line of sight range, and could penetrate foliage sufficiently well to allow the retrieval of data from tags by aircraft, or from the ground at ranges less than 200 m.

5.1 Future Work

The first step towards integrating the μPT module with the current wildlife tags is the completion of the software.

In future work, the PCB would be redesigned to more effectively reduce the transmission loss. This could be done by improving the impedence of the microstrip so it is as close to 50 Ω as possible.

The CC1120 supports a number of other amplitude and frequency shift modulation formats. So, the range and foliage tests would be repeated, this time with the intention of employing the modulation format that best suits both scenarios.

References

[1] N. Beausoleil, D. Mellor, K. Stafford, *et al.*, "Methods for marking New Zealand wildlife: amphibians, reptiles and marine mammals.," *Methods for marking New Zealand wildlife: amphibians, reptiles and marine mammals*, 2004.

[2] ETSI, "Short Range Devices." `http://www.etsi.org/website/technologies/shortrangedevices.aspx`.

[3] M. of Economic Development, "Short range devices." `http://www.rsm.govt.nz/cms/licensees/types-of-licence/general-user-licences/short-range-devices`.

[4] T. Employee, "CC1120 RSSI threshold." `http://e2e.ti.com/support/low_power_rf/f/155/t/131734.aspx`.

[5] J. Richards, *Radio Wave Propagation: An Introduction for the Non-specialist*. Springer, 2008.

[6] M. Weissberger, "An initial critical summary of models for predicting the attenuation of radio waves by trees," tech. rep., DTIC Document, 1982.

[7] Y. Meng, Y. Lee, and B. Ng, "Empirical near ground path loss modeling in a forest at VHF and UHF bands," *Antennas and Propagation, IEEE Transactions on*, vol. 57, no. 5, pp. 1461–1468, 2009.

Bachelor of Engineering Re-design with a Project Spine at Massey University

Amal Punchihewa

School of Engineering and Advanced Technology
Massey University, Palmerston North
g.a.punchihewa@massey.ac.nz

Abstract: This short paper presents an outline of project spine introduced by School of Engineering and Advanced Technology, Massey University in New Zealand. Though it is too early to draw any conclusions about the effectiveness of proposed curricula for Bachelor of Engineering honours programmes, this paper attempts to introduce some key attributes and background to the project spine. Though there were mixed reactions from students and staff about the project based learning and outcome of the proposed scheme, generally both stakeholders were satisfied with many positive outcomes. Some of the issues such as access that the students have to resources and staff workload have been identified and they will be addressed in future.

.

Keywords: Engineering education, project led education, project spine, capstone project

1 INTRODUCTION

Engineering Education in general has been undergoing many changes. Large number of universities around world including the ones in top 50 have been addressing topic of education in Science, Technology and Engineering including Mathematics (STEM). This short paper wishes to inform the electronics community in New Zealand the initial steps in re-design of Bachelor of Engineering (BE) degree programme by the School of Engineering and Advanced Technology (SEAT). Students are exposed to a selective range of principles and concepts in five areas; namely,

1. Projects (Design & Build)
2. Mathematics
3. Engineering & Technology Principles
4. Computing & Electronics
5. Chemistry & Biology

Projects are core component and common across all engineering disciplines.

It is important to remind what the focus of engineering educations is. Though there can be many definitions and descriptions by various scholars, generally, an engineer is a technical problem solver - designing large systems such as road, power, gas, water, communications networks or large industries, utilities; building large systems like power plants, dams, server banks, etc.; operating large engineering systems like communications services, wide area networks, power systems, etc.; or managing systems and services. Therefore, an engineer should be a designer, thinker, innovator and systems integrator. Hence, the educational system should inculcate into a student various aspects like engineering principles, standards and practices, design methodologies, modelling and optimization capabilities, systems analysis and integration techniques. An engineer can become a thinker, creative person or innovator only if he/she is allowed to independently put together all aspects of learning to solve a practical problem and look at alternatives [1]. Figure 1 depicts the conceptual framework that SEAT used while developing the new curricular.



Figure 1: Conceptual framework

In each semester, as in most of the universities, four papers or subjects are offered to the students. In this document, the term paper will be used here after.

2 PROJECT SPINE

At the School of Engineering and Advanced Technology of Massey University in New Zealand, in February 2012, projects were introduced for the first time in BE programmes that run in all four years of Bachelor of Engineering honours programmes. Figure 2 summarises the focus in each year. There will be an independent project in each semester of first year. In year one the focus is on developing self, explore creativity, with an emphasis on design; all within a global context.

There is no constraint on a company or product. It was decided that there should be emphasis in assessment on discovering why students' ideas did/didn't work rather than assessing if the solution's not technically capable. It was identified that time is the only constraint.

In year two the focus is on developing their skills in product design and development, and manufacture. The context is within a company & industry environment.

They will be exposed to how companies work i.e. products, markets and financials. There would be an expectation of demonstration of technical competence, albeit with simple products or processes or sub-systems.



Figure 2: Project spine over four years

2.1 General Objectives of Project Work

a. Capability to work independently - to think, to conceptualize, to design, to operationalize, to diagnose on failures and to innovate.
b. Learn to work as a team - sharing work amongst a group, learning human behaviour, etc.
c. Learn leadership qualities.
d. Learn to solve a problem through all its stages by understanding and applying project management skills.
e. Learn to do various implementations, and learn problems of system integration, fabrication, testing and trouble shooting.
f. Learn about some specific technology that is under development (like mobile communications, blue tooth, etc).
g. Learn about issues and problems in prototyping and applications of technologies.
h. Learn communication skills.
i. Learn report writing skills.

2.2 Stages in a Project Work Activity

The following are typical steps that can be found in a project work activity. Though many stages are listed, some may not be needed for some types of project work whereas may need another step.

I. Project topic selection.
II. Literature survey and analysis.
III. Problem definition, including inputs, analysis, hardware / software / testing components required, functions, requirements, responses, outputs and deliverables, diagnostics, constraints, scope limitations, types of users, etc.
IV. Rapid prototyping initiated.
V. Design - preliminary, logical, physical and detailed. Design approaches.
VI. Rapid prototype seen.
VII. Implementation - fabrication, programming, system integration.
VIII. Testing
IX. Operationalization
X. Field testing

XI. Experimentation - performance measurements, diagnostics, etc.
XII. Project report preparation.



Figure : 3 Graduate Engineer as SEAT-Massey like to perceive - capstone and pillars of learning

Figure 3 shows the five key attributes that a graduate engineer should possess in the 21st century. Professionalism and approach to systematic problem solving abilities are key pillars of a globally recognised graduate engineer.

3 SEMESTER ONE:

In semester one, the project was based on a developing country and activities of organisation known as Engineers without borders. The region selected was Mekong Delta in Vietnam.

3.1 Engineers Without Borders (EWB)

EWBNZ is an organisation of professionals and students who share a vision to confront global challenges of poverty, sustainable development and social inequity by undertaking initiatives that will directly improve the quality of life in communities within New Zealand and in the South Pacific region [2] http://www.ewb.org.nz/.

3.2 Context of Mekong Delta Project

The Mekong Delta in Vietnam is an underprivileged area of Vietnam where many people struggle with the simple things such as food, water, shelter and electricity. Massey University, together with Engineers Without Borders (EWB) wants to provide engineering solutions which will improve the quality of life of people living in the Mekong Delta region. The region is very vast and a range of different ways of life is present. The region is among one of the largest producers of rice in Vietnam.
The people living in the Mekong Delta make their living as farmers and fishermen. Living on the edge of the rivers or canals is common in various structures built from whatever materials found. Consequently, the architecture along the delta varies from place to place.

Students were provided with a detailed context as shown in next paragraph as an example for the problems related water. For each of other problems such as housing, energy, cooling, lighting and food, students were asked to address different context and problems to find a solution in a selected area.

Access to clean drinking water is an important issue that needs to be addressed throughout the Mekong Delta region. Drinking water is currently collected from three sources: rain water, bore wells and surface water. Rainwater is collected in open cement / ceramic jars during the wet season and will typically provide a family water for approximately six months of the year. Impurity build up in the jars is a problem. 85% of rural households have access to borehole wells with a hand pump. Well water often contains heavy metals including arsenic and is often highly saline. Surface water is collected from rivers and ponds and is typically used for drinking and cooking, but is dirty, contaminated with chemicals and has a high salt content. Although the region does not experience severe flooding on the scale of a tsunami, it does experience daily tidal variations of about 7cm. For those living close to the water table these fluctuations present a challenge to the design of effective sanitation systems. Managing this issue is important for preventing breakouts of disease due to poor sanitation. Most people in rural areas do not have access to a toilet and go outside directly into the river, pond or field near the house.

3.2.1 Assessment criteria

The solutions were assessed by a panel of two engineers who have not been supervising the projects. This is an oral assessment involving all members of the team. The teams requested to be prepared to discuss their solution with a moderating panel and other staff interested in the projects.
The panel's assessments were based on how well the team specified the problem, and can demonstrate, discuss and justify its solution. The panel assessed the solution's potential for wealth creation and feasibility in three areas:

I. Engineering
II. Contextual
III. Financial, social, and environmental.

Some teams developed physical models to help demonstrate, justify and discuss their solution, and they equally used other forms of modeling, testing and communication techniques to demonstrate, justify and discuss their solution. The emphasis was placed on the evidence generated from the process of research, developing, testing, and improving the solution (such as: project books, calculations, concepts, iterations of development etc.) rather than one off outputs presenting the solution such as posters etc.

3.2.2 The team will be assessed on:

- How well the team has specified the problem and been able to develop a solution that meets these specifications. (30%)
- How the team's solution might potentially create wealth for the village stakeholders in terms of financial wealth, and/or social wealth, and/or environmental wealth. (10%)
- How well the team demonstrates / discusses / justifies engineering principles and technical feasibility. (20%)
- How well the team demonstrates / discusses / justifies the solution's contextual feasibility. (20%)
- How well the team demonstrates / discusses /justifies the solution's financial, social, and environmental feasibility. (20%)

3.3. Project charter template for Mekong Delta Project

Students were provided with a template for completion of the Project Charter for individual projects. They were referred to a Project Charter document that addresses the questions; what is it?, what does it achieve?, and what should it contain?. They were required to provide a description of each of the sections. Students are also provided with the start of a project charter based on a water pump project. Water pump project contains some information to start each section of their own charter. This was designed as a basic model for what should be included. Students were asked to start with their own specific project once this is defined. Throughout the project they were asked to fill out their Charter with more information. By Week 6, they were requested to have substaintially developed their Charter to a level where it proves sufficient content to assess the direction and potential for their project.

Following guide lines were provided to students through online site '*Stream*' for writing the report. Staff from communication department of Massey provided information, tutoring and evaluation of student written communication in form of a report.

1. Title page and table of contents
The name of the project
The client for the project – Engineers without Borders
Your team number and member names and student IDs

2. Project authorisation
Water pump project, Team x, project paper 228.111
Project Director – Professor Allan Anderson, School of Engineering & Advanced Technology Supervisor – Dr Who XXXX?

3. Project context

4. Key stakeholders
People of Anh Minh district within the Kien Giang province on the Mekong Delta – will use final solution
Habitat for Humanity Engineers – will implement final solution

5. Project outcomes
Access to clean drinking water to the people of village x in the Mekong Delta.

6. Project constraints
• Any chemicals used must not have an adverse effect on the environment.
• Solution must conform with Cultural beliefs and practices of the people living in the community.
• Cost associated with implementing any proposal must be less than US$1000.
• Solution is compatible with Infrastructure already in place for water and sanitation.

7. Project scope
Project solution to focus specifically on the needs and capabilities of Anh Minh district within the Kien Giang province on the Mekong Delta.

8. Project risks
• Lack of technical expertise among team members
• The time schedule of completing work will be insufficient, leading to failing to complete the project
• Inexperience of the team members at leading and completing projects
• Environmental factors and climate of the Mekong Delta region are not well enough understood.

9. Project management
A project plan – you will be introduced to the Fever chart method in weeks 2 and 3.
In this section they were asked to include your basic plan and how you will manage the project throughout.
10. Project deliverables

Intermediate deliverables:
I. Project plan
II. Project charter

End deliverables:
I. Final report
II. Recommended solution

11. Change control
Changes to the project must be approved by the project supervisor and key stakeholders must be contacted if required.

12. Health and safety
Ensure familiarity with safety requirements in all work spaces – workshops, labs etc.

13. Communication plan
I. Weekly meeting with supervisor
II. Presentation of Charter
III. Presentation of solution
IV. Submission of final report

4 SEMESTER TWO - Creative Solutions

4.1 Outline of the Content:

Engineering is the application of mathematics, science and technology principles, integrated with business and management, to develop and provide products and processes for industry and the community. This paper allows students to develop their creative skills and encourages a thirst for knowledge. The development of engineering design, prototyping, teamwork and communication through design form an integral part of this project focused paper.

4.2 Learning Outcomes (LOs):

Following are the learning outcomes expected from students by actively working on this creative project work.

