

ENZCON 2009

CONFERENCE PROCEEDINGS

18TH – 20TH NOVEMBER



THE COLLEGE BUILDING.

UNIVERSITY OF OTAGO
DUNEDIN, NEW ZEALAND

Published by:

Electronics Research
Department of Physics
University of Otago

www.physics.otago.ac.nz/nx/articles/enzcon-2009-at-otago-university.html

ISBN: 978-0-473-16099-9

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

Programme Committee:

Tim Molteno (Otago)
Colin Fox (Otago)
Melanie Johnston-Hollit (Victoria)
Dale Carnegie (Victoria)
Slava Kitaeff (Auckland University of Technology)
Michael Hayes (University of Canterbury)

Credits:

Cover design: Colin Fox, Hyuck Chung
Logo photo: Colin Fox
L^AT_EX editor: C. Fox, T.C.A. Molteno, H.J. Lawrence and H. Chung
using L^AT_EX's 'confproc' class by V. Verfaillie

Printed in uniprint (www.uniprint.otago.ac.nz) — November 2009

Welcome from Tim Molteno, Conference Chair

Welcome to the Sixteenth Electronics New Zealand Conference - ENZCon. This is the third time we have hosted ENZCon at the University of Otago, and as always the Electronics Group has pulled together to make this event possible. We have also had invaluable support from the Dept of Physics as a whole.

As always, we have attempted to make the programme a mixture of work and fun, balancing some fairly challenging and busy sessions with plenty of time to rest, network, or relax. Beer and nibbles are provided at the Poster session, and we hope this will encourage good attendance, and animated discussion. The conference dinner has always been a highlight of ENZCon, and this is something we do well here at Otago. But along with the social events, we have a busy programme, with first time and veteran presenters from many branches of the field. It is great to see the topics covered in ENZCon papers changing with the times.

We would like to thank Bev Reynolds for all the work she has done organising, and dropping all her other work to fulfil last minute requests. Hyuck Chung was co-opted in to do the layout of this Proceedings at almost the last minute, and has done an awesome job. We couldn't have an ENZCon without Gary Bold - we thank him for being persuaded to return to a University of Otago ENZCon. Colin Fox persuaded Hyuck to drop everything & work tirelessly for the ENZ-Con cause, and has otherwise used his humour to stop Organising Committee members stressing - and with this in mind, his contribution to this welcome is "to remind us not to eat too much cheese".

Finally, thanks to our sponsors RFTest Solutions, Hoare Software Research, Nichecom and Electrottest. The contribution made year after year by these Trade Sponsors help keep ENZCon affordable, and therefore accessible. We hope that all of you have an enjoyable and informative conference.

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

Conference Program

Day 1 : Wed 18th Nov.

Registration and Nibbles at St Margarets College, 12:00 - 13:00

Opening Address by Gary Bold, 13:00

Oral Session 1 (13:20): General Electronics, Chair: Tim Molteno

- 1 The Comparison of Analogue and Digital One- Cycle Control Feedback Methods around the Output Stage in a Digital Audio Power Amplifier
C.D. Benton, D.A. Carnegie and P. Gaynor
- 7 A Model Predictive Control toolbox intended for rapid prototyping
J. Currie, and D.I. Wilson
- 13 Fast Counters using a Modified Fibonacci Number System
R.Ward, and T.C.A. Molteno
- 19 Differential Evolution and its Application to Intelligent Spectral Design
S. Soltic and A.N. Chalmers

Poster Session (18:30), Chair: Colin Fox

- 25 Design of an Improved Fuzzy Logic Controller Micro-Chip for Washing Machine
P. Lohani, and S.M.R. Hasan
- 31 Speedometer Calibration Unit
S. Nalli
- 37 Power Characterisation of IEEE 802.15.4 and Zigbee Wireless Networks
A. Prince-Pike, and J.D. Collins
- 43 Low-Cost Temperature Sensor on a Modern Submicron CMOS Process
R.P. Fisk and S.M.R. Hasan
- 49 A toy example of electrical impedance imaging using computational inference
E. Ma and C. Fox
- 55 Low-cost contour and groundwater mapping
C. Fox
- 61 Precise and Highly Selective Tonal Frequency Tracking
A. Opie, M Hayes, and J. Cavers
- 67 OGRE Compute Client Architectures
P. Suggate, and T. C. A. Molteno
- 71 Dynamic linear scaling in genetic algorithm for antenna design
P.J. Williams, and T.C.A. Molteno

Day 2 : Thu 19th Nov.

Oral Session 2 (09:00): FPGA, Imaging & Inference, Chair: Slava Kitaev

- 75 Design of a FPGA-based multi-channel front-end for synthetic aperture sonar
B.C. Bonnett and M.P. Hayes

- 81 Digital Receiver for Transient Radio Emission Array Detector Prototype (TREAD-P)
V. Kitaev, T.A.S.A. Perera, and G. Soudlenkov
- 87 Minimum Hamming Weight Representations for Irregular Symbol Alphabets
W. Kamp, A. Bainbridge-Smith, and M. Hayes
- 95 Advantages of 3D Time-of-Flight Range Imaging Cameras in Machine Vision Applications
A.A. Dorrington, C.D.B. Kelly, S.H. McClure, A.D. Payne, and M.J. Cree
- 101 Circular Polarised Patch Antenna for Medical Microwave Holography
O. Sergieiev, and V. Kitaev

Oral Session 3 (13:00): Radio Astronomy & Instrumentation, Chair: Melanie Johnston-Hollitt

- 107 Current New Zealand Activities in Radio Astronomy: Building Capacity in Engineering & Science for the Square Kilometre Array
M. Johnston-Hollitt, V. Kitaev, C.P. Hollitt, N. Jones, and T.C.A. Molteno
- 113 Automatic Detection of Supernova Remnants using the Circle Hough Transform
C.P. Hollitt, and M. Johnston-Hollitt
- 119 A comparative study of ‘inversion’, ‘optimization’, and ‘inference’ as frameworks for imaging
C. Fox and N. Dudley Ward
- 125 WaterLogged, an assistive device for Waka Ama
V.A. Fletcher and D.A. Carnegie
- 131 Comparison of Optimized Low-Power LNA topologies for 866 MHz UHF RFID
J Li and S.M.R. Hasan
- 137 A Digital CMOS Sequential Circuit Model for Bio-Cellular Adaptive Immune Response Pathway Using Phagolysosomal Digestion: A Digital Phagocytosis Engine
S.M.R. Hasan
- 143 Channel Sounding with Software Defined Radio
H.W.H. Jones, P.A. Dmochowski, and P.D. Teal
- 149 Design of a Low-power Narrow Band amplifier for HF Radio Applications using CMOS FETs
M Khurramand, and S.M.R. Hasan

Day 3 : Fri 20th Nov.

Oral Session 4 (09:00): Robotics, Mechatronics & Wireless Sensing, Chair: Dale Carnegie

- 153 Life Sign Detection on a Disposable Robotic Platform as Part of a Three-Tier System for Urban Search and Rescue Operations
B.M.M. Drayton, and D.A. Carnegie
- 159 The design and implementation of a reliable robotic pipe inspection system
S.J. Winch, and D.A. Carnegie
- 165 Environmental Monitoring with Wireless Sensor Networks
A. Ghobakhlou, T.A.S.A. Perera, P. Sallis, O. Diegel and S. Zandi
- 171 Wireless Soil Moisture Sensor for Vineyard Soil Monitoring
T.A.S.A. Perera, and J.D. Collins
- 177 The Slot Car Stig: Performance and Consistency of a Slot Car Driven by a Heuristic Algorithm in an Embedded Microcontroller
S. Kane, and J. Scott
- 181 **Index of Authors**

The Comparison of Analogue and Digital One-Cycle Control Feedback Methods around the Output Stage in a Digital Audio Power Amplifier

Carl D. Benton¹, D. A. Carnegie¹ and P. Gaynor²

¹School of Engineering and Computer Science,
Victoria University of Wellington,
Wellington, New Zealand.

²Department of Electrical and Computer Engineering,
University of Canterbury,
Christchurch, New Zealand

Abstract – A digital amplifier has benefits such as reduced power consumption and a smaller, lightweight form factor. However, for the Hi-Fidelity audio market, the quality of this amplifier configuration is typically lacking as many sources of error have been identified around the output stage in a switching amplifier. The One-Cycle Control scheme is a solution which has the ability to eliminate many of these errors in a single cycle. Current implementations have been of an analogue form, but due to the nature of a digital amplifier, it is reasonable to include this feedback control in the digital domain. Therefore, a digital form of One-Cycle Control is being developed, and is being compared to the analogue version. If successful, this digital implementation will be incorporated into commercial high-power, high-fidelity audio amplifiers by Perreux Industries Ltd.

1. INTRODUCTION

Audio power amplifiers targeted at the Hi-Fidelity market have typically been of a Class-A or Class-AB topology as Hi-Fidelity audio quality is easily and readily obtained with these configurations. However, these amplifiers are usually bulky and very inefficient with a typical amplifier reaching only 40-70% efficiency [1]. With the market today demanding a higher power output with a focus on energy and resource conservation, a smaller and more efficient amplifier is desirable.

The switching (or digital) amplifier has a significantly improved efficiency compared to the Class-A and AB techniques. However, until recently it has not been feasible to use this type of amplifier in a Hi-Fidelity audio environment. With the advancement of digital processing techniques, it has been possible to improve the Pulse Width Modulated (PWM) signal generation process, and recent research around the output stage has seen an increase in the number of switching-type amplifiers available to consumers. However, many of

the available amplifiers are of a Class-D configuration or similar, and are not compatible with a digital audio format. Also, the digital amplifiers that are available still leave room for improvement.

It is proposed that if a digital amplifier can be developed to perform to an equal or better standard than that of the current linear amplifiers available (as far as the human ear is concerned) then this new digital amplifier technology would have significant market potential. In quantitative terms, this typically means an amplifier with a Signal-to-Noise Ratio (SNR) of better than 96 dB (CD Quality), and a Total Harmonic Distortion (THD) of less than 0.01 %.

2. BACKGROUND

The common Class-A and Class-AB classes of amplifier are an analogue form of amplification where the devices in the output amplification stage spend the majority of time operating in their linear region. This means that these devices are dissipating power for the majority of the audio cycle and thus the amplifier is rather inefficient and bulky due to the large heat sinks required.

A common solution is a lighter, cheaper and more efficient form of amplifier commonly called a switching amplifier. A switching amplifier is typically based around the concept of a PWM signal where the width of each pulse is proportional to the amplitude of the analogue control signal being amplified. In such an amplifier, this PWM signal is what is amplified, thus the devices in the output stage are ideally either on or off, and never operate in the linear region at any time. The resulting amplified PWM signal is then filtered with a low-pass filter to remove the switching component, and the resulting audio signal is delivered to the loudspeaker. Such an amplifier can reach efficiencies greater than 90 % [2,3,4].

The switching amplifier is commonly implemented as a Class-D type. A Class-D amplifier takes an analogue audio input signal and from this generates the PWM signal to be amplified. To achieve this, a reference triangle-wave signal is generated, and this is compared with the audio signal. These signals are shown in Figure 1, which shows that when the audio signal is greater than the triangle-wave, the PWM output is high (1), and when the audio signal is less than the triangle-wave, the PWM output is low (0).




Figure 1: Source, reference triangle wave and resulting PWM signals.

However, in a digital amplifier the input is taken from a digital representation of the audio, such as Pulse Code Modulation (PCM) encoded audio or Sony/Philips Digital Interconnect Format (SPDIF), and the PWM signal is generated in the digital domain. Digital amplifiers have become an increasingly active research topic in recent years, especially as its application to consumer audio has become apparent [5,6].

In this case, the PWM signal cannot be derived with an analogue circuit using the method above. A simplified block diagram for a digital amplifier is shown in Figure 2, and the focus of this paper is on feedback methods around the output stage.




Figure 2: Digital Amplifier Block Diagram

3. SOURCES OF ERROR

There are many sources of error in the output stage leading to an amplified PWM signal that is far from ideal. Sources of error, as shown in Figure 3, include power supply ripple, voltage drops from the switching devices (typically MOSFETs, and sometimes flyback

diodes), overshoot and undershoot, rise and fall times and turn on and turn off delays. Each of these unwanted artifacts is a source of THD in the resulting output audio signal.




Figure 3: Sources of Error in the Switching Converter

To overcome these sources of error, a feedback control is required to provide corrective measures. This can be achieved by adjusting the duty-cycle of the PWM signal accordingly. Although several feedback methods have been discussed previously [7,8,9], the focus of this paper is on One-Cycle Control [10,11] as it can correct for errors in a single PWM cycle period allowing for a faster correction and therefore resulting in a more accurate PWM signal.

4. ONE-CYCLE CONTROL

The feedback around the output stage is called One-Cycle Control [10,11] where the integration of each period of the PWM signal is forced to a value that is equal to, or proportional to, the control reference.

Previously, One-Cycle Control has been implemented in the analogue domain as shown in Figure 4. The circuit includes an integrator (*INT*) to integrate the area under the PWM output signal over time, and a comparator (*CMP*) to compare the integrator output to a control reference voltage. As shown in Figure 5, when the integrator output V_{int} reaches the reference value $-V_{ref}$, the PWM driving signal changes state and the integrator is reset. These signals are depicted in Figure 5 and show how the reference signal is negative for a PWM duty cycle of greater than 50%, as the output PWM signal is the inverse of the driving PWM signal Q .




Figure 4: One-Cycle Control




Figure 5: One-Cycle Control Waveforms

The act of resetting of the integrator circuit every cycle is performed by discharging capacitor $C1$ via switch Sr . It has been found that an error is introduced due to the finite time taken to discharge the capacitor [12]. This leads to a reduction in the area under the output PWM signal, and therefore results in an output that is non-linear. Another problem with an analogue One-Cycle Control circuitry are prone to noise coupled from the relatively high voltage output stage. Furthermore, due to the nature of a digital amplifier, the control reference signal resides in the digital domain, and thus a conversion to an analogue reference signal, via a Digital-to-Analogue Converter (DAC) is required for the One-Cycle Control circuitry.