On completion of this paper, a student will be able to:

I. Apply science and engineering principles to the solution of a complex engineering problem – where complexity is defined by incomplete and ambiguous information
II. Interpret a problem demonstrating analysis of uncertainty and creative thinking in the approach to solving a problem
III. Interpret and synthesise relevant information and data from written and oral sources.
IV. Use a systems thinking approach to define the problem
V. Recognise self-reflection to increase awareness of one's personal skills and attributes
VI. Design an effective team and demonstrate its successful operation
VII. Practice electronic/multimedia and graphical communication styles
VIII. Use project management tools to plan and manage activities and apply the design process to a relatively simple project

As the name, creative solutions, suggests, second semester paper is all about creativity. Students are encouraged to dream a little about what the future holds and design something (a product or process) that might be used in this future scenario. Students will be introduced to some key tools including computer aided design, visual communications, creative thinking, foresighting and creative design, all of which will contribute to their final project outcome - a truly innovative concept for the future that challenges, surprises and impresses visitors to an end of project exhibition.

4.3 Assessment

Key point to note in this paper is the knowledge and exposure given to the students on wide range of tools that can be used for designing and visualising conceptual ideas.

Table 1 Mark allocation and learning outcomes tested

Assessment	LOs	Individual	Group
1. Recap Quiz	1,2,3,4	10%	
2. Practical Exercises	1,3	20%	
3. Project portfolio	5,6	10%	
4.Exhibition/presentation	6,7		20%
5. Report	1,2,3,4,5,6,7,8	20%	20%

Industry experts evaluated the visual and presentation skills of students on the final day of project work followed by formal evaluation of student group reports by Massey staff. Prior to this, individual learning was assessed using a quiz, practical exercise and a project portfolio.

4 CONCLUSIONS

As interactions, engagements and involvements in project based learning offer many opportunities to the students for learning, there are many universities those have adopted project and problem based learning. The School of Engineering and Advanced technology (SEAT) at Massey University has rolled out the first year of project spine in 2012 smoothly. Staff involved in teaching and supervision in general noted that students have demonstrated adequate skills in using engineering tools and conceptual knowledge, analysing problems and communicating solutions in written, oral and visual media effectively.

5 REFERENCES

[1] Subramaniam, D. K., Project work-some issues, 2006.

[2] http://www.ewb.org.nz/ visited on 29th October 2012.

6. BIBLIOGRAPHY

[3] Diamond R. & Adam, B. 1993. Recognizing faculty work: Reward systems for the year 2000. San Francisco, CA: Jossey-Bass

[4] Boyer, Ernest L. 1990. *Scholarship reconsidered: Priorities for the* professoriate. Princeton, NJ: The Carnegie Foundation for the Advancement of Teaching.

[5] National Research Council. 2002. *Scientific research in education.* Committee on Scientific Principles in Education. Shavelson, R.J., and Towne, L., Editors. Center for Education. Division of Behavioral and Social Sciences and Education. Washington, DC: National Academy Press.

[6] Shulman, Lee S. 1999. Taking learning seriously. Change, 31 (4), 11-17.

[7] Smith, Karl A., Petersen, Renee P., Johnson, David W. & Johnson, Roger T. 1986. The effects of controversy and concurrence seeking on effective decision making. *The Journal of Social Psychology*, *126* (2), 237-248.

[8] Wankat, P.C., Felder, R.M., Smith, K.A. and Oreovicz, F. 2001. The scholarship of teaching and learning in engineering. In Huber, M.T & Morreale, S. (Eds.), Disciplinary styles in the scholarship of teaching and learning: A conversation.

[9] Schachterle, L., and Vinther, O., Introduction: The Role of Projects in Engineering Education, European Journal of Engineering Education, Volume 21, Issue 2, 1996, DOI: 10.1080/03043799608923394, pp 115-120.

[10] McDermott, K. J., and Machotka, J., Enhancing final year project work in Engineering programmes, Global joournal of Engineering education, Vol.10, No.2 Australia, 2006.

[11] Hudson, C. & Harding, P., Improving project work oj degree courses, The Institution of Engineering and Technology (IET), 2002.

Swap moves for MCMC applied to electrical impedance tomography

Erfang Ma[1] and Colin Fox[2]

[1] Xi'an Jiaotong-Liverpool University, Suzhou, China
[2] Electronics Group, Department of Physics, University of Otago, PO Box 56, Dunedin, New Zealand

Email: affinema@yahoo.com.cn & fox@physics.otago.ac.nz

Abstract: Electrical impedance tomography (EIT) is a non-invasive means of imaging the conductivity (or permittivity) distribution within an object. In the Bayesian framework, computation of 'solutions' is performed using Markov chain Monte Carlo (MCMC). Despite 20 years of development, the simplest single-site random walk update still is viewed as most efficient in some present MCMC implementations for EIT. We report on experiments utilizing swap moves in MCMC that significantly improve on the single-site update.

Keywords: Electrical impedance tomography, inverse problem, computational inference, Markov chain Monte Carlo, single-site update, swap moves

1. INTRODUCTION

Electrical impedance tomography (EIT) is a canonical inverse problem in which the electrical property within an object is inferred from electrical measurements on the surface of the object. This leads to many potential applications, when structures of interest have a contrast in electrical properties, such as scanning for breast cancers and measuring flow in milk-drying plants.

We denote the unknown electrical conductivity by x and the (noisy) surface measurements of potential at electrodes by u^{n}. EIT is known to be a severely ill-posed inverse problem [1]; that is, direct inversion to give x in terms of u^{n} gives a solution that is highly sensitive to errors on measured data, and to model errors. This can be understood in terms of the physics of electrical currents, since each measurement of surface potential depends on all of the unknown conductivity; this follows from Ohm's law that states that current follows *all* paths, in proportion to conductance of each path[1]. Hence, measurements are most sensitive to bulk properties of the conductivity, such as the mean conductivity, while localized or 'high spatial frequency' variations in conductivity have little influence on measurements.

This sensitivity to bulk processes actually makes EIT a good imaging mode for those situations where we wish to non-invasively measure bulk properties. One such application is the measurement of *total* flow of the various constituents in milk powder flowing in a pipe, which is an application we are developing at Otago. In a sense, we can think of EIT as a set of scales – that measures the total mass of an object on the scales – though with some low resolution spatial discrimination.

Here we consider a simple observation process, in which noise-free measurements at electrodes u are a known function of the conductivity x and measurement noise e is additive and independent of x, that is

$$u^{\mathrm{n}} = u + e. \tag{1}$$

The functional relationship $x \mapsto u$ is called the *forward map* for EIT. In this paper we address the problem of computing statistical answers to questions about x given measurements u^{n}, a known forward map, and knowledge of the (statistical) distribution over noise e.

2. A TEST PROBLEM

Figure 1 shows a phantom conductivity in a square region that we will attempt to recover, from simulated noisy measurements. The conductance is constant within each square pixel in a 24×24 pixel grid. Measurements consist of injecting (unit) current into each of 16 point electrodes in turn, with current removed uniformly on all electrodes, and measuring the potential ate each electrode with respect to the average potential on electrodes. This stylized problem has been studied previously in [2, 3].

[1] The common saying that that electricity travels "by the path of least resistance" is incorrect.



Figure 1: Phantom conductivity distribution, having conductance 3 (white) and 4 (black) in arbitrary units. Sixteen point electrodes on the boundary are numbered 1,...16.

We use the same numerical computation of the forward map for simulating measurements and performing the inverse problem; hence we are certainly committing every possible 'inverse crime' and so recovered conductivity distributions need to be viewed as overly optimistic. However, the speed of performing the inference is our primary concern here, and that is not affected by the inverse crime.

3. COMPUTING THE FORWARD MAP

Simulation of forward map $x \mapsto u$ involves solving the boundary value problem (BVP)

$$-\nabla \cdot (x(s)\nabla u(s)) = 0 \qquad s \in \Omega, \tag{2a}$$

$$-x(s)\frac{\partial u}{\partial n}(s) = j(s) \qquad s \in \partial\Omega. \tag{2b}$$

We solve for $u(s)$ that is the potential at location s, while $x(s)$ is the that is a given conductivity distribution. We have not yet stated a potential reference, though that is necessary to ensure the solution of this BVP is unique. The current crossing the boundary $\partial\Omega$ is denoted $j(s)$, which must satisfy $\int_{\partial\Omega} j(s)\mathrm{d}s = 0$ since current is conserved.

We solve the BVP (2) using a standard finite element method (FEM) in which the pixels are taken as elements, and potential is interpolated using bilinear interpolation (see [2, 3] for details). FEM discretization results in a symmetric sparse system of linear equations

$$Ku = f \tag{3}$$

where the global stiffness matrix K is assembled (by elements) as

$$K(x) = \sum_{i=1}^{n} x_i K^i \tag{4}$$

where x_i is the conductivity in pixel i and K^i is the local stiffness matrix for pixel i. We solve the sparse system (3) using a sparse Cholesky factorization of the stiffness matrix K followed by solution of the resulting triangular system.

4. PRIOR MODELING AND POSTERIOR DISTRIBUTION

We follow [3] and model the conductivity as continuous valued in each pixel, restricted to the interval $[2.5, 4.5]$. The 'blocky' nature of the image is modeled by the prior distribution

$$\pi_p(x) \propto \exp\left\{\beta \sum_{i\sim j} v(x_i - x_j)\right\} \mathbb{I}_{[2.5,4.5]^n}(x) \tag{5}$$

where

$$v(s) = \begin{cases} \frac{1}{w}\left(1 - |s/w|^3\right)^3 & \text{if } |s| < w \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

and $\mathbb{I}_{[2.5,4.5]^n}(x)$ is the indicator function for each component of x being in the interval $[2.5, 4.5]$. The sum is over all nearest neighbours on the pixel lattice, denoted by $i \sim j$. This prior encourages neighbouring x_i and x_j to have similar values, but once x_i and x_j are more than w apart, the penalty does not grow. This allows occasional large shifts between neighboring pixels. Thus, the effect is to prefer images with connected regions with constant value, separated by distinct jumps in conductivity. The lumping parameter β is set to 0.5 and the separation value w is set to 0.3.

The likelihood function is

$$l(u^{\mathrm{n}}|x) \propto \exp\left\{-\frac{1}{2\sigma^2}\|u^{\mathrm{n}} - u\|^2\right\} \tag{7}$$

when the noise e in Eqn 1 is iid (independent identically distributed) zero-mean Gaussian with variance σ^2. As in [3], we set σ so that the resulting SNR (signal to noise ratio) in the measurement process is close to 1000:1 to mimick practical EIT measurements.

By Bayes' rule, the posterior distribution over conductivity images is

$$\pi(x) \propto l(u^{\mathrm{n}}|x)\pi_p(x). \tag{8}$$

It is not possible to summarize the posterior distribution by analytic means, so we employ MCMC (Markov chain Monte Carlo) to sample from this distribution and evaluate summary statistics.

Figure 2 shows five samples drawn from the posterior distribution. These samples are close to each other and all resemble the original conductivity distribution, indicating that the posterior distribution is supported around the true conductivity distribution. Figure 3 shows the posterior



Figure 2: Five samples of the conductivity x drawn from the posterior distribution.



Figure 3: Posterior mean (left), posterior variance (middle), and marginal posterior mode (right).

mean, posterior variance, and the marginal posterior mode (MPM) (see [4] for a discussion on the MPM). As can be seen, the marginal posterior mode gives a useful estimate of the true conductivity distribution. The variance image shows that the primary uncertainty is in the location of boundaries of the regions of constant conductivity, while the conductivity values are well determined.

The results presented in this section were evaluated using samples drawn by the standard single-site Metropolis MCMC as in [3] that had run for $1.2 \times 10^5 \times 576$ iterations. In the following section we consider improved proposals.

5. SWAP MOVES

As the name implies, the swap moves simply swap the conductivity value in a pair of nearby pixels. We consider three specific swap moves as shown in Figure (4). They are (b) 1×2 swap between the two-nearest neighbours, (c) swap between the diagonal pixels within a 2×2 block, and (d) swap between two components within a 3×3 block as shown. Swap move (b) was used previously in [5, 2], while moves (c) and (d) are special cases of moves used in [5, 6], though in the context of a discrete representation. Also shown is (a) the single site update move.

Each of these swap moves provide a localized update that leaves approximately invariant the posterior distribution in Eqn 8. When these moves are chosen with probabilities p_1, p_2, p_3, p_4, and the Metropolis-Hastings (MH) accept/reject rule applied, the resulting transition kernel is



Figure 5: Traces of conductivity at three representative pixels for the swap move (with probabilities $(0.1, 0.1, 0.4, 0.4)$) and single-site Metropolis.

$$A(x, y) = \sum_{i=1}^{4} p_i A_i(x, y) \tag{9}$$

where $A_i(x, y)$ is the transition kernel resulting from a MH algorithm using just the i-th proposal. Since kernel $A_1(x, y)$ associated with the single-site move is irreducible and we set $p_1 > 0$, the resulting chain is also irreducible.

We tested eight sets of move probabilities listed in Table 1 and ran the resulting algorithms for $1.2 \times 10^5 \times 576$ iterations. During each run, the state-variable was recorded every 10×576 updates, while log-likelihood and log-prior were recorded every 576 updates. As in [3], three representative pixels were selected (marked by red, blue and green circles in Figure 3) and their trace plots were used for qualitative assessment of how well the Markov chain mixed. From the traces shown in Figure 5 it is clear that using the swap moves significantly improves performance of the MCMC. As in [3], we also evaluated the integrated autocorrelation time (IACT) [7] at the 'blue' pixel as a quantitative measure of the performance of the



Figure 4: Probability tree for moves: (a) random scan single-site update, (b) 1×2 swap, (c) 2×2 swap and (d) 3×3 swap, that are selected with probabilities p_1, p_2, p_3, p_4, respectively.

Table 1: Estimated IACT at 'blue' pixel τ_{blue} with various swap move probabilities.

(p_1, p_2, p_3, p_4)	$\tau_{\text{blue}}(\times 10 \times m)$
(0.1,0.1,0.7,0.1)	9.5 ± 2.1
(0.1,0.1,0.1,0.7)	8.3 ± 1.8
(0.1,0.7,0.1,0.1)	7.5 ± 1.5
(0.1,0.1,0.4,0.4)	5.6 ± 1.0
(0.1,0.4,0.4,0.1)	6.5 ± 1.2
(0.1,0.4,0.1,0.4)	11 ± 2.6
(0.1,0.3,0.3,0.3)	7.3 ± 1.5
(0.1,0.9,0.0,0.0)	20 ± 6



Figure 6: Estimated normalized ACF at x_{blue} using swap move and single-site Metropolis.

MCMC sampler, also listed in Table 1.

We were not able to accurately estimate IACT at the 'blue' pixel using the Markov chain with just the single-site update, even after $1.6 \times 10^5 \times 576$ iterations. Instead we show in Figure 6 the normalized ACF at at the 'blue' pixel using singl-site update and the swap move (as in Figure 5). As can be seen, the normalized ACF decays to zero about 10 time faster for the swap move algorithm compared to the single-site algorithm. This confirms the qualitative assessment, and shows that the swap move is about ten times more efficient than the single-site Metropolis.

Acknowledgements

EM thanks the NZIMA for its generous support of his PhD study.

6. REFERENCES

[1] C. Fox, *Conductance Imaging.* PhD thesis, University of Cambridge, 1989.

[2] J. D. Moulton, C. Fox, and D. Svyatskiy, "Multilevel approximations in sample-based inversion from the Dirichlet-to-Neumann map," *Journal of Physics: Conference Series*, vol. 124, no. 1, 2008.

[3] D. Higdon, C. S. Reese, J. D. Moulton, J. A. Vrugt, and C. Fox, "Posterior exploration for computationally intensive forward models," in *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC, 2011.

[4] C. Fox and G. K. Nicholls, "Exact map states and expectations from perfect sampling: Greig, Porteous and Seheult revisited," in *Bayeisan Inference and Maximum Entropy Methods in Science and Engineering, 20th International Workshop*, pp. 252–263, 2001.

[5] G. K. Nicholls and C. Fox, "Prior modeling and posterior sampling in impedance imaging," in *Bayesian Inference for Inverse Problems*, 1998.

[6] J. A. Christen and C. Fox, "MCMC using an approximation," *Journal of Computational and Graphical Statistics*, vol. 14, no. 4, pp. 795–810, 2005.

[7] U. Wolff, "Monte Carlo errors with less errors," *Computer Physics Communications*, vol. 156, pp. 143–153, 2004.

MechBass: A Systems Overview of a New Four-Stringed Robotic Bass Guitar

J. McVay, D.A. Carnegie
School of Engineering
Victoria University
Wellington, New Zealand
jameshmcvay@me.com
dale.carnegie@vuw.ac.nz

J.W. Murphy, A. Kapur
New Zealand School of Music
Victoria University of Wellington
jim.w.murphy@gmail.com
ajay@karmetik.com

***Abstract*—The Faculty of Engineering at Victoria University of Wellington, in Collaboration with the New Zealand School of Music, has constructed a four-stringed modular robotic bass guitar, MechBass. This paper presents a systems overview of the instrument, focusing on the fretting mechanism, picking mechanism, actuator control electronics, and control software. Upon evaluation and testing, MechBass has been deemed a successful project, having met its goals of speed, accuracy, and repeatability.**

Keywords- Bass Guitar, Musical Robotics, Stringed Instruments

I. Introduction

Historically, robotic stringed instruments have implemented relatively few degrees of freedom. Their often-simplistic string plucking and fret positioning systems result in a lack of musical expressivity. The Faculty of Engineering at Victoria University of Wellington seeks to address these limitations by building MechBass, a multi-stringed modular robotic bass guitar. Ultimately, MechBass will be used in collaboration with the New Zealand School of Music as a platform for research into robotic musical composition and performance techniques.

MechBass can be divided into several subsystems. These subsystems are the plucking system, the fretting system, and the damping system. The subsystems are controlled by an actuator control board which communicates with a PC via the MIDI protocol. A major goal in the design of MechBass was to endow each subassembly with, to the authors' best knowledge, more degrees of freedom than equivalent subassemblies in prior robotic guitar and bass-playing systems.

The aim of this paper is to present a systems-level overview of MechBass. Following a brief history of musical robotics in general and robotic guitar and bass systems in particular, each of the aforementioned subsystems is described. Where applicable, evaluation information is presented.

II. Background

A. Automatic Music: A Brief History

Inventers and composers have long been interested by automatic musical instruments. Throughout the middle ages and Renaissance, automatic musical instruments were handcrafted devices: [1] presents a detailed history of such instruments, many of which stored musical information as pins or nails on rotating drums. The Industrial Revolution saw the Pianola and other such instruments mass produced on a wide scale: for a time they were the sole means by which music could be reproduced.

The 1970's saw resurgence in interest in automatic music. Pioneers Trimpin [2] and Godfried-Willem Raes [3] cite the kinetic aspect of their musical sculptures as a primary motivating factor in their work. Both Trimpin and Raes are among the groundbreaking workers to employ computer-aided musical composition techniques to drive actuators on musical robotics.

With increasing computing power in the 1990's and 2000's, musical robots began to be provided by their creators with artificial intelligence techniques. N.A. Baginsky's "The Three Sirens" [4] is an early example of an AI-driven musical ensemble. Gil Weinberg [5] and Eric Singer [6] have further explored combining mechatronic musical assemblies with AI-driven performance software.

B. Robotic Guitars: Motivation for Work



Figure 1 - Trimpin's Jackbox fixed fret mechatronic guitar

Several workers have created robotic guitars and basses. These fall into two categories: those employing a fixed fretboard and those equipped with a sliding, variable-position fretboard. Trimpin's Jackbox, shown in Figure 1, and EMMI's

PAM[1] are prominent examples of the former, while Singer's GuitarBot is one of the best examples of a sliding fretboard robotic guitar. Both systems have drawbacks: while fixed fretboard systems allow for rapid transitions between notes, they do not allow for small sub-chromatic changes in string pitch. Sliding fretboard systems can play subchromatic notes but have not to date been able to lift away from the string; every note transition becomes a sliding portamento.

In addition to their limited fret positioning mechanisms, most robotic guitar and bass systems utilise one of two general string plucking methods: solenoid-based pluckers and rotary pluckers. Neither approach allows for "dynamic" (or loudness) control. Methods to improve upon these two approaches are discussed in more detail in the authors' prior work [7].

MechBass attempts to address these issues, employing a variable-volume plucking mechanism and a sliding fretboard capable of being clamped to the string or removed from it.

III. MECHBASS SYSTEM

The MechBass is composed of four single string units. Each single string unit is electronically independent from the others. Each of the four units has a string plucking mechanism, a string fretting [2]mechanism, and a bass guitar string. Each of the strings are tuned to the notes of G, D, A, and E, to correspond with standard tuning of a bass guitar.

A. Single String Units

i. Structural Framework

The chassis of the MechBass is built from T-slot aluminium extrusion. Prior to physical assembly, the chassis was designed in 3D CAD software, as shown in Figure 2.



Figure 2 - A CAD rendering of MechBass

Each of the subassemblies described below are attached to the chassis via laser cut 6mm Perspex.

ii. Fretting Mechanism

The fretting mechanism consists of linear motion solenoids attached to a carriage which is positioned along the bass string through the use of a belt drive. Commercial solutions for a linear motion system were researched and dismissed due to their high costs of approximately $2000 NZ. A lower-cost solution was developed: this solution utilises a NEMA 23 stepper motor with an attached timing belt. The timing belt is attached to the solenoid carriage; the solenoid carriage, shown in Figure 3, rides along the T-slot aluminium extrusion.



Figure 3 - The solenoid carriage fretting mechanism

The NEMA 23 stepper motor and idler pulley are attached to laser cut 6 mm acrylic brackets which are in turn attached to the T-slot aluminium frame. The idler pulley is connected to the bracket with a central shaft and is fixed in position along the shaft with circlips.

The solenoids on the carriage are connected to the system's 24 V DC power supply. Cable guides are employed to prevent tangles and wire-related obstructions to the linear motion system.

The solenoid carriage can be moved along the length of the string, positioning the fretting mechanism at different pitches. Position information is stored in the form of a lookup table. The lookup table contains step number values: upon receipt of a MIDI NoteOn event, the note's pitch is converted to a corresponding step number away from the home position. The step numbers correspond to pitches playable on a normal bass guitar and are determined by the equation $x = \frac{L}{K^n}$, where L is the string length (822 mm in the case of MechBass), x is the fret number, and $k = \sqrt[12]{2}$, the ratio of the spacing of two consecutive frets. Thus, the distance in millimetres for each fret n from the home position is given by $x = \frac{822}{1.05946^n}$ and, in turn, the distance from the home position to move to a desired fret is found to be $822 - x$. The last step in creating the pitch-to-steps lookup table is to convert these distances from the home position to be in steps instead of millimetres. Given that the stepper motor driving a pitch shifter takes 200 steps to rotate 1 revolution, and the attached pulleys driving the belt of the pitch shifter have a diameter of 50 mm, the linear displacement of the pitch shifter for each step of the stepper motor is found to be $\frac{50*\pi}{200}$.

iii. String Plucking Mechanism

A string plucking mechanism similar to that developed by R. Vindriis in [7] is implemented on each of the MechBass string units. The string plucking system consists of a NEMA 17 stepper with a pickwheel attached to the motor's shaft. The 3D

[1] www.expressivemachines.org

[2] Fretting refers to the act of positioning a bridge at a point on the string, changing the string's pitch.

printed pickwheel holds five bass guitar picks. The guitar picks are clamped to the pickwheel with laser cut clamps, as illustrated in Figure 4.

To adjust the loudness of each string pluck, the plucking mechanism is mounted on a servo-driven pivot. The servo's arm moves the motor around the pivot point, adjusting the relative height at which each pick can strike the string.



Figure 4 - String plucking mechanism: a pickwheel is attached to a stepper motor which rests on a servo-driven pivot.

To pick the string, the stepper motor must simply be rotated the correct number of steps. As the motor rotates 1.8 per step, then it takes $\frac{360}{8} = 200$ steps per entire revolution. With five picks on the pick wheel, the motor must turn $\frac{200}{5} = 40$ steps to pick the string once.

iv. Damping Mechanism

To supress string vibrations, a damper system was added to the MechBass. The design goal of the damping system is to fully damp the string within one second of a damping instruction's transmission. The damper consists of an RC Servo with a felt-padded arm. When a MIDI NoteOff event is received by the MechBass unit, the servo's padded arm is moved in contact with the string, dampening the string's vibrations. While many existing robotic guitars and basses use either no damper or a simple on/off solenoid-based damper [6], the system implemented on the MechBass can vary the degree of damping by applying more or less pressure to the string with the RC servo.

Figure 5 - The damper mechanism on MechBass

The damper mechanism is shown in Figure 5, illustrating the bracket-mounted servo with laser cut acrylic arm and felt pad.

v. Optical Pickup

A typical bass guitar uses a magnetic pickup to capture the vibration of the string. However, due to the stepper motors and other electronics that are involved in this project, a traditional magnetic pickup cannot be used due to the amount of electromagnetic noise these devices generate. Instead, an optical pickup is employed to capture the vibration of the string for amplification.



Figure 6 - Optical pickup, with emitter LED visible above string and phototransistor circuit mounted on adjustable springs below the string.

The optical pickup, shown in Figure 6, consists of two plates: the top plate houses an infrared LED. The bottom plate houses a phototransistor. The vibrations of the string, mounted between the two plates, vary the amount of infrared light reaching the phototransistor. The amount of light varies in correspondence with the string's vibrations.

Upon initial evaluations, the pickup was found to be very susceptible to power supply noise. To address this sensitivity, it was powered separately with a 3.3V supply.

vi. Actuator Control Electronics

To control the actuators on each of the single string units of MechBass, a custom microcontroller-based system was designed and built. The board, based on the Arduino and dubbed the JM2, receives MIDI messages. The JM2, shown in Figure 7, employs an ATMEGA328 AVR microcontroller and responds to MIDI messages. MIDI, a serial protocol popular with musical instruments, is broadcast from a master device to a bus. The MIDI message contains channel data, pitch data, and velocity (loudness) data. Each JM2 board listens on the bus to messages that correspond to its specified channel.