5. THE DIGITAL IMPLEMENTATION

A digital implementation of the One-Cycle control method is proposed in an attempt to eliminate the problems associated with the analogue version and is shown in

Figure 6. As the control reference in a digital amplifier is already in the digital domain, it makes sense to implement the One-Cycle Control algorithm in the digital domain also. As shown, the switched output signal Vs is divided down to a suitable voltage via a resistor divider consisting of $R1$ and $R2$, and fed into the Analogue-to-Digital Converter (ADC). The digital values are then read from the ADC by a logic circuit in the Programmable Logic Device (PLD) which performs a digital summation and comparison, and generates the signals for the Driver.




Figure 6: Digital form of One-Cycle Control

The primary benefit of the digital implementation is the elimination of the requirement to reset an analogue integrator. The digital equivalent of this function is simply to clear or reset a digital counter back to zero or a predefined value. This can be achieved by a PLD logic circuit in a single clock cycle, and with a conservative clock frequency of 150 MHz, this equates to a time of less than 7 ns.

Another benefit of the digital implementation is that any noise generated from the higher voltages and large transient currents involved in the output stage will not affect the operation of the logic circuit and thus the control results are essentially immune. There is also a potential reduction in component count as much of the control can be performed inside a single PLD chip. Alternatively, this function can be performed within a digital processor that performs other signal processing functions, which is already present in many digital amplifier systems.

A further benefit of a digitally implemented One-Cycle Control is the ability to subtract the reference control value from the digital integrator rather than resetting the sum to zero. This means that any time domain errors associated with the sampling of the PWM signal can be accounted for, and corrected in the subsequent cycle. For example, if the integrator had a sum of 850, and a value of 100 was read from the ADC, while the target value was 900, then resetting the integrator to zero will cause an error of 50. However, if the target value was subtracted from the sum, the counter will begin the following cycle with a value of 50 ($850 + 100 - 900$), and the average of the two cycles will be correct. Not only is this almost impossible to achieve with an analogue implementation, but information is also lost during the time taken to reset the integrator.

However, the digital implementation of the One-Cycle Control method also has two disadvantages. Firstly, as the PWM output signal must be converted to a digital form via an Analogue-to-Digital Converter (ADC), a quantization error is introduced. Secondly, the limited clock frequencies of a digital circuit also introduce an error as the PWM output signal can only be sampled a finite number of times per second, and thus the potential

to lose error information present in the PWM output signal is introduced.

Work is currently underway to examine whether the benefits of a digital form of One-Cycle Control will outweigh the disadvantages. If it does, then this will lead to an overall reduction in the error component of the output PWM signal, and an improvement in the THD and SNR of the resulting output audio signal that is delivered to the loudspeaker is achieved.

6. PROGRESS

So far several prototype versions have been developed. The latest revision is two separate 4-layer circuit boards for a comparison between both the analogue and digital versions which are shown in Figure 7 and Figure 8. Both circuits include a Xilinx XC2C256 complex programmable logic device (CPLD) running at 150 MHz, and development and testing is being performed with an output PWM frequency of between 250 kHz and 800 kHz.



Figure 7: Prototype Analogue (Top) and Digital (Bottom) One-Cycle Control Boards - Top View

The analogue prototype has a standard One-Cycle Control feedback implementation, where a digital reference signal is converted to an analogue reference voltage using a 16-bit ADC. A high speed op-amp with a sufficient slew rate has been chosen for the integrator function. Also, a small value for the capacitor (1 nF) has been chosen, combined with a low impedance analogue switch, to minimize the time to reset the integrator.




Figure 8: Prototype Analogue (Top) and Digital (Bottom) One Cycle Control Boards - Bottom View

The digital implementation uses a high frequency 10-bit ADC to convert the output PWM signal to a digital value at a rate of 75 million samples per second (MSPS). These values are fed into a digital integration counter and comparator circuit in the CPLD which generates the resulting PWM signal to be amplified. A tradeoff exists between the sampling frequency and resolution of the ADC. Of course, to get the most accurate feedback control, the PWM signal should be sampled at the highest possible frequency with the highest possible resolution (number of bits). However, the current technology available is somewhat limiting, and it is difficult to obtain a reasonably priced ADC that can operate above 100 MHz, with a resolution of 16 bits. Those that were able to be sourced cost over NZ\$150 ea.

The two feedback circuits are operational and the investigation into comparison of the two amplifiers is currently ongoing. One issue that has arisen is an instability that occurs when the negative power supply rail is greater than the positive, which has been identified and discussed in [12]. The result is a sub-harmonic oscillation at half the switching frequency as shown by the pink trace in Figure 9. As the solution for the analogue implementation is to add a common offset to V_s and $-V_{ref}$, an equivalent solution should exist for the digital implementation that can be included in the CPLD. However, the issue can be avoided for the purpose of comparative testing and is therefore not a concern.




Figure 9: Sub-harmonic Oscillation in Analogue Feedback

Once some definitive comparative results have been obtained from the two prototypes, the effect of constraints such as the resolution and sampling frequency of the ADC, the clock frequency of the CPLD and the time taken to reset the integrator, will be clearer. At that stage, the feasibility of a digital implementation of One-Cycle Control will be determined, and its place in a commercial digital audio amplifier will be known.

7. REFERENCES

- [1] Elliot, Rod. (1999). "Power Amplifier Efficiency Explained." from <http://sound.westhost.com/efficiency.htm>.
- [2] Simon C. Li, Vincent Chia-Chang Lin. "A High Efficiency 0.5W BLT Class-D Audio Amplifier with RWDM Technique", Asia and South Pacific Design Automation Conference, 27-30 Jan 2004, pp. 307-309.
- [3] Soo-Hyoung Lee, Jae-Young Shin, Ho-Young Lee, Ho-Jin Park, Kristian L. Lund, Karsten Nielsen and Jae-Whui Kim. "A 2W, 92% efficiency and 0.01% THD+N class-D audio power amplifier for mobile applications, based on the novel SCOM architecture", Custom Integrated Circuits Conference, 3-6 Oct 2004, pp. 291-294
- [4] Soo-Chang Choi, Jun-Woo Lee, Woo-Kang Jin, Jae-Hwan So and Suki Kim. "A Design of a 10W Single-Chip Class D Audio Amplifier with Very High Efficiency using CMOS Technology", IEEE Transactions on Consumer Electronics, **Vol. 45**, No. 3, Aug 1999, pp. 465-473.
- [5] Sandler, Mark. "Digital Amplification for Consumer Audio", International Conference on Consumer Electronics (ICCE), June 2000, pp. 402-403.
- [6] H D Farrar, R E Hioms. "Development of a Commercial Digital Amplifier", Audio DSP - Circuits and Systems, **219**, Nov 1993, pp. 3/1-3/5.
- [7] Pallab Midya, Bill Roeckner and Steve BergStedt. "Digital Correction of PWM Switching Amplifiers", IEEE Power Electronics Letters, **Vol. 2**, No. 2, June 2004, pp. 68-72.
- [8] Jwu-Sheng Hu, Keng-Yuan Chen. "Implementation of a Full Digital Amplifier Using Feedback Quantization", American Control Conference, June 2006, pp. 2759-2764.
- [9] Nielsen, Karsten. "PEDEC - A Novel Pulse Referenced Control Method for High Quality Digital PWM Switching Power Amplification", Power Electronics Specialists Conference, **Vol. 1**, May 1998, pp. 200-207.
- [10] Zheren Lai, Keyue M. Smedley. "One Cycle Control of Bipolar Switching Power Amplifiers", US5617306, 1 April 1997, pp. 12.
- [11] K. Smedley, S. Cuk. "One Cycle Control of Switching Converters", IEEE Transactions on Power Electronics, **Vol. 10**, No. 6, November 1995, pp. 625-633.
- [12] Zheren Lai, Keyue M. Smedley. "A New Extension of One Cycle Control and Its Application to Switching Power Amplifiers", IEEE Transactions on Power Electronics, **Vol. 11**, No. 1, January 1996, pp. 99-105.

A Model Predictive Control toolbox intended for rapid prototyping

Jonathan Currie* and David I. Wilson†

*Electrical & Electronic Engineering
Auckland University of Technology
Auckland, New Zealand

†Industrial Information & Control Centre
Auckland University of Technology
Auckland, New Zealand

Abstract:

Model predictive control (MPC) is a powerful control strategy that handles multi-variable constrained dynamic systems. It is however very different from classical feedback controllers, and demands different design skills. Furthermore, until recently MPC, due to its internal constrained optimiser, was ill-suited for embedded controllers designed to tackle high-speed applications. This paper introduces an object-orientated MPC framework, an MPC graphical-user interface intended for testing and tuning MPC applications, and an MPC implementation in SIMULINK suitable for subsequent compiling and auto-code generation for embedded devices.

Keywords:

Model Predictive Control, embedded, development tools.

1 INTRODUCTION

Model Predictive Control, or MPC, is one of a family of predictive controllers that have seen widespread use in the chemical processing plants over the last few decades, not the least because of its elegant constraint handling abilities, [1]. Today MPC is regarded as one of the most successful advanced process control strategies available, [2]. In addition to constraint handling, it possesses attractive features such as intuitive tuning, the ability to optimally handle non-square systems, and it must be admitted slightly less convincingly, nonlinear plant dynamics, [3].

However MPC, or at least the current commercial implementations such as those surveyed in [4], suffer a serious shortcoming; being an optimal controller, they are computationally demanding. Consequently to date MPC applications have been mainly in the chemical processing industries in part due to three reasons, [5]. First the industrial plant time constants are relatively slow, in the orders of minutes to hours; second, operational constraints on a chemical plant are often safety-critical and therefore should be an integral part of the control design, and finally the complexity of the MPC design means that it is more economical to implement it once on a plant-wide or unit-wide application. These reasons may also explain the paucity of MPC applications outside the process control community. The MPC algorithm is described in detail in numerous texts, [2, 6, 7], all of which give an appreciation of the level of computation required.

Notwithstanding the computational complexity of MPC, we believe that there is no reason why an MPC

could not work adequately on embedded hardware such as field programmable gate arrays (FPGAs) at high speed. If this is achievable, as recent work in [8] seems to suggest, then MPC could be applied to much faster systems such as unmanned vehicles, auto-pilots, intelligent sensors where the same control benefits could open up new opportunities in robotics and intelligent systems.

Currently however, a generic and flexible implementation of MPC requires substantial processing even to achieve sampling rates of around 1kHz, far more for example than classical optimal controllers such as LQG. For smaller, power or weight critical applications (such as those mentioned above), we need an alternative, high speed and small footprint processing platform, such as an FPGA. Developing an MPC in embedded hardware is not unique, (see for example [9–11] for FPGA applications and [12] for PLC), but to date, substantial algorithm modifications such as shortening the prediction horizons, or replacing the online optimisation with a table-lookup are required for it to be practical, [13].

2 THE QUADRATIC PROGRAM INSIDE AN MPC

The model predictive controller delivers a future sequence of control moves, $\Delta \mathbf{u}$, over the immediate future control horizon, N_c , such that the cost function

$$J = \sum_{j=1}^{N_p} \|\hat{\mathbf{y}}_{k+j|k} - \mathbf{y}_{k+j|k}^*\|^2 + \sum_{j=1}^{N_c} \|\lambda \Delta \mathbf{u}_{k+j|k}\|^2 \quad (1)$$

is minimised. The vectors $\hat{\mathbf{y}}, \mathbf{y}^*$ are the estimated future output of the plant, and future setpoints, and λ is a weighting factor. The prediction and control horizons, N_p and N_c respectively, specify how long the optimisation is to take place, and are important tuning parameters. The objective function is constrained by the plant dynamics, in this paper assumed linear, but not necessarily stable, and possible linear output, input, and input rate constraints. With some algebraic manipulation (see [2, 7] for the details), the objective function and constraints can be re-written as

$$\begin{aligned} \min_{\Delta \mathbf{u}} J &= \frac{1}{2} \Delta \mathbf{u}^T \mathbf{H} \Delta \mathbf{u} + \mathbf{f}^T \Delta \mathbf{u} \\ \text{subject to: } &\mathbf{A} \Delta \mathbf{u} \leq \mathbf{b} \end{aligned} \quad (2)$$

which is a standard quadratic program to be solved each sample time.

2.1 MPC QP dimensions

The characteristics of MPC problems that until recently were dominated by the process control community were that the sampling rate was relatively sparse when compared to the dominant time constant; the systems were linear, or at least possessed differentiable nonlinearities; and the plants of interest were dominated by deadtimes or unusual responses such as non-minimum phase or inverse responses. The latter characteristic being a consequent of the fact that the default PID control did not suffice.

For a MIMO system with n states, m inputs and p outputs, the number of decision variables of the QP formulation of MPC in Eqn. 2 is $n_d = mN_c$. From this it follows that the dimensions of \mathbf{H} are $(mN_c \times mN_c)$, and the dimensions of the constraint matrix \mathbf{A} are $((2pN_p + 4mN_c) \times mN_c)$. Thus for a typical problem in process control, say $n = 10, m = p = 2$, with horizons $N_p = 50, N_c = 25$, the transpose of the constraint matrix, \mathbf{A} , of dimensions (50×400) is shown in Fig. 1. These dimensions, admittedly on the large size, are consistent with the recommendations given in [14, p555].

Clearly the \mathbf{A} matrix exhibits considerable structure, and this structure could be exploited in the QP algorithm. The diagonal lines are the upper and lower constraints on $\Delta \mathbf{u}$, the triangular sections are the upper and lower constraints on the \mathbf{u} , while the trapezium structures are the output maximum and minimum constraints. One option not utilised at present is that the straight decision variable bounds are not removed from \mathbf{A} to be treated separately.

2.2 QP algorithms

Currently there are two approaches for solving the QP in Eqn. 2, one is solving the full classical QP at each iteration perhaps exploiting solutions obtained at previous iterations. The other approach, popular for embedded applications, is known as parametric MPC is where the solution

is computed *a priori* thereby reducing the online computation to a table look-up, [15, 16]. Attractive though this may be, the approach does not scale well, so consequently control designers typically employ very simple models with very short time horizons.