Figure 7 - The JM2 actuator management board

Upon receipt of a message, the JM2 board interfaces with Allegro A4988 stepper motor drivers, driving the fretting mechanism to a position corresponding with the received MIDI note's pitch. Subsequently, the velocity servo is raised to a position corresponding to the note's velocity and a second A4988 stepper motor driver moves the pickwheel, resulting in a string pluck at desired volume level.

To help compensate for slippage, the fretting mechanism is returned to its home position on system start-up, resulting in the tripping of a limit switch connected to the AVR microcontroller's interrupt line and resetting its home position.

The JM2 board is divided into three main subsections: The communications subsection, the microcontroller subsection, and the motor driver subsection. The communications subsection contains the necessary support circuitry [3] to implement a MIDI input and MIDI through port. It interfaces with the microcontroller subsection through the AVR microcontroller's USART.

B. Four-string System

Four single string units of MechBass are combined to form a four-stringed assembly. This four-stringed system is tuned in the same manner as a bass guitar. While each single string unit has its own electronics and mechanical systems, all four strings share a single power supply module consisting of a 24 V supply for the stepper motors, a 24 V supply for the solenoids, and a 5 V supply for the electronics and servos. The power supplies are attached to the aluminium frame in a laser cut enclosure, shown in Figure 8.



Figure 8 - The power supply module, attached to the aluminium frame of MechBass

[3] http://www.midi.org/techspecs/index.php

To allow for all four separate single string units to communicate with a MIDI host device, a MIDI bus is created. The bus is created by connecting the output from the MIDI host device to the MIDI In port of any of the JM2 boards. The subsequent board on the bus is connected to the previous JM2 board's MIDI through port.

Each MIDI message is then addressed to a single channel, with the single string units not on that channel ignoring the message. This technique allows for easy setup and use with existing MIDI software: popular music composition and performance tools such as Ableton Live and Apple Logic can be easily configured to function appropriately.

The full four-string system is attached to the t-slot aluminium frame. Also attached to the frame is the power supply and a Firewire audio interface for multitrack recording of the output of individual strings. The fully-assembled four string unit is shown in Figure 9.



Figure 9 - The fully-assembled MechBass system with four string modules.

IV. Evaluation

Upon completion of design and assembly, the subsystems of the MechBass were evaluated. Five musically important parameters were selected for evaluation: fretting speed, fretting accuracy, plucking velocity with varying input parameters, damper effectiveness, and optical pickup performance. In addition to profiling the performance of MechBass, the outcome of these evaluations will prove useful in comparing the system's performance to future mechatronic systems as well as human players and physical modelling synthesizers.

A. Fretting Speed and Accuracy

To observe the strings' pitch error at different frets, each single string unit is instructed to move to each of its preset fret positions. The pitch at each position is then extracted, and the error (in cents, or hundredths of a musical semitone) calculated and shown in Figure 10. The maximum observed error is 5.2 cents, lower than the six cents perceivable by humans as described in [8].



Figure 10 - Each string's error at different frets in cents

The pitch shifters' speed is evaluated by recording the time taken to shift from each fret to every other fret. The results are shown in Figure 11, which illustrates time taken to move to each other fret from a given starting fret. The shifter is shown to achieve top speed in long shifts, as it has time to accelerate to its maximum speed. The speed was found to be significantly higher than in previous prototype versions of the bass [7].



Figure 11 - Time taken to shift from each fret from different starting frets

B. *Plucking Dynamics and Consistency*

To determine how the signal's velocity varies with changes in the height of the velocity control servo, the signal's power at 20 different servo heights is tested. The results of this evaluation, shown in Figure 12, demonstrate a clear trend in increasing signal power as servo height increases. It is clear from Figure 12 that the goal of controllable velocity has been met; further linearisation of the velocity curve could be accomplished with a lookup table or a linearisation function.



Figure 12 - Effects of velocity control on RMS of signal amplitude

To determine how consistent the picks are at a given specified velocity, the power of ten picks (two complete pickwheel revolutions) is recorded. The results are shown in Figure 13, with deviations from the mean likely resulting from inconsistent pick placement on the pickwheel. These deviations could be addressed by refinements in pick mounting techniques.



Figure 13 - Signal RMS for two complete revolutions of the pickwheel

C. *Damper Effectiveness*

To evaluate the damper system, the string is plucked for one second followed by a NoteOff event, instructing the damper to activate at full intensity. The resultant waveform is evaluated. From the waveform, shown in Figure 14, it is clear that the damper first makes contact with the string approximately 50 milliseconds after the NoteOff command is sent. Significant damping occurs after 300 milliseconds, with full damping occurring after 750 milliseconds. The design goal was to create a damping system capable of full string damping within 1 second of the instruction being sent. As shown, the goal is met. Improvements in damping could be made by further padding the felt damping arm.



Figure 14 - Waveform demonstrating the effects of a damper after string pluck

D. *Optical Pickup Output*

The output of the optical pickup is tested by plucking the string at a frequency of 103.83 Hz and analysing the resultant waveform generated by the pickup. The note, G#, is picked and the output is recorded for one second. Figure 15 shows the output which, as expected, is clearly periodic. No obvious signs of clipping or excessive noise are present.



Figure 15 - The optical pickup's output waveform

To further analyse the waveform, the signal's frequency spectrum is observed as shown in Figure 16. The note's harmonic series is clearly evident with minimal noise evident. While further analysis of the pickup's output could be

conducted with the aid of audio feature extraction, these preliminary tests accompanied by listening tests indicate that the pickup's behaviour is acceptable.



Figure 16 - Frequency spectrum of G# note from optical pickup

V. Conclusions And Future Work

MechBass is the first four-stringed mechatronic bass playing system featuring positionable fretting mechanisms, variable dynamic range, and variable-intensity damping mechanisms. Upon building and evaluating the subsystems, it is evident that MechBass will be applicable in research toward machine musicianship and musician-robot interaction. Further research will focus on system optimisation to minimise latency and automatic adjustments for detuning in the strings.

A companion robot to MechBass, dubbed Swivel, is also in development. Future work will involve comparisons between different string fretting techniques on Swivel and MechBass. Such comparisons will allow for further refinement in future robotic string players.

The original project goal of designing, building, and testing a four-stringed robotic bass system was accomplished. Based upon the evaluation undertaken, MechBass performs sufficiently well to be used as both a research and musical tool.

References

[1] A. W. J. G. Ord-Hume, “Cogs and crotchets: A view of mechanical music,”*Early Music*, vol. 11, pp. 167–171, April 1983.

[2] S. Leitman, “Trimpin: An interview,” *Computer Music Journal*, vol. 35, no. 4, pp. 12–27, 2011.

[3] T. R. Laura Maes, Godfried-Willem Raes, “The man and machine robot orchestra at logos,” *Computer Music Journal*, vol. 35, no. 4, pp. 28–48, 2011.

[4] N.A. Baginsky “Nicolas Anatol Baginsky” Website. Retrieved 16 October 2012.

[5] G. Hoffman and G. Weinberg, *Musical Robots and Interactive Multimodal Systems*, ch. 14, pp. 233–251. No. 74 in Springer Tracts in Advanced Robotics, Springer, 2011.

[6] E. Singer, J. Fedderson, and D. Bianciardi, “Lemur guitarbot: Midi robotic string instrument,” in *Proceedings of the 2003 International Conference on New Interfaces for Musical Expression*, (Montreal, Canada), 2003.

[7] R.G. Vindriis, D.A. Carnegie, and A. Kapur, “A Comparison of Pick-Based Strategies for Robotic Bass Playing,” in *Proceedings of the 2011 ENZCon*, (Palmerston North, NZ), 2011.

[8] D. Loeffler, “Instrument timbres and pitch estimation in polyphonic music,” Master’s thesis, Department of Electrical and Computer Engineering, Georgia Tech, 2006.

Adaptive Noise Cancellation for the Electroencephalogram

Anusha Kandasamy[1], Richard D. Jones[2], and Stephen J. Weddell[1]

[1]Department of Electrical and Computer Engineering, University of Canterbury
Christchurch, New Zealand
[2]New Zealand Brain Research Institute
Christchurch, New Zealand
steve.weddell@canterbury.ac.nz

Abstract: Electroencephalography is a technique for recording the brain's electrical activity – the electroencephalogram or EEG. It is an important procedure in the diagnosis of several brain disorders, as well being a valuable physiological tool for studies of normal brain function. However, the EEG is often contaminated by numerous artefacts such as eye-blinks, muscle activities, and eye-movements. This paper describes the *Emotiv EEG Neuroheadset* to record EEG time-series ensembles and correct the inevitable ocular artefact, the eye-blink. In order to classify and remove eye-blink artefacts from the EEG, a multilayer perceptron, feed-forward artificial neural network was used. In MATLAB, an multilayer perceptron network was trained to classify and remove the eye-blink artefact. Preliminary results show that eye-blink artefacts can be substantially reduced in the EEG when trained and used on a single subject. The generalised case to eye-blink removal is also addressed.

Keywords: EEG, Artefacts, Eye-blinks, Artificial Neural Network, Multilayer Perceptron.

1 INTRODUCTION

Electroencephalography is a commonly used technique to measure, record, and monitor electrical activities of the brain by attaching electrodes to the scalp. The purpose of an electroencephalogram (EEG) is to diagnose epilepsy and other brain function disorders such as dementia and brain death. It is also used to conduct studies on sleep disorders and microsleeps [1]. Artefacts such as eye movements, eye-blinks, electrode movements, etc., can dominate, distort, and mask individual EEG signal response. This condition makes it difficult to detect brain function disorders [2].

An article by James *et al.* in 1997 discussed an artificial neural network (ANN) developed for the enhancement of epileptogenic spikes [3]. The article illustrated that the non-linear multilayer perceptron (MLP), in the form of an adaptive filter can be used to improve the *noise-to-signal ratio* (SNR) of non-stationary signals. This helps to improve the detection of specific disorders, in contrast to multi-reference adaptive noise cancelling (MRANC), which is a linear implementation [3].

The aim of this research was to implement and further explore James' technique for classification and removal of the eye-blink artefact from EEG time series ensembles. This was accomplished using a multilayer perceptron and this serves as a basis for further development.

1.1 Artefacts

Artefacts cause distortion in brain wave recordings and in terms of signal processing can be considered as noise. An artefact is divided into physiologic and extraphysiologic artefacts. Extraphysiologic artefacts arise from factors external to the body, such as equipment and environment, i.e., alternating current (50 or 60 Hz), electrodes, whereas the internal activities of the body are the reasons for physiologic artefacts, i.e., eye-blinks, eye-movements, and ECG artefacts [4].

Eye-blink is one of the most frequently occurring ocular artefacts in the EEG signal. The eye related signals such as eyes-open, eyes-close, looking left/right or an eye-blink will be dominate in the frontal and prefrontal regions of the brain [5]. Eye-blink also has a unique feature, where the peak amplitude of the eye-blink is significantly higher compared to the other components in EEG. The eye-blink can be distinguished by its positive and negative peak occurrences [5].

1.2 Neural Networks

The artefacts affecting the EEG signal need to be classified and removed during the preprocessing stage of the EEG recording, which is typically done manually by a human expert. Ideally, this can be semi- or fully-automated. Several methods have been proposed to remove artefacts from EEG recordings [6]. However, as outlined, ANNs have been used for signal enhancement [3].

ANNs are an information-processing paradigm. Inspired by the way biological nervous systems such as the brain processes information, it uses a large number of neurons, which are interconnected to learn to recognise patterns or classify data [7]. Artificial Neural Networks have been used in diverse applications such as image and data compression (telecommunication), real estate appraisal (financial), autopilot enhancement (aerospace), and EEG and ECG analysis (medical); all of which has provided excellent outcomes. Training an ANN depends upon a particular network-learning algorithm and the structure of the topology that supports the network.

Based on James *et al.* [3], it was deemed that the feed-forward multilayer perceptron would be the simplest approach to cancel common ocular artefacts.

1.2.1 Multilayer Perceptron

A multilayer perceptron network is a feed-forward artificial neural network model that maps sets of input data into a set comprising one or more outputs. The network consists of three main layers: an input layer, a hidden layer and an output layer. A model of the multilayer perceptron used in this research is shown in Figure 1.



Figure 1: Multilayer Perceptron Neural Network Model [8].

The values of vector **x** form the input layer and are distributed to each node in the hidden layer. The values are then multiplied by weights (w_{ji}) and bias constant, which are then summed together (u_j). This sum is then fed into an activation function σ, and the resultant value is h_j. In simple terms, an activation function scales the values to within a corresponding range. Numerous activation functions are available for ANN implementation. Three most commonly used are hard-limiting, linear, and log-sigmoid. Each activation function complies with specific requirements depending on the problem the neural network is solving. Lastly, the output of the hidden layer neuron is multiplied by a weight (w_{kj}), and combined with other hidden node weights to give v_k. This is then fed again into another activation function σ, which outputs $y_{1..m}$, forming output vector **y** [8].

1.3 Emotiv EEG Neuroheadset

The Emotiv EEG Neuroheadset is a recent addition in neuro-technology. This wireless neuroheadset is a high-resolution innovative interface for man-machine interaction [9]. The electrodes on the headset observe the electrical activity of the brain to sense a user's expression, thoughts, and feelings. As shown in Figure 2, the neuroheadset has 14 saline electrodes, compared to the usual 19 electrodes on a standard medical EEG, and supports a sampling rate of 128Hz. It also has a built in low-pass filter (5th order Sinc filter) and a bandwidth between 50Hz – 60Hz.

The Emotiv Neuroheadset was used to acquire EEG data and the resulting data ensembles were stored to a file. Matlab was used as an analysis tool, supported by the Artificial Neural Network toolbox.



Figure 2: The 14 saline electrodes and their positions [13].

2 DISCUSSION

The implementation of the feed-forward multilayer perceptron in MATLAB has four primary steps:

1. Collect data.
2. Create a network.
3. Train the network.
4. Verify and test the network.

2.1 Collecting Data

The first step included collecting EEG data using the neuroheadset and successfully plotting data using the MATLAB plot function. Collection of data involved obtaining the data and converting to *csv* format using the inbuilt converter. The data file was then loaded into Matlab, which converts the file to *mat* format. Matlab code was used to plot a time-series ensemble of 14 corrupted signals obtained from the headset.

The collected data were preprocessed by normalising and removing the DC offset. This will ensure data are in an appropriate form for the neural network to learn patterns within each input ensemble.

2.2 Network Initialisation

The network did not require initialisation and the default settings were used, i.e., the Levenberg-Marquardt training algorithm (*trainlm*), the log-sigmoid activation function for the hidden layer, and linear activation function for the output layer.

2.3 Training the Network

Figure 3 is a schematic of the training phase and details inputs and the training signal that were used in this process. The *input* consists up to seven signals, while the eighth signal (the signal of interest) is defined as Channel X; this is also referred to as the *target* signal. When the target signal (Channel X) and input signals are fed into the MLP, the network should produce an output ideally identical to the target signal artefact.



Figure 3: The system diagram of the network during training phase.

2.3.1 Input Signal

The input signal included channels that primarily showed dominant eye-blink artefact. The prefrontal region on the skull associated with electrode signals F7, F3 FC5, AF4, F4, F8, and FC6 is shown in Figure 2.

2.3.2 Target Signal (Channel X)

The target signal (AF3) used to train the network is a template of the eye-blink artefact and these were used for classification. The target signal was created by replacing a null wherever there was a non-artefact signal present and the eye-blink artefact was left unaltered. The signal traces replaced by nulls are represented by rectagular blocks in Figure 4.



Figure 4: The original signal (Channel AF3) and resultant target signal after modification.

Using the input signals, the network adjusts the weights in the hidden neurons accordingly to match the target signal. The MLP is trained by repeatedly submitting the inputs and related training signals to the network until the error between the target signal and the output is minimal.

2.3.3 Output

The upper plot in Figure 5 shows the original EEG data used to obtain the target data. The uncontaminated part of the signal was equated to null, thus ensuring the MLP will only classify the artefact. Comparing the target data and output signal, the artefact looked somewhat similar. However, when training the network a eye-blink artefact was inadvertently missed from the training set at 7 seconds. The network output predicted an eye-blink artefact, as shown by the lower plot in Figure 5, using the input signals as reference. This is an astonishing fact since it shows that the neural network is not only able to classify the eye-blink artefact but will also predict it.



Figure 5: Comparison between original EEG data (top), target signal (centre), and the output (bottom).

The network was trained again by using the same data from Figure 5, except the eye-blink artefact missed at seven seconds was included in the target data.

The resultant plot in Figure 6 shows that the output plot has the corresponding eye-blink artefact at 7 seconds.



Figure 6: The original data (top), target data with the missed eye-blink artefact (centre) and the output (bottom).

2.4 Testing the Network

After training, the network was verified and tested using: 1) a different input sequence from the same set of EEG data used to train the network, and 2) a completely new set of EEG data from a different subject. The testing procedure used for both tests is shown in Figure 7.



Figure 7: The system diagram in testing phase using trained MLP.

As stated in Section 2.3.1, signal input comprised seven EEG data channels, where an eighth EEG channel, Channel X, was the signal of interest that required artefact removal. When these input signals were fed into the trained MLP, eye-blink artefacts, characteristic to Channel X, were produced. The signal shown in the upper plot of Figure 8 was then subtracted from Channel X. The product of this operation was the output signal, which is equivalent to Channel X with attenuated eye-blink artefacts. This is shown in the lower plot of Figure 8.



Figure 8: Top: Predicted eye-blinks from the network. Bottom: The output from the system after the subtractor.

2.4.1 Testing with same EEG data

To evaluate network performance, the same EEG data initially used to train neural network was employed to test the network. This can be referred to as the non-generalised case.

Based on the procedure outlined in Section 2.3.2, if the network was trained and tested with the same data then the predicted eye-blinks artefacts should look identical to the eye-blink artefacts in Channel X. Tests showed these were highly correlated however some residual error was evident.



Figure 9: Comparison between Channel X (top), Predicted Eye-blinks (centre), and the Output (bottom) between 30 to 40 seconds.

The central plot in Figure 9 shows six individual eye-blink artefacts that were predicted by our MLP network. These closely matched the eye-blink artefacts in Channel X (Top). The decrease in the peaks can be found by taking the difference between the original peak and the diminished peak.

Table 1: MSE of five different non-contaminated signal portions in Figure 9.

Time Start (s)	Time End (s)	Error
11.47	12.63	0.0214
21.20	22.54	0.0082
30.62	31.98	0.0189
41.45	42.44	0.0076
50.53	51.59	0.0099
	Average	0.0132

Table 2: Diminished eye-blink artefact peak calculation; Figure 9.

Peak Time (s)	Performance (%)
30.44	101.29
32.09	97.38
33.60	95.83
35.18	102.04
36.66	100.87
38.28	101.48
Average	99.82

The performance of the network configuraton, based on the results shown in Figure 9, can be summerised using two tables. Table 1 shows the degree to which Channel X was affected by the network, were the average mean squared error (MSE) was calculated to be 0.0132. Table 2 showns that average eye-blink attenuation was 99.82%. Thus, the network performed as expected.

2.4.2 Testing with the same subject

The network was then tested with a different set of EEG data, but from the same subject. This illustrated how well the network was able to classify eye-blink artefacts that it was not trained for. The resultant plots are shown in Figure 10.



Figure 10:Comparison between Channel X (top), Predicted Eye-blinks (centre), and the Output (bottom) between 30 to 40 seconds for different set EEG data from the same subject.

From Figure 10 it can be seen that the amplitude of the eye-blink artefacts from Channel X has diminished. The amount the peaks have diminished by 111.9% is shown in Table .

Table 3: Diminished eye-blink artefact peaks calculation;Figure 10.

Peak Time (s)	Performance (%)
31.89	86.97
34.11	126.5
36.10	83.22
38.00	141.66
39.86	122.31
Average	111.9

Table 4: MSE of five different non-contaminated signal portions; Figure 10.

Time Start (s)	Time End (s)	Error
11.7	13.24	0.0192
22.06	24.02	0.0139
32.04	33.96	0.0295
40.13	42.06	0.0226
50.42	52.19	0.0112
	Average	0.0193

A comparison of the peaks shown in Figure 10 is given in Table 3 and the average was 111.9%. There results suggest that the predicted eye-blink artefacts were slightly over attenuated.

The MSE given in Table 4 was calculated using five randomly selected signal periods to evaluate the degree Channel X was modified by the MLP. The average was found to be 0.0193. Even though the MSE ideally should be zero, the network can be improved by increasing the number of nodes or even the number of hidden neurons.

2.4.3 Testing with a different subject

This test quantified the generalised performance of the network by using another subject's EEG data, i.e., another subject's EEG input and training set ensemble were used to train the network. This was done by preprocessing the newly obtained signal and using the *sim* function to simulate the trained network.

The resulting plots are shown in Figure 11. The central subplot shows that the network is able to accurately classify the occurance an eye-blink artefact.

Table 5 shows he degree of eye-blink artefact removed, in terms of the average performance of the output from Figure 9, and was calculated to be at 29.17%. Thus, the attenuation was only marginal compared to non-generalised result. In terms of how the network affected the non-artefact portion of Channel X, the average MSE shown in Table 6 was calculated to be 0.0218. This increase is in contrast to the result obtained from the same subject's EEG data given in Table 4. Evaluating both the average results and the MSEs, a generalised result was not obtained.



Figure 11: Channel X of another subject's EEG data (top), Predicted Eye-blinks from the simulated network (centre), the output of the system (bottom).

Table 5: Diminished eye-blink artefact peaks calculation; Figure 11.

Peak Time (s)	Performance (%)
30.23	32.85
32.91	23.53
34.86	38.20
37.25	22.11
Average	29.17

Table 6: MSE of five different non-contaminated signal portions: Figure 11.

Time Start (s)	Time End (s)	Error
5.461	9.195	0.0077
16.93	22.51	0.0175
24.50	28.42	0.0581
30.45	32.8	0.0113
34.99	37.17	0.0143
	Average	0.0218

However, we felt that a generalised result was possible if the network was trained with more than one subject's EEG data.

3 FURTHER WORK

Reservoir Computing (RC) is a new paradigm to train, validate, and test recurrent neural network (RNN). A generic RC network consists of *input* neurons, a *reservoir* and *readout* mechanism, which receives data

from some or all of the neurons in the reservoir [10]. The *reservoir* is a random dynamical system, which maps the inputs to a higher dimension. The readout mechanism is then trained and the resulting vector is mapped to the output. Readout weights are trained using either linear regression or ridge regression [11]. Two major theoretical models associated with RC include the Echo State Network and Liquid State Machines [12].

4 CONCLUSION

In conclusion, the aim of this research was to record brain signals using an *EmotivEEG Neuroheadset* with eye-blink artefacts, classify, and remove the eye blink artefacts from a channel of interest. The methodology used to implement this was a feed-forward artificial neural network, specifically, a multilayer perceptron. MATLAB was used to implement, train and simulate the feed-forward network to classify and effectively remove the eye-blink artefact using EEG signals from several subjects.

This process was successfully implemented and the results obtained were as expected. It was remarkable that the neural network was able to predict a missing eye-blink artefact from the training data set using the input signals as a reference. The results obtained from our test showed that the trained network worked well with only one subject's EEG data. However, when simulating the network on another subject's EEG data, the result were not as accurate. We conclude this was due to limited training data presented to the network from a variety of subjects and a wider study is planned.

5 REFERENCES

[1] R. S. Benbadis, "EEG Artifacts," 25 April 2012. [Online]. Available: http://emedicine.medscape.com/article/1140247-overview. [Accessed 30 April 2012].

[2] P. Sadovsky, D. Cermak and J. Rozman, "Removing Artifacts from EEG Signal," 2003. [Online]. Available: http://www.urel.feec.vutbr.cz/ra2007/archive/ra2003/papers/197.pdf. [Accessed 28 August 2012].

[3] C. J. James, M. T. Hagan, R. D. Jones, P. J. Bones and G. J. Carroll, "Multireference adaptive noise cancelling applied to the EEG," *IEEE Transactions on Biomedical Engineering,* vol. 44, pp. 775--779, 1997.

[4] C. J. James, "Detection of epileptiform activity in the electroencephalogram using artificial neural networks," Ph.D. Thesis, University of Canterbury. Electrical and Electronic Engineering, Christchurch, 1997.

[5] B. Chambayil, R. Singla and R. Jha, "EEG Eye Blink Classification Using Neural Network," *Proceedings of the World Congress on Engineering,,* vol. 1, June 2010.

[6] J. Bhat, "Signal Processing Techniques for Artifact Removal in Electroencephalogram (EEG)," Arlington, 2009. [Online]. Available: http://dspace.uta.edu/handle/10106/2047, [Accessed 28 August 2012].

[7] C. Stergiou and D. Siganos, "Neural Networks," *Surprise 96* Journal, vol. 4, no. 11, 1996.

[8] P. Sherrod, "Multilayer Perceptron Neural Networks," DTREG, [Online]. Available: http://www.dtreg.com/mlfn.htm. [Accessed 12 September 2012].

[9] S. Hughes, "Emotiv Emerges from Stealth to Revolutionize Human Computer Interaction," 2007.

[10] R. Chaudhuri, "Reservior Computing: A New Hope?," 3 Quarks Daily, January 10 2011. [Online]. Available: htp://www.3quarksdaily.com/3quarksdaily/2011/01/reservoir-computing-a-new-hope.html. [Accessed 14 September 2012].

[11] D. Verstraeten, "Introduction to Reservoir Computing and the RC toolbox," 24 September 2009. [Online]. Available: http://reslab.elis.ugent.be/sites/csl/files/RCT.pdf. [Accessed 10 September 2012].