A recent example highlighting the difficulties of implementing MPC on FPGA hardware is given in [11] where the largest report application is a model helicopter application with 6 states, 2 inputs, 3 outputs and only input constraints. With a prediction horizon of 1, they achieve a loop time of $6.45 \mu\text{s}$ with 202kB of FPGA memory. (To put this in perspective, our intended target FPGA has only 45kB on chip.) Increasing the prediction horizon to 2 (which equates to a QP with 4 decision variables), can be solved in $10 \mu\text{s}$, which while extremely fast, requires an unrealistic 62MB to store the parametric representation of the system. Such poor scaling properties seriously questions the reason for employing MPC in the first place, so for this reason our approach has been to stay with the full QP, but to streamline and optimise the algorithm to deliver a reasonably fast, reliable solution.

Our MPC toolbox allows the user to choose one of five QP solvers. Only one, `quadprog` from the Optimisation toolbox for MATLAB, [17], is designed to solve non-convex QP problems. However, numerical round-off issues aside, the MPC QP problems are always positive definite. The second QP solver is an adaption of the Hildreth algorithm in [6]. While this algorithm can be expressed succinctly in MATLAB code, it regularly fails to converge on modestly sized well-posed problems. The GUI also contains two algorithms from an earlier toolbox from [18]. One of these is an interior point, the other is an active set method. The final QP algorithm is one adapted from [19] which is our solver of choice.

2.3 Performance of QP algorithms

Fig. 2 shows the execution timing for solving the QP for a randomly generated MPC problem with a 15 state system with 3 inputs, 3 outputs and varying horizon lengths of $2 < N_c < 30$ with $N_p = N_c + 10$. The timings in Fig. 2 are relative to Wright's QP algorithm. All three algorithms generated solutions of comparable accuracy.

One heuristic we have added to the QP solution step is that we first solve the relaxed optimisation problem assuming no constraints are active,

$$\Delta \mathbf{u}^* = -\mathbf{H}^{-1} \mathbf{f}. \quad (3)$$

As shown subsequently in Fig. 4, for the majority of the samples, this is the case so we deftly avoid the full QP solution step.

As an aside, it is interesting to note that automatically compiling the MATLAB Wright QP algorithm to C slows down the code by an order of magnitude for the sizes of




Figure 1: The sparsity pattern and structure of the constraint matrix, \mathbf{A}^T in the MPC formulation given in Eqn. 2. The blue symbols are positive values; red negative.




Figure 2: Timing of the Hildreth and Matlab's `quadprog` QP solvers compared to Wright's QP for a series of MPC problems with random dynamics.

problems of interest as shown in Fig. 3. Carefully hand-coding a C implementation and exploiting the appropriate BLAS and Lapack subroutines is slightly faster for small problems, but again drops in speed to only 80% of the carefully optimised MATLAB implementation.

This is, we believe, due to recent advantages in the just-in-time compiling of the MATLAB code, the fact that the profiling reveals that the majority of the time spent is in key linear algebra components, and that we have spent considerable effort optimising the MATLAB code itself. We also note that MATLAB will exploit the modern architecture of today's gigahertz processors.

2.4 Performance of the MPC algorithm

There is a subtle difference between the performance of the algorithm solving a constrained QP, and the overall MPC. One of the interesting features of the MPC controller is that there is a huge difference in computation load between solving a constrained optimisation QP with no prior information, and solving an unconstrained optimisation problem which is simply a matter of solving Eqn. 3 where the Hessian \mathbf{H} can be factored offline. In practice, one pre-computes the Cholesky factor of $\mathbf{H} = \mathbf{R}^T \mathbf{R}$, and then online one must simply execute one forward substitu-




Figure 3: Comparing the timing of an auto-generated C code (`mcc`) implementation of Wright's QP, a hand-coded C implementation (`mex`), and an optimised MATLAB m-file implementation.

tion followed by a backward substitution, or in MATLAB notation, $\mathbf{d}u = \mathbf{R} \setminus (\mathbf{R}' \setminus \mathbf{f})$.

For cases when the optimisation problem is constrained, then some savings can be made by using the previous optimal solution, appropriately shifted, as an initial start guess. As it turns out, this warm start procedure is not as beneficial as it seems cutting down less than a quarter of the iterations. Fig. 4 shows the near-perfect MPC control and timing breakdown for a second order non-minimum phase example, $G(s) = (-5s + 1)/((3s + 1)(s + 1))$. (Note that due to the significant right-hand plane zero, PID control of such an inverse response plant is a challenge.) In this example, the limiting constraint is the rate limiter on Δu , and the controller is allowed to take advantage of future setpoint changes.

In this example, one can see that the seemingly unimportant assembly of the vectors \mathbf{f} and \mathbf{b} in Eqn. 2 take around 0.6 ms which is comparable to the 1ms the QP takes to solve the constrained optimisation problem using around 10 iterations. We also note that 10 iterations lies at the upper end of the range recommended in [8] after in-




Figure 4: Timing breakdown of an MPC implementation for a SISO non-minimum phase (inverse response) application.

investigating 12 MPC problems with different dimensions.

3 THE MPC OBJECT ENVIRONMENT

The design of an MPC controller encompasses a wide variety of topics from model identification, optimisation algorithms, tuning, and testing. For this reason, we have chosen to develop an MPC object which in turn exploits the advantages of the relatively recent developments in the object-orientated capabilities of MATLAB. This object not only contains the dynamic plant (in any one of various forms), but also plant models (which may not necessarily be the same), controller specifications such as prediction horizons, sample time, controller weights etc, and the input, output and rate constraints.

The development of such a tool was not prompted for undergraduate university education, but rather industrial use. In our activities running industrial short courses in control, we find that those most interested in MPC are control engineers who are comfortable with PID controllers, and are very curious about the potential of MPC. Specifically they want to play with the controller, varying horizons, introducing increasing amounts of model/plant mismatch, adjusting constraints, and especially comparing it to controllers in their comfort zone such as PID and variants. Even those companies that do run a commercial MPC such as Pavilion Technologies' Pavilion8, Honeywell's RMPCT or Perceptive Engineering find that this tool is easy to use and most of all encourages safe exploration.

At present, the tool does not allow the possibility of time-varying constraints, time-varying weightings, or blocking, although that functionality could be easily added. While the internal model used in the MPC controller is linear, one can use a nonlinear plant model, perhaps to explore model/plant mismatch, or even control an actual plant using the SIMULINK implementation described in section 3.2 with a data acquisition card. It is also possible to enable the MPC controller to exploit future setpoint changes if known.

Such a toolbox is not unique of course, the commercial MPC tool box, [20] is one example, but we believe our version has a number of advantages. First it simplifies some of the less-used MPC options such as time-varying constraints, and the scrolling screen with the future predictions is a useful addition. Second the underlying object design means that we can re-use this tool in many other controller design studies. But most importantly, is that this MPC tool is designed to support our parallel work in developing development tools for MPC to run on FPGAs. And finally, there are some cost advantages for us.

3.1 The MPC graphical user interface

To aid the design of model-predictive controllers, and to assist new users in MPC we have developed a graphical user interface (GUI) for the design and testing of MPC controllers. Fig. 5 shows the interface demonstrating the MIMO control of a 3 degree of freedom model helicopter.

The GUI design shows the input/output trends on the left, complete with green/red lights showing if constraints are active. The right panels allow the user to interactively adjust the various horizons, controller weights, and constraints.

The tool accepts a wide variety of models: linear discrete or continuous (or a mixture), transfer function or state-space, with, or without deadtime. One can even specify a nonlinear plant, although the model internally used within the MPC will be linearised. The tool also allows one to specify a different dynamic system for the plant than for the model thus enabling the exploration of model/plant mismatches.

This is no limit to the size of the models that can be used, but the practicalities of the GUI plotting are such that systems with more than half a dozen i/o rapidly get confusing. However in these cases, one can readily export the data for plotting externally.

The screen capture in Fig. 5 also demonstrates the power of MPC. In this case we have a non-square model helicopter plant with 2 input rotors and 3 outputs: azimuth, elevation and pitch. Also depicted are the output constraints (dotted lines), and the future predictions for both the inputs and outputs to the right of the solid line two thirds along the scrolling trend.

It is also possible to exchange the system with raw




Figure 5: The Model Predictive Control graphical user interface controlling a 3 DOF Quanser helicopter. Note the depiction of the future input and output predictions.

MATLAB thereby utilising all the existing controller design tools within that environment as well as explicitly constructing the MPC object. For example, for the SIMO plant and model

$$G = \begin{bmatrix} \frac{s-1}{s^2+4s+5} \\ \frac{s+1}{s+2} \end{bmatrix}, \hat{G} = \begin{bmatrix} \frac{1.2s-1.1}{1.2s^2+4.4s+5.1} \\ \frac{1.4s+1.7}{1.5s+2.2} \end{bmatrix} \quad (4)$$

with various constraints is coded with

```

1 % Specify MIMO plant & model
  Gp=tf([1,-1],[1 2]), ...
        [1 1],[1 4 5]); % Plant
  Gm=tf([1.2 -1.1],[1.5 2.2]), ...
        [1.4 1.7],[1.2 4.4 5.1]); % Model

6
  Ts=0.1;Np=10;Nc =5; % Ts, Nc, Np
  setp = 1;
  % Specify i/o constraints, & controller weights
  Constraints.u = [-5 5 2.5;
                  -5 5 2.5];
11
  Constraints.y = [-5 5];
  Weights.uwt = [0.1 0.2]';
  Weights.ywt = [0.1]';

16 Ex1=jGUI(Gp,Gm,Weights,Constraints, ...
           Np,Nc,Ts,State_Est,setp);

```

The MPC object, Ex1 created in the last line of the script can now be directly loaded into the GUI.

3.2 The MPC controller block for Simulink

Fig. 6 shows the MPC controller as implemented in SIMULINK. One needs only to drop our MPC object into the MPC controller block, and, provided the plant is not too dissimilar, the simulation will reproduce a controlled response. Since the SIMULINK tool uses the same MPC object as the GUI described in section 3.1, one could easily export the models and associated properties (constraints, controller horizons etc, and algorithmic details) to the SIMULINK block.




Figure 6: Implementing the MPC controller block for an arbitrary plant in the SIMULINK environment.

The jMPC toolbox is available from www.auckland.ac.nz/i2c2.

4 CONCLUSIONS

The intuitiveness, the ease of tuning, and the optimality are persuasive features of Model Predictive Control which in part account for the increasing interest in industry. However the underlying philosophy of the controller is quite a departure from the classical feedback PID which industry is comfortable with, so to sell the idea of advanced control, we need elegant, easy to use tools. This paper describes one implementation of an MPC controller that is suitable for real-time control, and is the first stage in the development of predictive controllers designed for high-speed embedded applications such as those implemented on FPGA hardware.

The maximum obtained sampling frequency for a challenging SISO model is over 1kHz on a 3GHz desktop PC, while sampling rates of 200Hz are possible for the heavily constrained MIMO helicopter case. These numbers, being similar to those reported in [8], suggest that this implementation qualifies for the ‘fast’ MPC label.

ACKNOWLEDGEMENTS

Financial support to this project from the Industrial Information and Control Centre, AUT University, New Zealand is gratefully acknowledged.

REFERENCES

- [1] C. Garcia, D. Prett, and M. Morari, “Model predictive control: Theory and practice — A survey,” *Automatica*, vol. 25, no. 3, pp. 335–348, 1989.
- [2] J. M. Maciejowski, *Predictive Control with Constraints*. Prentice–Hall, 2002.
- [3] D. I. Wilson and B. R. Young, “The Seduction of Model Predictive Control,” *Electrical & Automation Technology*, pp. 27–28, Dec/Jan 2006. ISSN: 1177-2123.
- [4] S. J. Qin and T. A. Badgwell, “A survey of industrial model predictive control technology,” *Control Engineering Practice*, vol. 11, pp. 733–764, July 2003.
- [5] R. Ross, “Revolutionising model-based predictive control,” *Computing & Control Engineering Journal*, vol. 14, no. 6, pp. 26–29, 2004.
- [6] L. Wang, *Model Predictive Control System Design and Implementation using Matlab*. Springer, 2009.
- [7] J. A. Rossiter, *Model-Based Predictive Control: A Practical Approach*. CRC Press, 2003.
- [8] Y. Wang and S. Boyd, “Fast model predictive control using online optimization,” in *Proceedings of the IFAC World Congress*, (Seoul, Korea), pp. 6974–6997, July 2008.
- [9] P. Dua, K. Kouramas, V. Dua, and E. Pistikopoulos, “MPC on a chip — Recent advances on the application of multi-parametric model-based control,” *Computers & Chemical Engineering*, vol. 32, no. 4-5, pp. 754 – 765, 2008.
- [10] L. G. Bleris, J. Garcia, M. V. Kothare, and M. G. Arnold, “Towards embedded model predictive control for system-on-a-chip applications,” *Journal of Process Control*, vol. 16, no. 3, pp. 255–264, 2006. Selected Papers from Dycops 7 (2004), Cambridge, Massachusetts.
- [11] T. A. Johansen, W. Jackson, R. Schreiber, and P. Tøndel, “Hardware synthesis of explicit model predictive controllers,” *IEEE Transactions Control Systems Technology*, vol. 15, no. 1, pp. 191–197, 2007.
- [12] G. Valencia-Palomo and J. Rossiter, “Auto-tuned predictive control based on minimal plant information,” in *International Symposium on Advanced Control of Chemical Processes ADCHEM 2009*, (Istanbul, Turkey), pp. 582–587, International Federation of Automatic Control, 12–15 July 2009.
- [13] K. Ling, B. Wu, and J. Maciejowski, “Embedded Model Predictive Control (MPC) using a FPGA,” in *IFAC Proceedings of the 17th World Congress*, (Seoul, South Korea), pp. 15250–15255, 6–11 July 2008.
- [14] D. E. Seborg, T. F. Edgar, and D. A. Mellichamp, *Process Dynamics and Control*. Wiley, 2 ed., 2005.
- [15] E. N. Pistikopoulos, V. Dua, N. A. Bozinis, A. Bemporad, and M. Morari, “On-line optimization via off-line parametric optimization tools,” *Computers & Chemical Engineering*, vol. 24, no. 2–7, pp. 188–188, 2000.
- [16] E. N. Pistikopoulos, M. C. Georgiadis, and V. Dua, “Parametric programming & control: from theory to practice,” in *17th European Symposium on Computer Aided Process Engineering* (V. Plesu and P. S. Agachi, eds.), vol. 24 of *Computer Aided Chemical Engineering*, pp. 569–574, Elsevier, 2007.
- [17] Anon., “Optimization Toolbox User’s Guide,” tech. rep., MathWorks, 3 Apple Hill Drive, Natick, MA, USA, 2009.
- [18] J. Åkesson, “MPCtools: A toolbox for simulation of MPC controllers in Matlab,” tech. rep., Lund Institute of Technology, Lund, Sweden, January 2006.
- [19] S. Wright, “Applying new optimization algorithms to model predictive control,” in *Fifth International Conference on Chemical Process Control CPC V*, pp. 147–155, CACHE Publications, 1997.
- [20] A. Bemporad, M. Morari, and N. L. Ricker, “Model predictive control toolbox 3,” tech. rep., The Mathworks, 2009.