[12] "Reservior Computing: Shaping Dynamics into Information," Reservior Computing, 30 March 2009. [Online]. Available: http://reservoir-computing.org/. [Accessed 14 September 2012].

[13] S. Manuel, "Sousa Rodrigues: Signal Processing and Evaluation," July 2011. [Online]. Available: http://sousa.net23.net/pages/signal_processing.html. [Accessed 11 September 2012].

Use of the Kinect Sensor in a Robotics Development Project

P. F. Nieuwoudt[1], R. Mackill[2] and R. Kumar[2]

[1]Engineering Centre of Excellence, Manukau Institute of Technology, Auckland, New Zealand.
[2]School of Electrical Engineering and Trades, Manukau Institute of Technology, Auckland, New Zealand.
Contact email: *pnieuwou@manukau.ac.nz*

Abstract: This paper deals with the implementation of a Kinect Xbox360 as a sensor on a fixed sorting robot in a development project. The system used C# software for interfacing and control and Matlab for object recognition, colour analysis and height detection. Kinect position data was converted into robotic x, y and z co-ordinates for use by a Trio Motion robotic controller. The paper outlines challenges that were faced, and the steps taken to achieve satisfactory performance. The project demonstrated the feasibility of using the Kinect as a sensor for an educational robotic system.

Keywords: Kinect, Xbox360, object recognition, colour detection, depth sensor.

1 INTRODUCTION

The purpose of this work was to develop an educational tool for the field of machine vision in a robotic application. The task of developing the vision system was set as a final-year project [1] for a two-person group of BEngTech students, who were given the task of investigating the feasibility of applying machine vision for the control of a pick-and-place robot.

To give the project a measure of realism, the robot was to be made capable of sensing the colours and positions of a number of different-coloured circular work-pieces, with the added possibility of being able to discriminate work-pieces of different heights. Hence, the sensing system was required to have full-colour capability together, if possible, with depth sensing features.

A recent consumer-electronic sensor to have appeared on the market is the Microsoft Kinect for Xbox 360. The Kinect was designed primarily for user interaction in a video game playing environment. However, the characteristics of the image and depth data captured by Kinect have attracted the attention of researchers from many other fields. The decision was made, therefore, to adopt the Kinect as the sensor for this project, and to investigate its feasibility for the application described above.

Two automated robotic sorting applications were developed; first, a research version using the Matlab Image Processing toolbox with its analytical facilities; and second, a C# version that more closely resembles an industrial prototype application. Both were interfaced to the same Kinect sensor and the same Trio Motion multi-axis robotic controller. This paper focuses on the Matlab application since this will form the basis of envisaged future student projects.

2 KINECT FEATURES

The Kinect is equipped with an RGB colour camera located at its centre, as shown in Figure 1. This provides an 8 bit VGA resolution (640 x 480 pixels) video stream at 30 frames per second. In this project the Kinect's RGB camera is used for colour detection of circular objects.

The monochrome depth sensing video stream is in VGA resolution (640x480 pixels) with 11 bit depth data, however Kinect's firmware reduces the resolution to 320 x 240 pixels [2]. The Kinect has a stated working range from 0.85m to 4m, however, the later Windows-version can sense closer ranges [3].

The depth sensor was used in this project for depth acquisition of circular pieces on a robotic work table, above which the Kinect sensor was mounted at a height of 1.3m. Taking into account the range of movement of the robotic vertical axis, this gave a minimum range of 0.919m below the sensor to the highest position of the robotic gripper. This was chosen so as to be greater than the minimum sensor range of 0.85m.



Figure 1. Kinect for Xbox 360.

3 DEPTH SENSING DETAILS

The Microsoft Kinect for Xbox 360 range measurement is based on a continuous projected infrared structured field in the form of a speckle pattern, where depth

mapping is performed on the captured image using triangulation. The pattern is captured by an infra red camera and is compared against a reference pattern in memory. The reference pattern is obtained by capturing a plane at a known distance from the sensor. When an object is closer or further than the reference plane, the speckle image will be shifted. This is caused by the position offset at Kinect's baseline between the infrared projector and the infrared camera and is related to the change in the object range. These shifts are measured for all speckles by a simple correlation procedure and deliver a disparity image. For each valid captured pixel, the distance to the sensor can then be calculated by the onboard processor. Invalid pixels, in shaded areas not illuminated by the projector, are returned as black with no range measurement.

The depth resolution refers to the minimum depth difference that can be measured and is determined by the number of bits per pixel used to store the disparity image. The Kinect disparity measurements are stored as 11-bit integers, where one bit is reserved to mark the pixels where no disparity is measured. Thus the disparity image contains 1024 levels of disparity.

The depth resolution falls off with the inverse square of the distance from the sensor; this starts off as 1.3mm at the minimum range of 0.85m, and deteriorates to 70mm at the maximum range of 5 metres. Random errors of depth measurement increase quadratically with increasing distance from the sensor and can reach 40mm at a range of 5 metres [4].

4 MC206 MOTION CONTROLLER

The Trio Motion MC206 is a robotic controller capable of handling up to 5 axes. In this project we only use 3 axes: X, Y and Z. Because the stepper motors used in this robot give no feedback, the gripper needs to be homed to a fixed datum position before beginning a run. This is done with the aid of datum switches which provide the datum coordinates to the controller. The open loop control has the disadvantage that any accidental mechanical displacement can result in all subsequent motion suffering an offset. An alternative would be to use servo motors which can feed back and control the position of the gripper, however, these were not adopted due to cost constraints.

5 INTERFACING

This was achieved through the use of the Kinect SDK (July 2011 beta stage 1 prototype version) [3,5] which provides sample programs to implement its features. Much of the sample code is given in C#, hence it was decided that Visual C# would be the language to use for the interface.

The design of the interface sets up the Kinect to give two output streams: the first being the RGB stream and the second the depth data stream. Refer to Figure 2 for the system block diagram. The two data streams are both in the form of image frames at 30 frames per second.



Figure 2. Block Diagram of System.

5.1 Extracting Depth

The depth data is in the form of an image containing 2 bytes per pixel. This depth data is calculated in millimeters from the disparity image and is placed in two bytes by the Kinect onboard processor.

Three requirements had to be met in order to access the required distance and object-identification data:

1) The depth frame (given in a single integer sequence) had to be ordered to provide a frame similar to the RBG video frame.

2) An image coordination process is necessary since the RGB frame is 640 x 480 pixels while the depth frame is only 320 x 240 pixels.

3) An offset correction is necessary. There is a slight offset due to the positioning of the depth camera and the RGB camera: the depth frame is offset by a slight angle and there appear to be shadows around some of the objects, producing undetected regions.

Equations were derived to match the RGB pixels to the correct depth pixels by loading and analyzing all the distances from a frame capture of the 320 x 240 depth stream in a Matlab array, then placing circles in selected places on the RGB frame and finding their corresponding centroids in the Matlab depth array. The frame disparities were corrected using linear approximations which linked the two frames at the operational surface depth. These were used for depth calculation on a selected RGB image from the RGB video stream.

6 MATLAB PROCESSING

Matlab code was developed, using Image Acquisition and Image Processing Toolboxes to acquire and process the required images. The RGB image data and distance data from Kinect are respectively stored in an image file named 'capture.bmp' and a text file named 'distances.txt'. When the application is run, Matlab reads the RGB and depth data and starts image processing.

The processing begins with the circle detection process, and then the detection of the colour of the circular pieces. Finally, combined with the depth data, all the processed information is manipulated into meaningful commands that the Trio Motion controller can interpret, to instruct the robot to operate according to its pick and place commands.

A suitable GUI was designed using Windows Presentation Foundation, which is a relatively new

graphical API released by Microsoft, and developed to improve the way an application handles 2D and 3D graphics. The GUI enables the user to choose between Manual and Auto modes and switch between Depth and RGB streams, as well as home the gripper. In Manual mode the user can run individual scripts, which helps in troubleshooting when problems may arise.

6.1 Interfacing Matlab with C#

This application uses the Matlab API created for the .NET framework [6]. In order to initialize a command window, it is declared as an object in the C# program. This object can then be used to send commands to the Matlab command window.

The buttons in the GUI are linked to commands that run script files which enable us to process our data. Each pixel's distance data is saved to a text file in the Matlab script folder. This is done using the text writer built into the .Net framework, which can write data to a .txt format. A function in Matlab will read the text file and place all the distance values into a 320 x 240 array.

6.2 Image Processing and Circle Detection

The colour image file 'capture.bmp' is loaded into Matlab using the "imread" function and is stored as a variable in a 640 x 480 x 3 RGB data array. This image is then converted into grey-scale for circle detection using the Hough Transform. The image is then further reduced to select only the region of interest – a black area in the original image that approximates pixel rows 178 to 316, and pixel columns 205 to 380. An averaging filter is used to smooth the image to produce a clean local gradient in grey-scale intensity transitions.

The circular Hough Matlab function [7] is called in the command

```
[ac, cen, rad] = Circular Hough Grd (imgfltrd, [5 15],…
0.01, 5, 1)
```

which detects circular shapes in a grayscale image, 'imgfltrd', and resolves their centre positions and radii. Output variable 'ac' provides an accumulated array of the circular Hough Transform (in the dimensions of the image), 'cen' gives each circle's (x, y) centroid position in pixels and 'rad' gives the radius (pixel units).

It should be noted that the four input arguments of this function specify the circle to be detected. The specifier '[5 15]'sets the radius range; *i.e.* the process will search for circles with a radius in the range of 5 to 15 pixels. The next three inputs are for gradient threshold setting, radius filtering (finding the radius of an imperfect circle) and multi-radii tolerance (detecting the radius of concentric circles: the higher the tolerance value the more sensitive it will be). The parameters were adjusted to best suit the requirements in the field.

Running the circle detection routine produces the radii and centroid values (in pixels) for the circles. By plotting these values on the original RGB image and by calling a "DrawCircle" function, Matlab will plot the circle with its respective centroid, as indicated in Figure 3.



Figure 3. Image Displaying Detected Circles



Figure 4. Original RGB Image

Figure 4 shows the full field of view captured by the coloured camera.

6.3 Image Processing and Colour Detection

The image detection was performed on six distinctive colours (Red, Orange, Yellow, Green, Blue, Violet).

The system uses a 'for loop' for automatic detection of the six defined colours in sequence. The code for colour extraction is same for all six, the only changes being different settings for the colour thresholds and colour enhancing parameters.

As an example, the following code was used to detect RED circles:

```
if number == 1;  (For RED)
H = fspecial('disk',2);
rgbImageBlur = imfilter (rgbImageRaw,H,'replicate');
```

Here "fspecial" is a 2-D circular averaging filter. When used with "imfilter", the image is blurred to smooth out any noise on the image. This noise could be white, black or any other unwanted coloured spots. Blurring reduces the influence of the noise. For the detection of object boundary pixels, the image can be sharpened to increase

the prominence of the boundaries. This is achieved by using a Laplacian filter in Matlab's 'fspecial' function.

```
S = fspecial('unsharp');
rgbImageSharp=imfilter(rgbImageBlur,S,'replicate');
```

The function "imadjust" is used to enhance the intensity of the colour. This is achieved by changing the specific RBG upper and lower parameters in the function. The following code was used to enhance the red colour.

The parameters change according to the lighting system, as well as the original object colour and saturation. If applied properly, the colours can be enhanced to their purest forms, shown in Figure 5 below; even a dull looking red can be enhanced to a pure red with RGB intensity of (255,0,0), for example

```
rgbImage=imadjust(rgbImageSharp,[.4 .5 .5; .6 .9 .9],[]);
```



Figure 5. Enhanced Red Image

Using similar processes for each colour, it was possible to correctly identify work-pieces of all six colours.

The work table of this robot was exposed to variable lighting conditions which affected the operation of the system and sometimes led to errors in colour identification. Parameter changes were then required to meet the prevailing lighting conditions.

After extraction from the original, the six colour images are acquired. As an illustration, Figures 6 and 7 show the extracted red and green images.



Figure 6. Binary Images after processing: Red (left) and Green (right).

Figure 7. Red plane (left) and Green plane (right) binary images after conversion to black and white.

6.4 Importing Distance Data

The text file contains all the depth information of the image that was captured by C#. However, it is saved as a 76800 x 1 matrix. This had to be converted into a 320 x 240 matrix. By taking the first 240 values from the 76800 long column and inserting these values into the first column and recursively repeating into all 320 columns, a filled matrix of 240 rows and 320 columns is produced.

The depth values for each work-piece were obtained by providing the x and y coordinates of their centroids to the depth calculation procedure. This enabled the extraction of depth information (*i.e.* the required z coordinate for the specified object) from the 320 x 240 depth data matrix:

```
inputx = circen(k,2);
inputy = circen(k,1);
y =  int16((( 320 - ( (inputy-5.45) / 2 )) - 10 ) * 1.12);
x =  int16((((  inputx / 2 )) - 24 ) * 1.12);
z = ( matrix( x , y ) ) - 919;
```

After the identification of all the centroids, colours and depths of the objects in the image, the robotic pick-and-place routine can start.