Fast Counters using a Modified Fibonacci Number System

Roy Ward, Dr T. C. A. Molteno*

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

Abstract: We describe a fast counter design that uses a modification of the Fibonacci number system. This design achieves high speed with a relatively small amount of logic. We compare this with other common counter designs including radix-2 counters, redundant radix-2 counters and linear feedback shift registers, investigate extensions of this design to saturating counters and counters using k-Generalised Fibonacci numbers, and describe some applications.

Keywords: Counter, Number Representation, Fibonacci

1 INTRODUCTION

Fast, power-efficient event-counters are widely used in processor power management [1, 2], network management [3] and high-speed measurement in particle physics experiments [4].

One issue with radix-2 (the conventional ‘binary’) counters is that there is a requirement to propagate the carry bit each cycle, so that operations such as

$$01111111_2 + 1_2 = 10000000_2$$

are performed correctly. This requires carry logic [5] which is either slow or expensive. The work that has been done on this includes designs for anticipating the carry to give a constant time counter [6] and fast carry lookahead adders [7].

Various schemes for redundant counters [8] which have more than one register per bit have also been proposed that enables only one stage of carry propagation is required. Such a scheme is known as *carry free* addition.

In some circumstances, Linear Feedback Shift Registers (LFSRs) [9] can be used for extremely fast counters where reading the value of the counter is not required.

In this paper we propose a design for a carry free counter based on the Fibonacci number system.

1.1 Fibonacci Number System

The Fibonacci number sequence, $\{F_i\}$ is defined by the recurrence:

$$F_n = \begin{cases} 0 & n = 0 \\ 1 & n = 1 \\ F_{n-1} + F_{n-2} & n > 1 \end{cases} \quad (1)$$

The first few Fibonacci numbers are 0, 1, 1, 2, 3, 5, 8, 13, 21, 34, ... It can be shown that $F_n \simeq \frac{\phi^n}{\sqrt{5}}$, where $\phi = \frac{\sqrt{5}+1}{2}$ is the Golden Ratio.

Zechendorf’s Theorem[10] states that any non-negative integer n can be uniquely represented as

$$n = \sum_{j=1}^r F_{i_j} \quad (2)$$

where $\forall 1 \leq j \leq r, i_j \geq 2$ and $\forall 1 \leq j < r, i_j \geq i_{j+1} + 2$.

The Fibonacci number system [11] uses this to represent any non-negative integer, n , as a sequence $b_{m-1}b_{m-2} \cdots b_1b_0$ where each b_i is either ‘0’ or ‘1’ and the subscript F denotes that position i has the value F_{i+2} , so

$$n = \sum_{i=0}^{m-1} b_i F_{i+2} \quad (3)$$

In the Fibonacci number system there are never two adjacent 1’s.

1.2 Modified Fibonacci number representation

The Fibonacci number representation in the previous section does not immediately make the system carry free. For example

$$010101010101_F + 1_F = 100000000000_F \quad (4)$$

If the requirement that there are never two adjacent 1’s is relaxed, changing it to a requirement that there are never more than two adjacent 1’s (which will be referred to as the *adjacency constraint*), we obtain a *modified Fibonacci*

representation in which an efficient counter can be designed. Using this, Equation 4 becomes

$$010101010101_F + 1_F = 010101010110_F \quad (5)$$

Note that the uniqueness of the representation is lost, so this is no longer a number system [11], so we call it a number representation. For instance the number 16 can be represented by 100100_F , 100011_F or 011011_F . 011100_F would also represent this number, but would violate the adjacency constraint.

2 THE FIBONACCI COUNTER

If we consider an integer n represented in the modified Fibonacci number representation, we can define an increment operator $\sigma(n)$ with the following two steps on each group of three adjacent elements in the representation of n :

1. Partially perform any outstanding carries by examining once each sequence of adjacent three bits (we allow overflow by temporarily adding a leading 0 and dropping that bit afterwards — if that bit was 1, it overflowed), and replace 011 by 100.
2. To perform the increment, we substitute the lower two bits in the following way: $00 \rightarrow 01$, $01 \rightarrow 10$, $10 \rightarrow 11$.

The first (partial carry) step can be reduced to replacing $z11$ by 100 (where z represents either a 0 or a 1) as because of the adjacency constraint, z must always be 0. It can also be done in parallel with the second step, as we are not propagating the carries.

In the second step, we never need to substitute 11 in the lower two bits, because it will have been removed in the first step. The steps in this algorithm can be composed to yield Equation 6 and Figure 1. It is worth noting that each bit only depends on the value of 4 previous bits. Table 1 shows the increment sequence if a counter is started at $n = 0$, compared to the standard Fibonacci number system.

$$b'_{m-1}b'_{m-2} \cdots b'_1b'_0 = \sigma(b_{m-1}b_{m-2} \cdots b_1b_0) \text{ where}$$

$$b'_i \leftarrow \begin{cases} \overline{b_i b_{i-1}} \vee b_{i-1} b_{i-2} & i = m - 1 \\ \overline{b_{i+1} b_i} \overline{b_i b_{i-1}} \vee b_{i-1} b_{i-2} & 1 < i < m - 1 \\ \overline{b_2 b_1} \overline{b_0} \vee \overline{b_1} b_0 & i = 1 \\ b_1 \vee \overline{b_0} & i = 0 \end{cases} \quad (6)$$




Figure 1: Fibonacci counter

n	Standard	Counter	n	Standard	Counter
0	000000 _F	000000 _F	10	010010 _F	010010 _F
1	000001 _F	000001 _F	11	010100 _F	010011 _F
2	000010 _F	000010 _F	12	010101 _F	010101 _F
3	000100 _F	000011 _F	13	100000 _F	010110 _F
4	000101 _F	000101 _F	14	100001 _F	011001 _F
5	001000 _F	000110 _F	15	100010 _F	100010 _F
6	001001 _F	001001 _F	16	100100 _F	100011 _F
7	001010 _F	001010 _F	17	100101 _F	100101 _F
8	010000 _F	001011 _F	18	101000 _F	100110 _F
9	010001 _F	001101 _F	19	101001 _F	101001 _F

Table 1: Standard Fibonacci number system and modified Fibonacci increment sequence.

3 RADIX-2 COUNTERS

There has been much research into radix-2 counters [6, 5].

A *carry ripple counter* is the simple implementation of a counter where each clock cycle the carry has to propagate through each bit of the counter using a carry chain as shown in Figure 2. This requires very little logic but is very slow (delay $O(m)$ for an m bit counter).




Figure 2: Radix-2 counter with carry chain (clock and reset signals omitted)

A *carry lookahead counter* is a refinement of the carry ripple counter, where the full carry chain is avoided by using higher fan-in gates as shown in Figure 3 or a tree. The best case for using a tree brings the delay of an m -bit counter down to $O(\log m)$, but at the expense of increased logic — there are now $O(m \log m)$ AND gates. This technique is usually used for fast radix-2 adders [7].

A *constant time counter* [6] can be constructed that di-




Figure 3: Radix-2 counter with lookahead carry (clock and reset signals omitted). The multiple input AND gates could be replaced by a tree

vides a counter into chunks to amortise through all cycles the bursts of long yet rare carries, by using the carry from each block as the clock to the next block as shown in Figure 4. While the counter as a whole is synchronous, his has some signals propagating over multiple clock cycles, and makes setting or resetting the counter with a particular value a slow operation. This counter design is small and fast, requiring a similar amount of logic to a carry ripple counter, and having a propagation delay of no more than twice the propagation delay of a three input gate.




Figure 4: Constant time radix-2 counter. The carry out from each stage is ANDed with the clock signal to form the clock for the next stage

In comparison with this, Fibonacci counters have a similar size per bit to the carry ripple counter and constant time counter, and similar speed to the constant time counter. However as will be seen in Section 4.1, a Fibonacci counter requires more bits to store a particular range of values.

4 COUNTER EFFICIENCY

It is useful to have some measures for how efficient the various options for counters are. In particular, we consider *representation efficiency*, which can be used to determineshow many bits are required to store a particular value, and conversion to and from radix 2. Other efficiencies such as area per bit, propagation delay and power use depend on the implementation technology, so are only compared qualitatively.

4.1 Representation Efficiency

Representation efficiency is a measure of how well a value is ‘packed’ into the available bits.

We first define the number of states an m -bit counter traverses before overflowing or reaching a terminal value as $N(m)$. The states of the counter will typically be made to correspond with the integer values $\{0, 1 \dots, N(m)-1\}$. Some values for $N_{fib}(m)$ are shown in table 2.

Representation efficiency of a counter can then be defined as

$$R(m) = \sqrt[m]{N(m)} \quad (7)$$

This has a limit

$$R = \lim_{m \rightarrow \infty} R(m) \quad (8)$$

$$= \lim_{m \rightarrow \infty} \frac{N(m+1)}{N(m)} \text{ (if this limit exists)} \quad (9)$$

The representation efficiency of several counters are given in Figure 5. In particular, the representation efficiency of radix-2 is 2 (the best possible), and the representation efficiency of a Fibonacci counter is ϕ .




Figure 5: Representation efficiency for counter sizes up to 30 for the radix-2 counter (Section 3), Fibonacci counter (Section 2), Saturating Fibonacci Counter (Section 7.1) and the counters in Section 6

5 CONVERSION TO AND FROM RADIX 2

For some applications, it may be desirable to convert numbers in the Fibonacci number system to or from a radix-2 representation. Conversion from the Fibonacci number system to radix-2 can be performed using the following algorithm.

m	1	2	3	4	5	6	7	8	9	10	m
$N_{fib}(m)$	2	4	6	10	15	24	37	59	93	149	$\sim \frac{\phi^{m+2}}{\sqrt{5}} + \lfloor \frac{m}{2} \rfloor$

Table 2: Number of states of an m -bit Fibonacci Counter

```

fibonacci_to_radix2( $b_{m-1} \dots b_1 b_0$ )  $\mapsto$  radix-2
   $n \leftarrow 0$ 
   $c \leftarrow 1$ 
   $l \leftarrow 1$ 
  for  $i$  incrementing from 0 to  $m-1$ 
    if  $b_i = 1$ 
       $n \leftarrow n+c$ 
       $(c, l) \leftarrow (c+l, c)$ 
  return  $n$ ;

```

For an m -bit Fibonacci number, this will require either $2m$ radix-2 additions, or m radix-2 additions and a size m lookup table.

Conversion from an m -bin radix-2 to the Fibonacci system requires a m radix-2 additions, m radix-2 comparisons with zero, and either m radix-2 subtractions or a size m lookup table.

```

radix2_to_fibonacci( $n$ )  $\mapsto$  fibonacci
   $c \leftarrow F_{m+2}$ 
   $l \leftarrow F_{m+1}$ 
  for  $i$  decrementing from  $m-1$  to 0
     $d \leftarrow n-c$ 
    if  $d \geq 0$ 
       $b_i \leftarrow 1$ 
       $n \leftarrow d$ 
    else
       $b_i \leftarrow 0$ 
       $(c, l) \leftarrow (l, c-l)$ 
  return  $b_{m-1} \dots b_1 b_0$ ;

```

This suggests that using the Fibonacci number system for counters is only useful where conversions to or from radix-2 are either not required or are infrequent.

6 COMPARISON WITH OTHER COUNTERS

There are other sorts of counters that can be used where a fast counter is required. We discuss redundant binary counters and LFSRs.

6.1 Redundant Binary

Redundant number representations [12, 8] are widely used to for carry-free addition, and so make efficient counters. In redundant binary representation, each binary digit is a pair of bits, which commonly have the values $\{-1, 0, 1\}$ or $\{0, 1, 2\}$. The redundancy means that addition (therefore increment) can be performed carry free. The representation efficiency is poor however, requiring 2 bits for each extra binary digit, or $R = \sqrt{2}$. Also, converting to or from ordinary radix-2 representation is very easy.

6.2 Hybrid Redundant Binary

There are many hybrid schemes [8], but the one considered here is having the alternate binary digits stored in 1 or 2 bits (which we call 1-2-redundant binary), where the least significant digit requires 1 bit. This is still very fast, with no more than two levels of carry being required, and a better representation efficiency of $R = \sqrt[3]{4}$




Figure 6: Hybrid redundant binary counter

6.3 Linear Feedback Shift Registers

LFSRs can be used as extremely efficient counters [9], requiring very little logic for the increment operation. However the transformation from LFSR state to a corresponding radix-2 value is difficult. For small n this can be accomplished using a look-up table, but this becomes difficult for large n . Some work has been done on this problem [13], but large lookup tables are still required. They have very good representation efficiency (with a limit of 2). The average number of bit transitions per increment is proportional to the counter width m , whereas the average number of bit transitions for all the other counters is bounded by a constant, so for some technologies (such as CMOS), LFSRs may consume more power than other counters.

7 EXTENSIONS TO THE FIBONACCI COUNTER

7.1 Saturating Fibonacci Counter

A variation of the Fibonacci Counter is to only replace 011 by 100 if all those bit positions exist, so instead of $11zz \dots zz$ being replaced by $00zz \dots zz$, it would be left as $11zz \dots zz$. The values for a 5 bit saturating counter is shown in Table 3 As a result of this are that the adjacency constraint is no longer satisfied, so we replace it with a *saturating adjacency constraint* that requires that there are never more than two adjacent 1's if a higher bit

position contains a 0. This is equivalent to specifying that the sequence 0111 never occurs. A saturating counter “saturates” at its maximum value in that $11 \dots 11 \rightarrow 11 \dots 11$, so an m bit saturating counter has a larger maximum value than a non-saturating counter. The equations for this are shown below

$$b'_i \leftarrow \begin{cases} b_i \vee b_{i-1} \overline{b_{i-2}} & i = m-1 \\ \overline{b_{i+1} b_i \overline{b_{i-1}}} \vee \overline{b_i b_{i-1} b_{i-2}} \vee b_{i+1} b_i & i = m-2 \\ \overline{b_{i+1} b_i \overline{b_{i-1}}} \vee \overline{b_i b_{i-1} b_{i-2}} \vee b_{i+2} b_{i+1} b_i & 1 < i < m-2 \\ \overline{b_2 b_1 \overline{b_0}} \vee \overline{b_1 b_0} \vee b_3 b_2 b_1 & i = 1 \\ b_1 \vee \overline{b_0} & i = 0 \end{cases}$$

n	Counter	n	Counter	n	Counter
0	00000 _F	7	01010 _F	14	11001 _F
1	00001 _F	8	01011 _F	15	11010 _F
2	00010 _F	9	01101 _F	16	11011 _F
3	00011 _F	10	10010 _F	17	11101 _F
4	00101 _F	11	10011 _F	18	11110 _F
5	00110 _F	12	10101 _F	19	11111 _F
6	01001 _F	13	10110 _F		

Table 3: 5 bit Saturating Fibonacci Counter

For an m bit counter, the number of states $N_{sat}(m)$ that can be represented is shown in Equation 10 and Table 4, and is approximately a factor of ϕ better than a non-saturating counter, although they have the same representation efficiency. Typically an saturating Fibonacci counter requires one less bit than a non-saturating Fibonacci counter. A non-saturating Fibonacci counter requires more complex logic to implement than a saturating Fibonacci counter, with each bit being determined by the value of 5 of the previous bits.

$$\begin{aligned} N_{fib-sat}(m) &= \sum_{i=2}^{m+1} F_i + 1 \simeq \sum_{i=2}^{m+1} \frac{\phi^i}{\sqrt{5}} + 1 \\ &= \frac{\phi^{m+3} - \phi^3}{\sqrt{5}} + 1 \simeq \frac{\phi^{m+3}}{\sqrt{5}} - 1 \end{aligned}$$

7.2 k -Generalised Fibonacci Numbers

Instead of using the Fibonacci sequence, we could generalised Fibonacci numbers[14] such as Tribonacci F_n^3 or Tetranacci F_n^4 . The first few values of the Tribonacci sequence are 0, 0, 1, 1, 2, 4, 7, 13, 24, 44, 81, 149, . . . , where $F_n = F_{n-1} + F_{n-2} + F_{n-3}$ for $n > 2$.

Now the equivalent of Equation 3 becomes

$$n = b_{m-1} b_{m-2} \dots b_1 b_0 \text{ }_F^k = \sum_{i=0}^{m-1} b_k F_{i+k} \quad (10)$$

For a Tribonacci counter, the adjacency constraint is that there are no more than 3 ‘1’s in a row, and the representational efficiency is $R \simeq 1.839287$ (the positive root of $x^3 - x^2 - x - 1$), which is considerably better than for Fibonacci counters. The disadvantage is more complex logic to perform the increment — in general each bit depends on the state of 6 of the previous bits. The equations for increment are:

$$b'_i \leftarrow \begin{cases} b_i \overline{b_{i-1}} \vee b_i \overline{b_{i-2}} \vee b_{i-1} b_{i-2} b_{i-3} & i = m-1 \\ \overline{b_i b_{i-1}} \vee \overline{b_{i+1} b_i b_{i-2}} \vee b_{i-1} b_{i-2} b_{i-3} & i = m-2 \\ \overline{b_{i+2} b_i b_{i-1}} \vee \overline{b_{i+1} b_i b_{i-1}} \vee \\ \quad \overline{b_{i+1} b_i b_{i-2}} \vee b_{i-1} b_{i-2} b_{i-3} & 2 < i < m-2 \\ \overline{b_4 b_2 b_1} \vee \overline{b_3 b_2 b_1} \vee \overline{b_3 b_2 b_0} \vee \overline{b_2 b_1 b_0} & i = 2 \\ \overline{b_3 b_1 b_0} \vee \overline{b_2 b_1 b_0} \vee \overline{b_1 b_0} & i = 1 \\ b_2 b_1 \vee \overline{b_0} & i = 0 \end{cases}$$

Otherwise the advantages and disadvantages are similar to those of Fibonacci counters.

8 CONCLUSION

Fibonacci counters offer a counter option where speed is important and representation redundancy is not an issue. These counters compare well with redundant radix-2 counters, the only disadvantage being that conversion to or from radix-2 representation is more expensive. They also offer a saturating counter with only a little extra logic, and related counters such as Tribonacci counters may also offer advantages.

8.1 Applications of Fibonacci Counters

Applications of Fibonacci counters depend on representation redundancy not being an issue, so for instance they are no use at addressing an array in memory. However, some applications might include:

- Event counters. These are often required to be very fast, and are widely used in power management. Other examples might be counting fast physical events [4]. A consideration here is that ‘readout’ (converting to radix-2) must not be a time critical operation.
- Timers. Counting a particular number of cycles is a good use for these counters, as they are fast and such applications typically don’t require a ‘readout’ stage. They have an advantage over LFSRs, as with the appropriate initial settings they can be set up so that only a single bit needs to be checked for expiry of the timer.

m	1	2	3	4	5	6	7	8	9	10	m
$N_{fib-sat}(m)$	2	4	7	12	20	33	54	88	143	232	$\sim \frac{\phi^{m+3}}{\sqrt{5}} - 1$

Table 4: Number of states for an m -bit Saturating Fibonacci Counter

Counter	Area per Bit	Speed	Representation Efficiency	Radix-2 Conversion
Radix-2 (carry ripple)	small	slow	2	-
Radix-2 (carry lookahead)	large	moderate	2	-
Radix-2 (constant time) [†]	small	fast	2	-
Fibonacci	small	fast	$\phi \simeq 1.6180$	moderate
Saturating Fibonacci	small	fast	$\phi \simeq 1.6180$	moderate
Redundant Binary	small	fast	$\sqrt{2} \simeq 1.4142$	fast
1-2 redundant binary	small	fast	$\sqrt[3]{4} \simeq 1.5874$	fast
LFSR	very small	very fast	2	very slow

[†] These counters have internal propagating carries, so are slow to set to a new value

Table 5: Comparison of Fibonacci and other counters

REFERENCES

- [1] R. Joseph and M. Martonosi, "Run-time power estimation in high performance microprocessors," in *Proceedings of the 2001 international symposium on Low power electronics and design*, pp. 135–140, ACM New York, NY, USA, 2001.
- [2] S. Chung and K. Skadron, "Using On-Chip Event Counters For High-Resolution, Real-Time Temperature Measurement," *Thermal and Thermomechanical Phenomena in Electronics Systems*, pp. 114–120, 2006.
- [3] Y. Ruan, V. Pai, E. Nahum, and J. Tracey, "Evaluating the impact of simultaneous multithreading on network servers using real hardware," *ACM SIGMETRICS Performance Evaluation Review*, vol. 33, no. 1, pp. 315–326, 2005.
- [4] J. Christiansen, A. Marchioro, and P. Moreira, "TTCrx, an ASIC for timing, trigger and control distribution in LHC experiments," in *Given at 2nd Workshop on Electronics for LHC Experiments, Balatonfured, Hungary*, pp. 23–27, 1996.
- [5] M. Stan, A. Tenca, and M. Ercegovac, "Long and Fast Up/Down Counters," *IEEE Transactions on Computers*, vol. 47, no. 7, pp. 722–735, 1998.
- [6] J. Vuillemin, "Constant time arbitrary length synchronous binary counters," in *10th IEEE Symposium on Computer Arithmetic, 1991. Proceedings.*, pp. 180–183, 1991.
- [7] Y. Pai and Y. Chen, "The fastest carry lookahead adder," in *Proceedings of the Second IEEE International Workshop on Electronic Design, Test and Applications (DELTA'04)*, pp. 434–436.
- [8] B. Parhami, *Computer arithmetic: algorithms and hardware designs*. Oxford University Press Oxford, UK, 1999.
- [9] M. Robshaw, "Stream Ciphers," *RSA Laboratories*, vol. 25, 1995.
- [10] E. Zeckendorf, "Representation des nombres naturels par une somme de nombres de Fibonacci ou de nombres de Lucas," *Bull. Soc. Roy. Sci. Liege*, vol. 41, pp. 179–182, 1972.
- [11] R. Graham, D. Knuth, and O. Patashnik, *Concrete mathematics: A Foundation for computer Science*. Addison-Wesley Reading, Mass, 2nd ed., 1994.
- [12] N. Takagi, H. Yasuura, and S. Yajima, "High-Speed VLSI Multiplication Algorithm with a Redundant Binary Addition Tree," *IEEE Transactions on Computers*, vol. 34, no. 9, pp. 789–796, 1985.
- [13] D. Clark and L. Weng, "Maximal and near-maximal shift register sequences: efficient event-counters and easy discrete logarithms," *Computers, IEEE Transactions on*, vol. 43, no. 5, pp. 560–568, 1994.
- [14] E. Miles Jr, "Generalized Fibonacci Numbers and Associated Matrices," *The American Mathematical Monthly*, vol. 67, no. 8, pp. 745–752, 1960.

Differential Evolution and its Application to Intelligent Spectral Design

S Soltic & A N Chalmers

EET School
Manukau Institute of Technology
Private bag 94-006, Manukau City 2240, NEW ZEALAND
Email: ssoltic@manukau.ac.nz, chalmers@manulauac.nz

Abstract:

This paper joins the quest for an efficient white-light source based on mixtures of multiple narrow-band monochromatic light emitting diode (LED) spectra optimised for high luminous efficacy and colour rendering properties. The complexity of the optimization process increases with the number of LEDs used to create a white-light source. Hence, this paper proposes an effective approach for mixing a number of narrow-band LED spectra based on differential evolution (DE), a simple and easy to use direct search method. A software program was designed to implement the DE algorithm. Optimal solutions for four, five and six different LEDs are analysed. Further research is underway to investigate white-light sources based on theoretical LED spectra and to construct physical implementations of practical LED-based white-light sources.

Keywords:

Optimization, Differential evolution, Spectral design, LED lighting, Colour rendering, Radiant luminous efficacy

1 INTRODUCTION

White LEDs have been labelled as the eco-friendly light source of the 21st century [1]. With LEDs white light can be produced either by using blue light to irradiate a yellow-emitting phosphor [2], [3], or by combining lights of multiple monochromatic LEDs [1], [4], [5]. The latter approach results in white LEDs with higher efficiency than the former approach since there is no loss of energy when down-converting [5], [6]. Both approaches aim to design LED-based white light sources characterized by high colour rendering properties and high luminous efficacy. Unfortunately, these two quantities are generally contravariant. Both colour rendering and luminous efficacy depend on the emitted spectrum of the light source. By fine tuning the spectrum intensity of the individual LEDs we are able to find a balance between the two properties of the light source. Typically, the optimization follows a trial-and-error search inside the n -dimensional space $\{I_1, \dots, I_i, \dots, I_n\}$ until a satisfactory spectral power distribution (SPD) of the LED-based white source is found [5], [7], [8].

Here we propose an approach to the optimization of SPD mixtures based on differential evolution (DE), a simple global optimization algorithm proposed by Reiner Storn and Kenneth Price in 1995 [9]. We optimise the SPDs of the mixtures of 3, 4, 5, 6 and 7 LEDs chosen from the Luxeon range [10]. For this purpose, a MATLAB program was developed that optimises the SPDs of an arbitrary number of LEDs.

This paper proceeds as follows. In Section 2, the principles of the DE algorithm are described in detail. In Section 3, the theory of colour rendering and luminous efficacy is given. Section 4 presents the experimental results and comparisons and, finally, suggestions for future work and conclusions are given in Section 5.

2 DIFFERENTIAL EVOLUTION

Differential evolution is a powerful population based evolutionary algorithm suitable for optimization of real-value multi-modal nonlinear and non differentiable objective functions $f_0(x_1, x_2, \dots, x_n)$ [9], [11]. DE is simple and has proven to be powerful in solving a number of benchmark problems [12][13].

The search for an optimal solution starts with a population of P randomly created solution vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_P\}$. The value of P is kept constant during the optimization process. The solution vectors undergo mutation, crossover, evaluation and selection over a number of generations G . Both the population size P and the number of generations G depend on the problem to be optimised. Storn and Price suggest $P \in [5 \times n, 10 \times n]$, n is the dimension of the objective function, but P must be ≥ 4 to provide enough mutually different solution vectors for the algorithm to function properly [11].

Mutation is the process of creating a new offspring vector $\mathbf{u}_{i,G+1}$ by adding the weighted difference between two

randomly chosen solution vectors, $\mathbf{v}_{r2,G}$ and $\mathbf{v}_{r3,G}$ to a randomly chosen $\mathbf{v}_{r1,G}$ [11]:

$$\mathbf{u}_{i,G+1} = \mathbf{v}_{r1,G} + \mathbf{F} \times (\mathbf{v}_{r2,G} - \mathbf{v}_{r3,G}) \quad (1)$$

where \mathbf{F} is a mutation weight $\in [0,2]$ and random indices $r1, r2, r3 \in [0, P-1]$ are chosen to be different from the index i . The mutation process is illustrated in Fig. 1.