6.5 Interfacing Matlab with the MC206

An ActiveX component available in Matlab is used to provide direct communication to the Trio Motion Controller via a USB link. An ActiveX component can be used in any Windows programming language that supports ActiveX (OCX) components, such as Matlab. Using the TrioPC ActiveX component the programmer can conveniently open a direct communications link to one or more Motion Coordinators.

7 CALIBRATION

Calibration of the robotic operational plane was obtained by measurement on the work table. Pieces were placed at various positions across the area of interest. Using Matlab's circular Hough functions the centres were obtained. These positions were then measured by robotic gripper at each piece to obtain the x and y centroid values. Using least squares fit algorithms, linear

approximations were derived to relate the work-table and image coordinates, expressed as follows in Matlab code:

```
xmm = (2.1324*(circen(k,1)))-362.06;

ymm = (2.1405*(circen(k,2)))-308.54;
```

8 PERFORMANCE ASSESSMENT AND FUTURE WORK

The Kinect was set up as a static viewing sensor for this system, and its position had to be kept stable relative to the work table. Any translational movement of the Kinect would invalidate the calibration equations for x and y (above). Similarly, any change in its height would affect the accuracy of the depth measurements.

Varying lighting conditions were found on occasion to cause errors in colour decoding for the various work-pieces, and resulted in colour selection errors. Similar problems were reported by Prodromou [8]. It was found that the problems could be minimized by installing:
1) Black background on the work space to counteract the effects of shadows which would otherwise affect the circle detection algorithm.
2) White surface around the periphery of the region of interest, to help control the auto-contrast on the Kinect's colour camera.

The system also faced depth resolution issues. The pieces used were only 10mm high and there was little room for error, given a resolution of 4.5mm at a sensor distance of 1.3m. This was remedied by rounding off the depth to specific predetermined values for each piece. The resolution values were similar to those predicted with the quadratic equations derived by Kourosh [4] and resolution measurements performed and published by Viager [9]. For larger objects and larger robotic grippers, or shorter ranges, resolution would be less of an issue.

The system demonstrated the feasibility of using a low cost Kinect sensor in an autonomous robotic sorting application to select objects at different positions, heights and colours.

Future work could improve the sensor resolution by considering a revised approach to achieving depth data.

The colour identification method used here was cumbersome and processor intensive. We intend instead to investigate the use of 3-dimensional colour systems which could be used for colour-difference representation. Examples are the CIELAB (L*a*b*) system and other similar systems currently under development in the CIE and elsewhere.

Better vector methods will be evaluated for the collinear registration of the two different RBG and depth images.

Alternative calibration procedures will be investigated as a more robust (and perhaps automated) method would be strongly preferred.

It is clear that variations in the ambient lighting were the cause of sub-optimal performance of the system. We intend to investigate the installation of an optimized task-specific lighting system for this robot, in order to avoid such issues in the future.

9 CONCLUSIONS

The system described above was successful in creating a viable prototype robotic vision system and demonstrated autonomous operation to scan and detect objects, sort them by colour, and find and select them using a gripper at different heights. The work shows that with a low cost prototype sensor, a well performing program can be made to pick and place objects based on colour and depth. However, the present prototype showed a restricted and variable resolution (from as little as 1.3mm at 0.85m to as much as 70mm at range 5m) and would limit the application to object sizes that would fit the particular resolution at a given operational depth from the sensor.

The development project exposed students to the full cycle of hardware and software integration. This included the design and integration of software between platforms and the required calibration to the robotic work space. Specific adaptations were implemented for resolution adjustment on the machine application.

This system will provide a sound basis for future robotics and vision projects.

10 ACKNOWLEDGMENTS

The project is supported by the RDTT Fund of the Manukau Institute of Technology. The practical results are based on the project work of Robert Mackill and Ralvin Kumar [1]. Thanks to Dr Andrew Chalmers of the Engineering Centre of Excellence at Manukau Institute of Technology for advice and assistance.

11 REFERENCES

[1] Kumar, R. & Mackill, R., (2011) Colour Processing and Depth Acquisition for Robotic Control. Manukau Institute of Technology Report for Development Projects 115.715.

[2] Shane McGlaun.(2010, December 20). Microsoft Working to Quadruple Kinect Resolution. Available online: http://www.dailytech.com/Microsoft+Working+to+Quadruple+Kinect+Resolution/article20436.htm.

[3] digitalerr0r.(2011, June 21). Kinect Fundamentals #3: Getting distance-data from the Sensor. Available online: http://digitalerr0r.wordpress.com/2011/06/21/kinect-fundamentals-3-getting-data-from-the-depth-sensor/

[4] Kourosh, K., & Elberink, S. (2012). Accuracy and Resolution of Kinect Depth Data for Indoor Mapping Applications. Sensors (14248220), 12(2), 1437-1454. doi:10.3390/s120201437.

[5] Microsoft Research.(2011, July 22). Programming Guide. Available online: http://research.microsoft.com/en-us/um/redmond/projects/kinectsdk/docs/ProgrammingGuide_KinectSDK.pdf

[6] MSDN.(2011). .NET Framework 4. Available online: http://msdn.microsoft.com/en-us/library/w0x726c2.aspx

[7] Tao Peng 2007. Detect circles with various radii in grayscale image via Hough Transform. Available online: http://www.mathworks.com/matlabcentral/fileexchange/9168-detect-circles-with-various-radii-in-grayscale-image-via-hough-transform (accessed September 2011)

[8] Prodromou, D., (2011). Robotas - development of a mobile robot for small goods transport. Annals Of DAAAM & Proceedings, 1479-1480.

[9] Mikkel Viager.(2011, March). Analysis of Kinect for mobile robots. Available online: http://www.scribd.com/doc/56470872/3/Light-Coding-with-the-Kinect

Real Time Monitoring of Cat Feeding Behaviour

Said Al Souti*, Donald Bailey*, and David Thomas^

* School of Engineering and Advanced Technology (SEAT)
^ Centre for Feline Nutrition (CFN)
Massey University
Private Bag 11 222, Palmerston North, New Zealand
al_sooti933@hotmail.com, D.G.Bailey@massey.ac.nz, and D.G.Thomas@massey.ac.nz

Abstract: Cat food manufacturers currently compete intensively against each other to gain and maintain market share. They are striving to produce a product that is palatable, healthy and preferred by the majority of the cats, by understanding more about cats' feeding behaviour. Therefore, in this paper, we propose an efficient and economically viable method to monitor cat feeding behaviour. The proposed method was based on a background subtraction algorithm to detect the cats if they were in the feeding area and trigger the video recording accordingly. Background subtraction was performed by calculating the absolute difference of the background and the new frame. Then the result was thresholded using Otsu's method to extract the cat's body from the background noise.

The proposed method was implemented and tested for one week in the Centre for Feline Nutrition, a total of 20 hours of monitoring, and it showed an extremely accurate performance, where a single i5 Processor was able to monitor and record videos of 4 cages at 25-30 fps simultaneously. In addition, the algorithm was able to provide a successful detection rate of 94%, with a 5% false acceptance rate and a 1% false rejection rate.

Keywords: Monitoring cats, running average background subtraction, template matching, cross correlation, load cell.

1 INTRODUCTION

Cats are New Zealand's most popular pet with 48% of households owning an average of two cats. For most of these cats, commercially prepared foods form the sole source of nutrition. Therefore, consumers expect high standards of quality, nutrition and value for their pets. To achieve this standard of quality, pet food manufacturers are continually working to increase their knowledge of pet nutrition, ensuring the food is palatable and healthy [1].

The Centre for Feline Nutrition (CFN) has the responsibility of assessing and quantifying the quality of different cat food products and ultimately deciding which formulation is preferred by the majority of a test panel of cats. Their current testing procedure is to place 8 cats in 8 cages for 2 hours a day for 5 days; with each cage having two bowls of different food products. The initial and the final weights of the food in each bowl are measured to determine the amount of diet consumed.

The system was recently upgraded to allow real time weighing of each bowl [2]; which introduced several improvements and benefits, such as:

1. identifying the number of meals each cat had;
2. identifying the periods between meals;
3. identifying the amount of food consumed in each meal;
4. identifying the relative rate of consumption of the 2 diets by each cat, which ultimately is a measure of the palatability of each diet.

In addition to these benefits, it was realised that there may be a possibility of extracting new information from the 2 hour testing sessions by monitoring the cats and studying their behaviours, while they consumed their meals. Cats have a complex feeding pattern; they consume small meals at intervals during a day, rather than finish the diet offered to them in one sitting. This type of behaviour makes it impractical for the analyser to directly observe their feeding behaviour because he/she would have to observe the cats for the entire session. In addition, a single person cannot track multiple cats simultaneously. Therefore, to facilitate these types of studies, we created a real time video-based monitoring of cat feeding behaviour system, which has the potential to:

1. Allow the analyser to determine how and where the felines have their meals.
2. Clarify some of the anomalies in the real time weighing system's results, such as the reasons behind the sudden drop or rise in a bowl weight, which cannot be determined without a visual proof; i.e. determine whether the cat may be playing with the food, spilling the

food, or removing the food from the bowls and consuming it somewhere else.

3. Enable the analyser on determining the role of olfaction (smell) in meal choice.

The proposed solution for monitoring cat feeding behaviour was to install a USB camera in front of each cage, which was held in place by a mechanical arm that provided the flexibility in adjusting the angle and the distance of view. Then a PC was used to run Lab View software that performed a background subtraction algorithm to each cage simultaneously at 25-30 frames per second to detect whether the cat was in close contact with the food or not. At the same time the program recorded the real time weight of each bowl and associated each weight reading with its corresponding frame while recording.

A separate Windows Presentation Foundation application was run on the client machine which presented the collected data to the client or the analyser with an elegant user interface; where the analyser has the ability to navigate through each recorded video by clicking on any value in the real time weight chart *(more details are given in sections 3 and 4).*

The challenging aspect of the monitoring system was the implementation of a motion detection algorithm that was computationally inexpensive and at the same time worked for cats of various colours *(apart from white cats, since the cages' background is white)* while keeping the number of false alarms to a minimum. The background subtraction algorithm was modified to account for slight changes in lighting, from cloudy to sunny, while not detecting small changes in background, such as a cat's tail swinging in the feeding area, small vibration of the camera and so on.

2 RELATED WORK

The feeding behaviour monitoring system consisted of an activity triggered video recording system and a customised media player to display the obtained results. During our research, it was found that there were several works and solutions already published using video-based tracking systems.

Chen et al. [3] developed a stand-alone video-based animal tracking system for noiseless application. Their approach for tracking rodents in an acoustic environment was to fix a coloured marker to each target object. Then a tracking algorithm was implemented, that used the Y (luminance), Cb (blue-difference chrominance) and Cr (red-difference chrominance) values of the coloured marker to track the object position and to allow multi-object tracking against a complex background. An FPGA is used to process the colour video signal and to convert the data into sets of coordinates representing the momentary positions of the targets. The information was then stored on an SD card using a microcontroller.

Using an FPGA module, in our project, to perform the detection and the video compression would dramatically reduce the strain on the PC allowing it to handle more cages at once. However, in our case the FPGA module must transfer the compressed video and weights rather than position coordinates, which then must be stored in a hard drive (saving onto a hard drive will limit the number of parallel video streams able to be stored at the same time, thus limiting the number of cages that can be handled using a single PC). Therefore, for the cat feeding behaviour monitoring system, it was realised that using a single PC to handle the detection and the video compression of 4 cameras simultaneously was more economically viable.

Yang et al. [4] tracked multiple workers on constructing sites using video cameras. They addressed the challenge of having multiple workers in an interactive workplace, by developing a tracking algorithm based upon machine learning methods. The algorithm required several sample templates of the tracking target and learnt a general model that could be applied to other targets with similar geometry.

Straw et al. [5] worked on multi-camera real-time 3D tracking of multiple flying animals. Their work was capable of tracking the position and body orientation of flies and birds. They tested their system by using 11 cameras to track 3 flies simultaneously at 60 frames per second in a 3D space using a gigabit network of 9 standard Intel Pentium 4 and Core 2 Duo computers.

The work presented by Yang et al. [4] and Straw et al. [5] have a limited application to our problem, because the degree of complexity employed is not necessary to detect and record the cats inside their cages. Furthermore, one of our objectives was to build an efficient algorithm that allows a single PC with a standard i5 processor to handle up to 4 cameras simultaneously.

3 DETAILED DESIGN

The major objective of the system for monitoring cat feeding behaviour was to record the cats only when they were in close contact with the food and to match the real time weight data with its corresponding frame number while recording. The algorithm had to be fast enough to enable a single i5 processor to handle 4 cameras simultaneously.

The first step as (shown in figure 1) was to recognise the connected USB cameras and to match each camera ID number with its associated cage number as described in section 3.1. Then a background image must be taken from each connected camera, with all of the bowls inside the cages, and all of the cage doors closed before the system was initialised. Then the system was started by the user, the cats were put inside their cages. This ensured that the initial meal was recorded.