Figure 1: Illustration of the mutation process for $P = 9$.

The offspring solution vectors $\mathbf{u}_{i,G+1}$ undergo crossover ensuring that offspring vectors differ from their parents [10]:

$$\mathbf{w}_{ji, G+1} = \begin{cases} \mathbf{u}_{ji, G+1} & \text{if } (j \leq \text{CR}) \mid (j = i) \\ \mathbf{v}_{ji, G} & \text{if } (j > \text{CR}) \mid (j \neq i) \end{cases} \quad (2)$$

where $i \in \{1,2,\dots, P\}$ is a randomly chosen integer, $j \in [0, 1)$ is a randomly chosen real value and $\text{CR} \in [0, 1]$ is a crossover constant influencing the number of elements to be exchanged. The crossover process is illustrated in Fig. 2.




Figure 2: Illustration of the crossover process, $n = 6$.

All offspring vectors $\{\mathbf{w}_{1,G+1}, \dots, \mathbf{w}_{P,G+1}\}$ are compared to their parent vectors $\{\mathbf{v}_{1,G}, \dots, \mathbf{v}_{P,G}\}$. If the $\mathbf{w}_{i,G+1}$ vector has a larger goodness function f_{good} , than its parent vector $\mathbf{v}_{i,G}$, then it becomes a member of generation $G+1$, otherwise the $\mathbf{v}_{i,G}$ vector is retained. Hence, only the fitter offspring vectors are moved to the $G+1$ generation.

Since 1995, different variants on the original DE algorithm have been proposed [11]. The main differences are in the creation of new solution vectors where:

1. the best solution vector from the current generation $\mathbf{v}_{\text{best},G}$ is mutated rather than $\mathbf{v}_{r1,G}$;
2. more than two difference vectors are used in mutation; and
3. different crossover schemes are employed.

Here, we used the original DE described in pseudo-code given below:

```

Require:  $G, P, F \in [0,2], \text{CR} \in [0,1]$ 
1: Create vectors  $\{\mathbf{v}_{1,0}, \dots, \mathbf{v}_{P,0}\}$ 
2: Evaluate  $\{\mathbf{v}_{1,0}, \dots, \mathbf{v}_{P,0}\} \rightarrow \mathbf{f}_{\text{good},0}$ 
3: for all  $G$  do
4:   Mutate  $\mathbf{v}_{i,G} \rightarrow \mathbf{u}_{i,G+1}$ 
5:   Crossover  $\{\mathbf{u}_{1,G+1}, \dots, \mathbf{u}_{P,G+1}\}$  and
      $\{\mathbf{v}_{1,G}, \dots, \mathbf{v}_{P,G}\} \rightarrow \{\mathbf{w}_{1,G+1}, \dots, \mathbf{w}_{P,G+1}\}$ 
6:   Evaluate  $\{\mathbf{w}_{1,G+1}, \dots, \mathbf{w}_{P,G+1}\}$ 
7:   if  $f_{\text{good}}(\mathbf{w}_{i,G+1}) > f_{\text{good}}(\mathbf{v}_{i,G})$ 
8:      $\mathbf{v}_{i,G+1} = \mathbf{w}_{i,G+1}$ 
9:   else
10:     $\mathbf{v}_{i,G+1} = \mathbf{v}_{i,G}$ 
11:   end if
12: end for
    
```

3 COLOUR RENDERING AND LUMINOUS EFFICACY

3.1 Colour rendering indices

Colour rendering is the characteristic of light sources that describes the effect of a light source on the colour appearance of an object. Typically, the colour rendering properties of light sources are evaluated using the CRI colour rendering index CRI [14]. The CRI values range from 0 to 100. The higher the CRI value of a light source the more “natural” the colours of objects look under this source.

The fundamental idea behind the CRI calculation is a comparison of the colours of eight test colour samples illuminated by the test and reference light sources (ΔE_i). The reference source has to be a source with the same correlated colour temperature (CCT) as the test source. The special colour rendering index R_i for each test colour sample is calculated as:

$$R_i = 100 - 4.6\Delta E_i, \quad i = 1, 2, \dots, 8. \quad (3)$$

Averaging those eight R_i values results in the general rendering index R_a . The reference illuminant is either a Planckian radiator for test sources having $\text{CCT} < 5000\text{K}$ or a phase of daylight for test sources having $\text{CCT} \geq 5000\text{K}$. In an attempt to improve the descriptive power of CRI, the number of test samples was extended to 14 by addition of six new test samples representing red, yellow, green and blue saturated colours, human complexion and leaf green. Averaging the additional six R_i values results in what we term index R_b , and averaging all 14 R_i values results in an “overall” index R_c . The acceptable value of CRI depends on the application of the light source. The applications where excellent rendering is important require the utilization of lamps with high CRIs. CRI values above 70 are considered reasonably good and values above 85 are believed to be excellent. But having a lamp with a high CRI does not guarantee

that all colours will be properly rendered. Yoshi Ohno [4] observed in his experiments that two 3-band light sources, both with $R_a = 80$, rendered the red test sample quite differently. Therefore, we report our results in terms of the smallest colour rendering index R_i (as suggested by Cuttle [15]) as well as in terms of R_a , R_b and R_c .

3.2 Radiant luminous efficacy

We rate the performance of our test sources in terms of radiant luminous efficacy, η_{rad} lm/W, which compares the amounts of luminous flux and radiant flux emitted by the source:

$$\eta_{rad} = \frac{683 \times \int_{\lambda} V(\lambda) S(\lambda) d\lambda}{\int_{\lambda} S(\lambda) d\lambda} \quad (4)$$

where $S(\lambda)$ is the spectral distribution of the light source and $V(\lambda)$ is the CIE spectral sensitivity function for photopic vision. The efficacy of the light sources is an important feature and we use η_{rad} as a term in the goodness functions in all our experiments.

4 RESULTS

An optimization program has been developed to perform DE and to provide a graphical visualization the SPD curve of the optimal LED combination (Fig. 3). The program calculates and displays the CRI indices R_a , R_b and R_c , the correlated colour temperature (CCT, in Kelvins), η_{rad} , the smallest R_i value and its index i_{min} to identify the poorest rendered test colour sample. The results, including the SPD values of the optimal solution are saved in a file for further analysis. The DE parameters, G, P, F and CR, can be entered by the user. The number of LEDs (N) to be optimised is automatically detected from the input data file.




Figure 3: The user interface of the optimization program.

The relative spectral power distributions of the individual LEDs with peak wavelengths at $\lambda_{rb} = 450$ nm (royal-blue), $\lambda_b = 460$ nm (blue), $\lambda_c = 505$ nm (cyan), $\lambda_g = 525$ nm (green), $\lambda_a = 590$ nm (amber), $\lambda_{ro} = 630$ nm (red-orange) and $\lambda_r = 640$ nm (red) used in this work are shown in Fig. 4.




Figure 4: The spectral power distributions of the seven LEDs used in our experiments. [10]

In order to optimize the LED mixtures, we experimented with different goodness functions:

$$f_{good} = a \times R_a + b \times R_b + c \times R_c + d \times \eta_{rad} + e \times R_{min} \quad (5)$$

where a, b, c, d and e are weights controlling the influence of R_a , R_b , R_c , η_{rad} and R_{min} on the optimization of the LED mixtures. Figure 5 shows an example of the f_{good} values in one of our experiments. It can be seen that f_{good} converges around the 1000th generation and exhibits only a few small changes in the next 1000 generations. Hence, the number of generations G was set to 1000 in all of our experiments for faster and still accurate convergence of DE. If G is too small, DE has no time to explore a sufficient number of solution vectors. However, for values of G that are too large the optimization process is unnecessarily slowed down without any optimization gain.

We also experimented with the DE parameters P, F and CR. The values of these parameters had a minor influence on the optimization results. According to our experience choosing the DE's control variables in order to get good optimization was easy. After some experimentation we set $P = 50$ and F and CR to the values suggested in [11], i.e. $F = 0.5$ and $CR = 0.1$.




Figure 5: Goodness function f_{good} in one of the experiments, $N = 4$ (blue, green, amber, red).

It is now recognised that four individual LEDs (blue, green, amber and red) covering the visible range must be used to achieve a satisfactory CRI [6]. Hence we optimised a 4-band mixture of the blue, green, amber and red LEDs. The optimization based on $f_{good1} = R_a + R_b + R_c +$

$\eta_{\text{rad}} + 10 \times R_{\text{min}}$ resulted in a light source with $R_a = 95$, $R_b = 86$, $\eta_{\text{rad}} = 339$ lm/W, CCT = 3268K (Table 1). The blue test sample was the poorest rendered of all samples with $R_{12} = 76$. The SPD of this particular LED combination is shown in Fig. 6. Lowering the influence of R_{min} on f_{good} ($e = 5$) resulted in slightly lower CRI values, $R_a = 93$, $R_b = 83$ and $R_{12} = 76$ (the blue test sample) but a higher efficacy $\eta_{\text{rad}} = 350$ lm/W. Removing R_{min} from the optimization ($e = 0$) and lowering the influence of η_{rad} ($d = 0.5$) resulted in an optimised mixture with a satisfactory $R_a = 92$, $R_b = 79$ but a much lower $R_9 = 58$ (the red test sample). These results highlight the importance of f_{good} and the weights a, b, c, d and e on the optimization of the LED mixtures. We found that having $f_{\text{good}} = f(R_{\text{min}})$ is important in the creation of a mixture with a satisfactory R_{min} value. Particularly, the colour rendering properties of the light sources optimised with R_{min} were better than of those optimised without R_{min} . Hence, according to our results the lowest R_i value should be somehow considered in the optimization process of the LED-based light sources to ensure higher CRI indices.




Figure 6: Relative SPD of the 4-band LED mixture (blue $\lambda_b = 460$ nm, green $\lambda_g = 530$ nm, amber $\lambda_a = 590$ nm and red $\lambda_r = 630$ nm), $R_a = 95$, $R_b = 86$, $\eta_{\text{rad}} = 339$ lm/W, CCT = 3268K, $R_{12} = 76$ (the blue sample).

Figure 7 and 8 show how the individual elements of f_{good} were fluctuating before converging to their final values. It can be seen that η_{rad} (Fig. 8) influences f_{good} considerably at the beginning of the optimization process ($G < 300$) before converging to just below 340 lm/W in the 336th generation. In this particular case, the population converges around $G = 400$ and it is only fine tuned in the next 600 generations.

We also observed minor performance changes from the LED-based light sources caused by an introduction of a 5th and 6th LED.

First, we introduced the red-orange LED ($\lambda_{\text{ro}} = 630$ nm)

and optimised the mixture (Fig. 9). The introduction of the red-orange LED did not have a significant effect on the overall characteristics of the light source: $R_a = 95$ (no change), $R_b = 85$ (down 1 unit), $\eta_{\text{rad}} = 343$ lm/W (up 4 units), CCT = 3269K (up 1 unit) and $R_9 = 75$ (the red test sample).




Figure 7: Behaviour of R_a, R_b, R_c and R_{min} while searching for the SPD shown in Fig. 6.




Figure 8: Behaviour of η_{rad} while searching for the SPD shown in Fig. 6.




Figure 9: SPD of the optimised mixture of the blue $\lambda_b = 460$ nm, green $\lambda_g = 530$ nm, amber $\lambda_a = 590$ nm, red-orange $\lambda_{\text{ro}} = 630$ nm and red $\lambda_r = 630$ nm LEDs: $R_a = 95$, $R_b = 85$, $\eta_{\text{rad}} = 343$ lm/W, CCT = 3269K, $R_9 = 75$ (the red test sample).

Table 1: Influence of different f_{good} on the optimization of LED mixtures.

Test no	Goodness weighting factors					Best results						
	a	b	c	d	e	R_a	R_b	R_c	η_{rad} (lm/W)	R_{min}	i_{min}	CCT (K)
1	1	1	1	1	10	95	96	91	339	76	12	3268
2	1	1	1	1	5	93	83	89	350	76	12	3074
3	1	1	1	0.5	0	92	79	87	384	58	9	3169

Second, the spectrum was broadened using the royal-blue LED with peak wavelength at $\lambda_{rb} = 450$ nm (Fig. 10). This resulted in a small improvement in R_a yielding: $R_a = 96$ (up 1 unit), $R_b = 86$ (no change), $\eta_{rad} = 338$ lm/W (down 1 unit), CCT = 3356K and $R_{12} = 75$.

Third, the combined influence of the royal-blue and red-orange LEDs is shown in Fig. 11 (the dotted lines): $R_a = 96$, $R_b = 86$, $\eta_{rad} = 338$ lm/W, CCT = 3356K, $R_{12} = 75$ (the blue test sample).

The results of the investigations with $N \geq 5$ are summarised in Table 2.




Figure 10: SPD of the optimised mixture of the royal-blue $\lambda_{rb} = 450$ nm, blue $\lambda_b = 460$ nm, green $\lambda_g = 530$ nm, amber $\lambda_a = 590$ nm and red $\lambda_r = 630$ nm LEDs: $R_a = 96$, $R_b = 86$, $\eta_{rad} = 338$ lm/W, CCT = 3356K, $R_{12} = 75$ (the blue test sample).




Figure 11: The 6-band light source with the royal-blue $\lambda_{rb} = 450$ nm, blue $\lambda_b = 460$ nm, green $\lambda_g = 530$ nm, amber $\lambda_a = 590$ nm, red-orange $\lambda_{ro} = 630$ nm and red $\lambda_r = 630$ nm LEDs: $R_a = 96$, $R_b = 86$, $\eta_{rad} = 338$ lm/W, CCT = 3356K, $R_{12} = 75$ (the blue test sample).

Table 2: Comparison of optimised mixtures with $N \geq 5$.

N λ	R_a	R_b	R_c	η_{rad} (lm/W)	R_{min}	i_{min}	CCT (K)
5 λ_{ro}	95	85	91	343	75	9	3269
5 λ_{rb}	96	86	91	338	75	12	3356
6 $\lambda_{ro}, \lambda_{rb}$	95	86	91	334	75	12	3541

5 CONCLUSION AND FUTURE WORK

We have shown that LED-based white sources can be optimised using an evolutionary computation method called differential evolution. This method is easy to understand and implement. According to our experience the DE's parameters, population size P, multiplication weight F and crossover constant CR are easy to choose. In fact we could not report any problems encountered with the use of DE.

However, the optimization of the LED mixtures was highly dependent on goodness function f_{good} . The ultimate goal of creating a light source with high-luminous-efficacy and high-colour-rendering is to find the optimum trade-off between these two properties. Therefore, we introduced a goodness function based on the value of the worst rendered colour test sample. Our white-light sources have comparable or better properties than those reported in much of the literature.

Two aspects of our results were somewhat unexpected and warrant further investigation in the future:

- In virtually all cases, the optimum spectrum turned out to have a CCT in the range $3000K \leq CCT \leq 4000K$.
- Decreasing the weighting factor d (for η_{rad}) from 1.0 to 0.5 led to our highest η_{rad} result and lowest R_{min} result.

Finding (a) was surprising in view of the fact that CCT was not used as an element of our goodness function f_{good} . The fact that the CCT values were all in a popular range of source colours for general lighting was an unintended consequence of the optimization process. The reasons require fuller investigation as part of our future work.

Finding (b) appears, on the face of it, to be counterintuitive, and will also be investigated further.

Given that narrow-band LED mixtures could be optimised to create a highly efficient light source with excellent colour rendering properties, it is unrealistic to expect that even better properties can be achieved with the currently available off-the-shelf LEDs. Hence, further research is underway to investigate white-light sources based on various theoretical LED spectra. In [5] the LED spectra were simulated using Gaussian curves with width at half magnitude (λ_{di}) of 30 nm. Our intentions are (1) to explore different bandwidths as well as a variety of shapes, such as triangular, rectangular or sine-square; and (2) to optimise the peak wavelengths λ_i , bandwidths λ_{di} and intensities I_i employing the differential evolution algorithm.

Moreover, further research is underway to investigate the practical implementation of multi-LED white-light sources.

6 REFERENCES

- [1] Z. Lei, G. Xia, L. Ting, G. Xiaoling, L. Q. Ming & S. Guangdi, "Color rendering and luminous efficacy of trichromatic and tetrachromatic LED-based white LEDs", *Microelectronics Journal*, volume 38, pp. 1-6, 2007.
- [2] R. Stevenson, "The LED's dark secret", *IEEE Spectrum*, August 2009, pp. 23-27.
- [3] A. Bergh, G. Craford, A. Duggal & R. Haitz, "The promise and challenge of solid-state lighting", *Physics Today*, December 2001, pp. 42-47.
- [4] Y. Ohno, "Simulation analysis of white LED spectra and color rendering", *CIE Expert Symposium on LED Light Sources*, Tokyo, 2004.
- [5] A. Žakauskas, R. Vaicekauskas, F. Ivanauskas, R. Gaska & S. Shur, "Optimization of white polychromatic semiconductor lamps", *Applied Physics Letters*, volume 80, number 2, pp. 234-236, 2002.
- [6] J. B. Protzman & K. W. Houser, "LEDs for general illumination: the state of the science", *Leucos*, volume 3, number 2, pp. 121-142, 2006.
- [7] A. Chalmers, C. Cuttle, L. Wang & P. Luong, "The quest for high-rendering high-luminous-efficacy light sources", *IESANZ Convention 2008*, Queenstown, New Zealand.
- [8] Y. Lu, Y. Gao, H. Chen & Z. Chen, "Intelligent spectra design and colorimetric parameter analysis for light-emitting diodes", *14th International Conference on Mechatronics and Machine Vision in Practice, M2VIP 2007*, 4-6 Dec. 2007, pp. 118-122.
- [9] R. Storn & K. Price, "Differential evolution – a simple and efficient adaptive scheme for global optimization over continuous spaces", *Technical report TR-95-012*, International Computer Science Institute, Berkeley, 1995.
- [10] Lumileds Lighting, Luxeon® K2 Emitter, *Technical Datasheet DS51*, Lumileds Lighting US LLC., San Jose, California, 2006.
- [11] R. Storn & K. Price, "Differential evolution – a simple and efficient heuristic for global optimization over continuous space", *Journal of Global Optimization*, volume 11, pp. 341-359, 1997.
- [12] Karaboğa & s. Ökdem, "A simple and global optimization algorithm for engineering problems: differential evolution algorithm", *Turkey Journal of Electrical Engineering*, volume 12, number 1, pp. 53-60, 2004.
- [13] J. Vesterstrøm & R. Thomsen, "A comparative study of differential evolution, particle swarm optimization, and evolutionary algorithms on numerical bench problems", *Congress on Evolutionary Computation*, 19-23 June 2004, volume 2, pp. 1980-1987.
- [14] CIE, "Method of measuring and specifying colour rendering properties of light sources", *Technical report 13.3-1995*, CIE Central Bureau, Vienna.
- [15] C. Cuttle, Private communication, 2009.

DESIGN OF AN IMPROVED FUZZY LOGIC CONTROLLER MICRO-CHIP FOR WASHING MACHINE

¹Pritesh Lohani and ²S.M. Rezaul Hasan

Center for Research in Analog and VLSI Microsystems dESign (CRAVE)
School of Engineering and Advanced Technology, Massey University
Albany, Auckland, New Zealand

¹ P.Lohani@massey.ac.nz, ² hasanmic@massey.ac.nz

Abstract: Washing machines are today a common household requirement. Wash time is one of the key factors that need to be taken into account in designing a washing machine. Washing machines which do not use fuzzy controller serves most of the purposes of washing function but wash time is somewhat not dealt properly. This paper describes the architectural design of an improved washing machine controller that uses fuzzy logic approach to specifically deal with the wash time in a much more efficient manner. Recent research shows that fuzzy logic approach responds much faster than any other conventional technique. This design of a fuzzy logic controller for Washing Machine has 37 I/O pins including two V_{dd} and two V_{ss}. On chip Fuzzification, Fuzzy Inference Engine, Defuzzification, ROM based Fuzzy sets and MIN-MAX array based Fuzzy rules implementation are the salient features of the design. With full CMOS compatible host interface it is suitable as a co-processor.

Keywords: fuzzification, defuzzification, wash time


1 INTRODUCTION

Washing machines are common household item and to have a washing machine that efficiently controls the wash time is vital. Conventional proportional, integral and differential (PID) controllers have proven to be less capable in such control situations. In recent years there has been a growing interest in applying fuzzy logic for control. Fuzzy logic is also easier to implement as problems such as this one could be solved using heuristics based on human “rule of thumb” approach incorporating a higher level of abstraction originating from human knowledge and experience. One principal feature of this “rule of thumb” is linguistic terms (concepts) which can be dealt with approximate reasoning. These linguistic terms (concepts) are fuzzy sets and the approximate reasoning is based on fuzzy logic [1]. Considering the high speed capabilities available through custom hardware design of fuzzy controllers, the approach is to design the improved fuzzy logic controller chip using CMOS technology.


2 THEORY

Three fundamental sections of Fuzzy logic principles are:

(a) Fuzzification of crisp input signals: This can usually be done in two different ways (i) Impulse Fuzzification and (ii) Triangular Fuzzification. Figure 1 shows the fuzzification of discrete crisp inputs using both ways.



(a) Impulse Fuzzification



(b) Triangular Fuzzification

Figure 1: Fuzzification of discrete crisp inputs

(b) Use of fuzzy sets and fuzzy rules for fuzzy inference: Fuzzy inference consists of two entities which are 'premise' and 'conclusion'. Both premise and conclusion consists of fuzzy logical operators and fuzzy linguistic variables (which could be either noun or a pronoun). The premise is connected to the conclusion by fuzzy conditional rules which are of the form: IF (premise) THEN (conclusion) [2]. The premise is made up of a statement which consists of the fuzzy predicates P_{ijk} (each of which is an adjective of the Linguistic variable) of the respective Linguistic Variable that are combined by different fuzzy logic operators such as AND, OR. The linguistic variables are defined in the Universe of Discourse. The following is an example of the fuzzy conditional rule using such operators:

IF (P_{111} AND P_{221} OR P_{321}) THEN (P_{411}) where, P_{111} AND P_{221} OR P_{321} are the input predicates and P_{411} is the output predicate for the first rule.

The predicates are of the form: $P_{ijk} = (LV_i \text{ is } LV_i_ADJ_j \text{ for RULE_k})$ where LV_i is the i^{th} Linguistic variable and $LV_i_ADJ_j$ is the j^{th} Adjective of the i^{th} Linguistic Variable for the k^{th} rule. The premise defines the conditions in which conclusions are to be applied; the conclusions define actions to be taken when conditions of premise are satisfied.

The Adjectives for the Linguistic Variable is expressed in terms of fuzzy subsets in a specific fuzzy set. Membership of fuzzy subsets is expressed in terms of degree of membership (varying from 0 to 10) of each element in the Universe of Discourse. The envelope of a fuzzy subset expressing the membership of different crisp values within the range of an adjective is known as membership function [3].

Fuzzy inference is the intermediate process in obtaining crisp output values [3]. MIN-MAX method is one of the well known techniques for fuzzy inference. The membership degree α_{ijk} of inputs of input predicates P_{ijk} is determined by an intersection between the fuzzified LV_i and $LV_i_ADJ_j$, and choosing the maximum degree of membership. This is the MAX part of the MAX-MIN inference technique. Next the degree of membership of the entire premise for the k^{th} rule θ_k is determined by completing the premise fuzzy logical operations (AND, OR) using the α_{ijk} 's which is usually the minimum value among the α_{ijk} 's of the premise. This is the MIN part of the MIN-MAX inference rule.

θ_k determines the firing strength of the premise into the conclusion part. The output Linguistic variable will have their Adjectives (given in the conclusion part of the k^{th}

fuzzy rule) truncated to a membership grade of θ_k of the premise.

There are a sizable number of rules on Fuzzy control, in the form:

RULE1: IF (Premise_1) THEN (Conclusion_1)

RULE2: IF (Premise_2) THEN (Conclusion_2)

.....

.....

RULEn: IF (Premise_n) THEN (Conclusion_n)

Crisp values for the linguistic output variables are then obtained from their membership-truncated adjectives (fuzzy subsets) derived from the MIN-MAX fuzzy inference process [3]. Fuzzy union (OR) is carried out between the membership-truncated fuzzy subsets representing the adjectives for the same output variable for all the rules using the MAX operator. This result is the final membership function and fuzzy subset representing the fuzzified output variable.

(c) Defuzzification and generation of crisp output signal: Defuzzification is finally carried out to obtain a crisp output value using a very popular technique known as the 'centre of gravity' method which is given by:

$$C_{out} = (\sum i.m_i) / (\sum m_i)$$

Where, m_i is the membership value of the i^{th} crisp output value for the fuzzy subset representing the fuzzified output.

3 HARDWARE DESIGN AND SYSTEM OPERATION

Keeping the above theory in mind we now define the fuzzy logic controller for wash time in washing machine. The linguistic inputs that determines the wash time are: **dirtiness of the clothes (Large, Small, Medium), type of dirt (Greasy, Medium, Not Greasy) and mass of the clothes (Heavy, Medium, Light)**. Figure 2 shows the basic overview of the problem. The fuzzy controller takes three inputs, processes the information and outputs the wash time. Working of the sensors to get these inputs is beyond the scope of this discussion.




Figure 2: Hardware system overview of a fuzzy wash_time controller

Prior to dealing with the details of the fuzzy controller, possible values for the input and output variables are determined. These values are the membership functions that are used to map the real world measurement values to the fuzzy values so that the operations could be applied to them.

Fuzzy rules are derived from common sense. The set of rules that are used here to derive the output are shown in Table 1.

Table 1: Rules for Fuzzy Wash Time control

RULE NUMBER	LINGUISTIC INPUTS			LINGUISTIC OUTPUT
	Dirtiness_of_clothes	Dirt_type	Mass_of_the_clothes	Wash_Time
1	<i>Large</i>	<i>Greasy</i>	<i>Heavy</i>	<i>Very_High</i>
2	<i>Large</i>	<i>Greasy</i>	<i>Medium</i>	<i>High</i>
3	<i>Large</i>	<i>Medium</i>	<i>Heavy</i>	<i>High</i>
4	<i>Large</i>	<i>Greasy</i>	<i>Light</i>	<i>Medium</i>
5	<i>Large</i>	<i>Not greasy</i>	<i>Heavy</i>	<i>High</i>
6	<i>Large</i>	<i>Not greasy</i>	<i>Light</i>	<i>Medium</i>
7	<i>Large</i>	<i>Medium</i>	<i>Medium</i>	<i>Medium</i>
8	<i>Large</i>	<i>Medium</i>	<i>Light</i>	<i>Low</i>
9	<i>Large</i>	<i>Not greasy</i>	<i>Medium</i>	<i>Low</i>
10	<i>Medium</i>	<i>Greasy</i>	<i>Heavy</i>	<i>High</i>
11	<i>Medium</i>	<i>Greasy</i>	<i>Light</i>	<i>Low</i>
12	<i>Medium</i>	<i>Not Greasy</i>	<i>Heavy</i>	<i>Medium</i>
13	<i>Medium</i>	<i>Not Greasy</i>	<i>Light</i>	<i>Low</i>
14	<i>Medium</i>	<i>Not Greasy</i>	<i>Medium</i>	<i>Low</i>
15	<i>Medium</i>	<i>Medium</i>	<i>Light</i>	<i>Low</i>
16	<i>Medium</i>	<i>Greasy</i>	<i>Medium</i>	<i>Medium</i>
17	<i>Medium</i>	<i>Medium</i>	<i>Heavy</i>	<i>Medium</i>
18	<i>Medium</i>	<i>Medium</i>	<i>Medium</i>	<i>Medium</i>
19	<i>Small</i>	<i>Not Greasy</i>	<i>Light</i>	<i>Very_Low</i>
20	<i>Small</i>	<i>Not Greasy</i>	<i>Heavy</i>	<i>Medium</i>
21	<i>Small</i>	<i>Greasy</i>	<i>Light</i>	<i>Low</i>
22	<i>Small</i>	<i>Not Greasy</i>	<i>Medium</i>	<i>Low</i>
23	<i>Small</i>	<i>Medium</i>	<i>Light</i>	<i>Low</i>
24	<i>Small</i>	<i>Medium</i>	<i>Heavy</i>	<i>Medium</i>
25	<i>Small</i>	<i>Greasy</i>	<i>Medium</i>	<i>Medium</i>
26	<i>Small</i>	<i>Medium</i>	<i>Medium</i>	<i>Low</i>
27	<i>Small</i>	<i>Greasy</i>	<i>Heavy</i>	<i>High</i>

These rules have been shown as a membership function in Figure 3 and Figure 4.