Figure 1: Block diagram of the real time monitoring of cat feeding behaviour system

When the monitoring system was started, a background subtraction algorithm was applied to 4 cages simultaneously and the weights from the 16 bowls (8 cages and each cage had 2 bowls) were recorded.

3.1 Cage identification

The cage identification numbers were determined through template matching using the cross-correlation method. The RGB image was first converted to a grayscale image (Luminance, Y) using the following equation:

$$\mathrm{Y} = 0.2126\,\mathrm{R} + 0.7152\,\mathrm{G} + 0.0722\,\mathrm{B} \qquad (1)$$

Then the search region, where the template matching was performed, was reduced (shown in figure 2) to speed up the process. The template images, which were the number plates, were retrieved beforehand and were stored on the hard drive as a part of the system data.

The monitoring system performs template matching 8 times (once for each template image) on each connected camera. The cage number associated with the camera is the number of the template with the highest score. However, if the maximum template matching score is below 50%, this indicates that the selected camera is not connected to any cage.

Figure 2: Search region of template matching

Template matching used correlation which treated the template image as a mask and moved the mask over the image (or the search region) while computing the sum of products at each location. The cross-correlation of a mask $\omega(x,y)$, with an image $f(x,y)$ may be expressed in the form:

$$c(x,y) = \sum_{s}\sum_{t} \omega(x,y)\, f(x+s, y+t) \qquad (2)$$

Equation 1 is very sensitive to scale changes in f and ω; so instead a *normalized correlation coefficient* was used, where values vary from -1.0 to 1.0, as described in Gonzalez and Woods [6]. The location where the coefficient is at its maximum and above 0.5, was where the match occurred.

3.2 Motion Detection

Motion detection was performed using a background subtraction approach. Background subtraction was applied to the luminance of both the background image (BG) and the luminance of the new frame (F). Both the BG and F images were cropped to include only the food area which was the region of interest. In addition, the process speed was improved by reducing the resolution of BG and F images by 2 in both horizontal and vertical directions. Then the absolute difference of F and BG was calculated:

$$\mathrm{Diff} = |\mathrm{F} - \mathrm{BG}| \qquad (3)$$

The absolute difference in F and BG intensities was an excellent method to separate the cat body from the background, because in our case there was a sufficient difference between the background (white) and the foreground intensities *(since there are no white cats in the CFN).* This difference was very sensitive to noise

and slight changes in BG or F intensities, so most of the small changes in intensity *(camera arm vibrations, small changes in lightings, shadows and reflections)* were avoided by thresholding the absolute difference.

However, a fixed threshold value was not applicable for our application; therefore, it was decided to use optimum global thresholding using Otsu's method to account for cats of different colours, such as black, brown, bright red and orange etc. Otsu's method is optimum in the sense that it maximizes the between-class variance [6]. In other words, when F has a bright red cat or orange cat, the total mean value of the absolute difference intensities will be low, so the Otsu's method will result on a low threshold value to separate the cat from the background noise, and will be the opposite for the cats with dark colour. Therefore, using an optimum global thresholding approach is like having an adaptive threshold that adapts for cats of different breeds and hair colour.

The next step was to decide whether the cat was in close contact with the food or not. This presented several challenges: 1. Cats often sat on the top of the feeding station and their tail occasionally dropped into the feeding area region which had the potential to cause false frame acceptance. 2. Cats occasionally sat or slept behind the feeding area in the cage, which was in the camera field of view and therefore within the region of interest.

The first issue was addressed by applying morphological filters to the thresholded image. A morphological closing of 3×3 structure element (SE) with 2 iterations was applied, followed by a morphological opening of the same SE with 5 iterations. This combination of morphological filters minimised the area of the tail while preserving the majority of the cat's body.

The second problem was solved mechanically by titling the camera's angle of view (shown in figure 3). This prevented the camera from capturing the cat at distances greater than 40 cm away from the food, and at the same time, it enabled the analyser to observe the cat's tongue and mouth while it is eating (shown in figure 3b).



(a) Wrong Angle of View

(b) Correct Angle of View

Figure 3: Modifying the cameras' angle of view to capture the cat when it is in close contact with the food

The decision of whether the cat was in close contact with the food or not was based on measuring the area (the number of detected pixels) of every connected component within the binary image. If the maximum area was greater than a certain threshold value, the cat was recognised as being in the feeding area otherwise it was not. The threshold value used in our case was 2000 pixels.

The background subtraction algorithm was enhanced by dynamically updating the BG when the cat is not in the feeding area:

$$BG_i = \alpha BG_{i-1} + (1-\alpha)F_{i-1} \quad (4)$$

This had the benefit of updating the BG when there was a gradual change in lighting or when a new small object was added to the background, thus minimising the possibility of having a false acceptance or a false rejection.

3.3 *Camera Arm*

The camera arm (shown in figure 4) provided flexibility in establishing the correct distance and angle of view, thus making it suitable for different types of camera. In addition, the camera arm was long enough so the camera was out of reach of cat's claws.

This design had the utility for permanently fixing the rotation of every joint; the first joint had 3 holes where the arm could be bolted, while the angles of the second and third joints were fixed using a bracket (shown in figure 4). This ensured that the ROI and the angle of view (as discussed in section 3.2) remained consistent.



Figure 4: 2D and isometric drawing of a Camera Arm

3.4 *Real-Time Weighing of the food bowls*

The weight of each bowl was measured using a load cell, where it's output signal was first amplified using *Instrumentation Amplifier INA2126 IC*. Then the

amplified signal was fed to a 12 bit analogue to digital converter. The digital signal was then passed to a microcontroller, where the input voltage was related to its corresponding weight in grams. After conducting several experiments on the load cells, it was found that there was a linear relationship between the amplified voltage and the real life weight (grams), as shown in the following equation [2]:

$$\text{Weight (g)} = 502.40 \times \text{Voltage (V)} - 570.38 \qquad (5)$$

A test was conducted to estimate the errors of the load cells, when equation 5 was used. The errors when plotted in x-y scatter plot appeared to be random and did not follow any particular pattern. In addition, the largest error was 0.782g, which was within the acceptable range of this process.

The microcontroller communicated with the host PC using the USB port. The PC accessed the weight of every cell, by transmitting the cage number to the microcontroller, and waiting for the microcontroller to process and send back the real time weight. This approach allowed the PC to access the weight of each load cell independently.

3.5 Presentation of the collected data to the client



(a)



(b)

Figure 5: Screenshot of the client software

The data collected from the experiment were the recorded videos and the bowl weights. This information was presented to the client with an elegant user interface application (shown in figure 5) with the following features:

1. The main window provides a quick overview of the result of each cage. A taskbar shows a summary of the amount of food consumed from each bowl and a complete 2 hour chart of the weight.

2. The user can click on any circular point on the chart to display the recorded video of that cat at that time.

3. The video player was designed to display the frames at their corresponding frame rate using the time stamp. In addition, it has a unique feature where the user can view the cat's behaviour associated with the bowl's weight by clicking on the corresponding point on the chart (shown in figure 5b).

The software was designed using Windows Presentation Foundation tools. The user first needs to select the folder that he/she wants to analyse (folder names contain the date and the time of the results). Then the software will read the weight charts for every cage and attach the recording time and the frame number for each point that is within the recorded period. These points are identified to the user by having big round circles (shown in figure 5a).

Clicking on one of these points will open up another window (shown in figure 5b). This window is the media player, where the video of the selected point will be played; the user has the option to stop, pause, replay and use the cursor to navigate through the video stream. Moreover, the system will display a detailed weight chart underneath the taskbar. Each point in the chart has the frame number attached to it, so it is possible to navigate through the video by clicking on any point in the chart.

4 EXPERIMENTAL RESULTS

4.1 Aim

The aim of the experiment was to assess the performance of the system and to understand how it will help the analyser assess visual components related to the cats' preference of the tested diets.

4.2 Method

The system was installed at the Centre for Feline Nutrition and it was used to monitor 4 cats simultaneously for 4 hours a day (2 sessions, with each session 2 hours long), for 5 days. Then the results were analysed by the authors on a daily basis.

4.3 Results of System's Performance



Figure 6: Pie Chart that shows how successful the system was in detecting the cats while they consumed their meals.



Figure 7: Frame rates versus the percentage of the total frames that were recorded. The system was tested in Intel i5 Processor with 8 GB RAM and 7200 RPM Hard Drive.

4.4 Discussion

4.4.1 System performance

The performance of the motion detection algorithm was evaluated (shown in figure 6). False rejection (FR) was identified by recording a single frame every second and using these frames to search for any movement that the system failed to detect. It was found that the system had 1% FR. This occurred when a cat had a meal of less than 4 seconds. This movement was recorded, but the system deleted it afterwards, since it was programmed to discard all recordings that were less than 4 seconds duration. The majority of the false acceptance was when a cat jumped from the bench to the back of the cage more than once during a 10 second period (the system was modified to have 10 second buffer after the end of each detection).

The average recording frame rate dropped from 29.76 to 26.70 (shown in figure 7), when the 4 cameras were recording simultaneously. This was due to the limitation on how fast the system could write to the hard drive and on how much CPU resources each JPEG compression required. The size of each JPEG image could be reduced by building a customised compression that exploited the temporal redundancy between frames. However, a customised video compression would increase the computational complexity of the system and a single PC might not be able to accommodate 4 cameras at once.

4.4.2 Analyser's experience with the new system

During these 5 days the analyser was able to watch the cats while they consumed their meals and to observe their behaviours. Moreover, there were several occasions when the cats rejected certain foods because of their smell, information the Centre had been unable to access previously.

4.5 Conclusion

The system performed well with a successful detection rate of 94%. It dropped an average of 3 fps when 4 cameras were recording simultaneously, which was an acceptable frame loss. In conclusion, the system served its purpose by enabling the centre to monitor the cats.

5 CONCLUSION

Real time monitoring of cat feeding behaviour system enabled the CFN to access new information from the 2 hours testing session that was unattainable before. The system will assist the Centre in understanding some of the reasons behind the cat's rejections and acceptance of certain food products, thus providing more feedback to the pet food manufacturer in relation to their products.

6 REFERENCES

[1] New Zealand Pet Food Manufacturers Association (2012). http://www.petfoodnz.co.nz/ (retrieved on 27 April, 2012)

[2] M. Cho. (2009). *Automated Analysis of Feline Feeding Behaviour*. Engineering Honours Project Report. Palmerston North: Massey University.

[3] Y.-J. Chen, Y.-C. Li, K.-N. Huang, S.-L. Jen, and M.-S. Young. (2009). Stand-Alone Video-Based Animal Tracking System for Noiseless Application. *Instrumentation Science and Technology*, vol. 37, no. 3, pp. 366-378.

[4] J. Yang, O. Arif, P.A. Vela, J. Teizer, Z. Shi. (2010). Tracking Multiple Workers on Construction Sites Using Video Cameras. *Adv. Eng. Info*. vol. 24, no. 4, pp. 428-434.

[5] A.D. Straw, K. Branson, T.R. Neumann and M.H. Dickinson. (2011). Multi-camera Realtime 3D Tracking of Multiple Flying Animals. *Royal Society Interface*, vol. 8, no. 56, pp. 395-409.

[6] R.C. Gonzalez, and R.E. Woods. (2007). *Digital Image Processing*. (3rd, Ed.) Prentice Hall.

Index of Authors

Al Souti, S., 163
Andreae, P., 7
Andrews, M., 67, 91
Atmojo, U.D., 103

Bailey, D., 163
Bainbridge-Smith, A., 49, 115
Balsom, T., 97
Barr-Smith, A., 13
Berber, S.M., 61
Brown, P., 129
Browne, W.N., 7, 31

Came, C., 13
Carnegie, D.A., 7, 31, 145
Chang, B., 13

Dorrington, A.A., 73
Duncan, S., 115

Eason, D., 67, 91
Evina, A.A., 109

Fox, C., 141

Hayes, M.P., 25, 85, 115
Hechei, C., 13
Ho, M-J., 61
Holmes, T., 121

Ivoghlian, A., 43

Jameson, C., 13
Jones, R.D., 151
Jull, H., 97

Kandasamy, A., 151
Kapur, A., 145
Kelly, S.G.J., 73
Kumar, R., 157

Lawrence, H.J., 55
Lee, W.Y.L., 67, 91
Lim, P.P.K., 37
Liu, R.SH, 37

Ma, E., 141
Mackill, R., 157
Manuel, I.L., 49
McLaren, D., 13
McVay, J., 145
Molteno, T.C.A., 55, 79, 121, 129
Murphy, J.W., 145

Nieuwoudt, P.F., 157

Panckhurst, B., 129
Payne, L., 13
Pink, N., 85
Punchihewa, A., 135

Rowe, G.B., 19

Salcic, Z., 43, 103
Scott, J., 13, 97
Shaw, C.F., 79
Shepherd, M., 19
Smaill, C.R., 19

Thomas, D., 163
Tsui, C., 109

Wang, K.I-K., 43, 103
Watterson, C.A., 7, 31
Weddell, S.J., 151
Wong, C.K., 37

Yu, K., 61

Electronics Group
Department of Physics
University of Otago
elec.otago.ac.nz