Figure 3: Membership function graphs for all three linguistic inputs




Figure 4: Membership function graphs for all five adjectives of linguistic output (Wash Time)

Since the fuzzy rules have been defined in imprecise sense they too are the fuzzy values (not crisp values) [7]. The three input parameters after being read from the sensors are fuzzified as per the membership functions of the respective variables. Fuzzy inference is then carried out using fuzzifier. The fuzzy rules are implemented by the MIN and MAX inference logic array according to the fuzzy rules mentioned above. After the membership functions and the fuzzy rules have been derived, the ordinary fuzzy inference process is then applied [4].

Defuzzification is carried out using Centre of Gravity method. Finally, the crisp value for the wash time is obtained as an answer. This crisp output value that we get is equivalent to Wash time (in minutes)/5. Thus, this value when multiplied by 5 (scaled by 5) gives the total wash time in minutes. The logic design for the MIN (or MAX) unit for the fuzzy inference is shown in Figure 5. All logic gates are static CMOS standard cells. The overall hardware design is shown in Figure 6. Figure 7 shows the hardware design of the defuzzifier. Figure 8 shows the I/O pin definition for the chip. There are 37 I/O pins. Fuzzy subsets (membership functions) representing adjectives for all 3 linguistic inputs- **dirtiness of the clothes, type of dirt and mass of clothes** as well as adjective for linguistic output – **Wash time** are stored in ROM modules, one adjective per ROM module. ROM modules for all the adjectives of Linguistic Inputs and adjectives of linguistic output are 8 bits in size.




Figure 5: Logic Design of a MIN (MAX) Circuit




Figure 6: Schematic of hardware architecture of Fuzzy Wash time controller for Washing machine




Figure 7: Defuzzification Hardware




Figure 8: Input/Output pin definition for the fuzzy microcontroller chip

4 SIMULATION RESULTS

The schematic design is prepared using S-Edit (a software tool for schematic design) and the functionality is tested using W-Edit (a software tool for simulation) and finally a complete chip layout is performed using L-Edit (a software tool for layout design).

Circuit simulation plays an important role during early design steps and final verification phases [5]. W-Edit is a waveform viewer that visualises the complex numerical data resulting from VLSI Circuit Simulation by presenting them in graphical form. Figure 9 shows an example of graphical results obtained upon simulating the MAX circuit used in this design.




Figure 9: Graphical representation of MAX Circuit Simulation result using W-Edit

In this example, two inputs A and B, where A=00000000 and B=11111111 were fed on the input ports of the MAX circuit. When the circuit is simulated, the resulting output 11111111 is obtained. In other words, all 8 bits of the resulting output are 1 (or 5V) as shown in Figure 9. This is correct because $11111111 > 00000000$ and MAX circuit is designed to compare the two sets of 8 bit binary numbers and output the one that is higher as the result.

Similarly, all the other circuits are simulated and the corresponding results are analysed to ensure that the correct circuit design is implemented

5 CONCLUSION

The functional and logic design for an improved Fuzzy controller chip has been described. The hardware requirements are quite modest and the entire design can be derived by using standard cells. Also, the memory (ROM) modules required are of minimal size and have been significantly reduced by using the symmetry of fuzzy rules and fuzzy subsets. Although this particular design controls only the Wash time of a washing machine, the design process can be extended to other control variables such as Water level and Spin speed. The formulation and implementation of membership functions and rules will be similar to the one described in this design.

5 REFERENCES

- [1] H. Watanabe et. al., "A VLSI Fuzzy logic controller with reconfigurable, cascable architecture," IEEE Trans. Solid State Ckts, vol.25, no.2, pp. 376-382, Apr.1990
- [2]V. Catania et. al., "A VLSI Fuzzy Inference Processor based on a discrete analog approach," IEEE Trans. Fuzzy systems, vol.2, no. 2, pp 93-106, May. 1994
- [3] S.M. Rezaul Hasan, "Design of a pH controller ASIC Fuzzy logic chip for industrial application," Proc. GlobalTRONICS Electronics Design Conference, WTC, Singapore, 1994.
- [4] P. Hong and J.B. Chen, "Processing individual fuzzy attributes for fuzzy rule induction," Fuzzy Sets and Systems, vol.112 , no.1, pp. 127-140, 2000.
- [5] M. Pedram and S. Nazarian, "Thermal modeling, analysis and management in VLSI circuits: Principles and Methods," Proceedings of the IEEE Special Issue on On-Chip Thermal Engineering, Vol. 94, no. 8, pp 1487-1501, August 2006.

Speedometer Calibration Unit

Sindhoor Nalli

Manukau Institute of Technology
Private Bag 94006, Manukau.
nallisindhoor@gmail.com

Abstract: This project aims at creating an automated process for checking the accuracy of a vehicle's speedometer. This reduces the time taken to test the speedometer and also makes the process more accurate than traditional handwritten reports of the inspection. The speedometer is used for maintaining an appropriate pace while driving a vehicle. Vehicle owners may wish to check their speedometer's accuracy after receiving a speeding fine, or when they fit different sized tyres to their vehicle.

Keywords: Speedometer, calibration, windows application, automatic process.

1 INTRODUCTION

Speedometers are a very important part of the vehicle's instrumentation. They help the driver to stick to speed limits, and so indirectly assist in avoiding accidents. Most of the speedometers in use today are not accurate. They generally read about 10% too high (ie when the vehicle is travelling at 50 km/h they will read 55). Different tyre pressures also affect the accuracy. Speedometer readings will change when the vehicles' tyres are changed from a smaller to a bigger or vice-versa.

The main aim of this project was to creating a client user interface to automate the vehicle checking process, create and store the test result as a report to give the clients, and store a copy of it in the system for future reference.

2 DISCUSSION

To measure a vehicle's speed a device known as a "rolling road" is used. It consists of a pair of in-ground rollers which the driving wheels of the vehicle are driven onto. As the wheels turn, they turn the rollers and the vehicle remains stationary. For this project, the rollers used to check the speedometer accuracy drive a toothed gear which interrupts an infrared light source as it spins. A light sensor generates a pulse for each interruption which is then amplified to produce an 8 volt square wave as the output. Various speeds are represented by varying frequencies in the output. The data is received via a LabJack U3 Data Acquisition module and sent to a computer to be processed. Since this process is automated, it saves time and is much more accurate than recording the results manually.

As a start, understanding the working of the rolling road is very important. In particular the circumference of the rollers and the number of pulse per revolution must be known in order to calculate the vehicle's speed. This information is contained in the following section.

2.1 Hardware Description:

A main device used in this project is the LabJack U3 Data Acquisition Module. This module acts as an interface between the signals from the rolling road to the computer. These are devices designed to measure and log some parameters that can be used for further analysis.



This module acquires signals and waveforms and processes them to obtain desired information for further analysis by a computer by using a vendor supplied software or various programming languages like Java, Visual Basic, C, C++, C# and many others.


Below are a few features that are available in this specific product which makes the data decoding easier and hassle free.

- 1) It has a 16 flexible I/O ports
- 2) 2 Counters and 2 Timers (32-Bits Each)
- 3) Very fast response times of less than 1ms.
- 4) USB 2.0/1.1 Full Speed Interface
- 5) Drivers available for a number of OS including Windows and Mac.

For establishing a connection with the LabJack module using any programming language the 'LabJackUD' reference library is provided by the LabJack manufacturer is used. The important feature that is we use in this project is the timers. The timers count the number of pulses in a certain time interval. This value after a some minor calculations can be used to get the actual speed of the vehicle. The timers are very close to being accurate. At 120km/hr the timer reads a value of **i dont know**. And at the minimum speed of 30 km/hr the timer reads a value of **i dont know**.

2.2 Software Description:

At the start of the project, the following process was used to calculate the frequency at a certain speed. Using this, the maximum frequency was calculated which helped while selecting the right Data Acquisition Module. The following is the flowchart of the frequency calculation at the start of the project.



2.1.1 Main Window




The Main Window consists of four buttons – 'Add New Client', 'Vehicle Test', 'Create Report', and 'Information' – and all these buttons open up new windows where they can perform tasks their tasks appropriately.

2.1.2 Adding New Client Window



The above is a snapshot of the Client Details Storage Window. Here the basic details like the Name, Registration Number of the vehicle, its make, the tyre size and the Customer ID is stored. When the 'Store' button is clicked, the details are stored into a table called 'client' in the database.

For the database, MySQL, a Relational Database Management System is being used.



2.1.3 Vehicle Testing




The figure above is a snapshot of the vehicle testing window. The label that says, 'Target Speed' is the speed that we want to check the vehicle's speedometer speed reading. For this, the vehicle is accelerated until it reaches that speed and when the vehicle is constant at that speed the button is clicked. During the time the vehicle is accelerating, at the label 'Actual Vehicle Speed', 'XXX' in the 'XXX Km/hr', the speed is updated 10 times every second so as to get a live display of the changing speed.

For calculating the speed, the counter feature that is available in LabJack UD is used which counts the number of falling edges in an 8V square wave which is the output from the rolling road. The frequency varies ac-

ording to the speed. This frequency can be calculated manually after considering the diameter of the roller which is explained in detail in figure 2 in the previous page.

When this button is clicked, the "Test Status" label in the Main form is changed from "Not Completed" to "Test Started". This reverts to a blank label when the test is completed.

The next figure is the flowchart for the vehicle's speedometer checking process.




The above is the basic working of the vehicle testing window. There is a 'BACK' button on the testing form that can be used to back to redo the previous speeds of the test.

When this button is pressed, it decrements the value of the variable 'i' and checks if there is a target speed below the current one and goes to it if it can otherwise shows a message to the user that 'it cannot go any lower'. This button makes the vehicle testing more efficient and accurate.

On the same window there is a live display of the speed. This is triggered when the start button is clicked the first time.

The function that stores the speed and the one that is used for displaying the live speed is the same and its process is described using the following flowchart.



2.1.4 Information Window



On this screen, all the important aspects of the program are mentioned. For instance, the different speeds at which the checking takes place, the default location of where the report is saved, the tags that are used in the sample report, which are then replaced by actual values from the database when the report is created.

The client report is created using a sample report that uses tags like '<name>', '<rege>', '<speed_1>', '<speed_2>' and so on. The tags are named exactly like the headings in the database table. Then we use the Find and Replace

method from Microsoft Word to find and replace the tags in the sample document to create a new document using the current date and the client number and vehicle registration number so that the report that is created is unique for that particular test.

2.3 Database Structure:

There are three data tables used in this project. One for client details, second one for the vehicle test results and the third one for the target speeds that the vehicle is tested at.

2.3.1 Client Table:

All the New client details are stored here. The client table is structured as follows:

Customer ID:	Varchar(20) (Primary key)
Name:	Varchar(45)
Rego:	Varchar(45)
Make:	Varchar(45)
Tyre:	Varchar(45)

2.3.2 "Test Results" Table:

Rego	Varchar(20)
Date	Varchar(45)
Speed_1	Varchar(45)
Speed_2	Varchar(45)
Speed_3	Varchar(45)
Speed_4	Varchar(45)
Speed_5	Varchar(45)
Speed_6	Varchar(45)
Speed_7	Varchar(45)
Speed_8	Varchar(45)

Speed_9	Varchar(45)
Speed_10	Varchar(45)
Inspector	Varchar(45)

3 FURTHER WORK

In the future, more work can be done for sending the report as an email attachment to the user straight from the User Interface.

On the Information Screen, the user should be able to change the default location for the storage of the final test report of the vehicle. The target speeds at which the vehicle is tested can be changed.

The buttons on the form should be triggered when a remote push-button is clicked by the vehicle inspector while being seated in the vehicle.

If the user wants to exit the vehicle test in the middle of the test because of various reasons, the test completed till that point should be stored properly in the database.

4 CONCLUSIONS

The project was successful and was tested positively at a vehicle testing station.

5 BIBLIOGRAPHY

- 1) www.labjack.com
- 2) http://en.wikipedia.org/wiki/Data_acquisition
- 3) The LabJack U3 User Guide
- 4) www.aspose.com

Power Characterisation of IEEE 802.15.4 and Zigbee Wireless Networks

A. Prince-Pike, J. D. Collins

Department of Electrical & Electronic Engineering,
Auckland University of Technology,
Auckland, New Zealand.
aprincep@aut.ac.nz, jcollins@aut.ac.nz

Abstract:

This paper analyses the power consumption and delivery ratio of Zigbee wireless networks using the mesh and star topologies for a single sink real time monitoring system. The analysis involves modelling the behaviour of a Zigbee network at varying traffic rates and examining the operation of the beacon and non beacon modes of its underlying IEEE 802.15.4 standard in the star network architecture. IEEE 802.15.4 networks have been analysed with NS-2 simulations and hardware experiments. Zigbee was only analysed using hardware experiments due to limitations with NS-2. The results show that non beacon mode has the lowest power consumption and best delivery ratio at all tested traffic rates.

Keywords:

IEEE 802.15.4, Zigbee, NS-2, power characterisation, wireless networks.

1 INTRODUCTION

Zigbee is a relatively new wireless mesh networking standard with emphasis on low cost and energy conservation. It is intended for use in wireless monitoring and control applications such as sensors and remotely operated switches where the end nodes are battery powered.

The Zigbee networking protocol was developed in response to the unsuitability of current wireless protocols like WIFI and Bluetooth for applications which require very long battery life, high reliability and low data rates such as remote monitoring and building automation. The first Zigbee specification was ratified in late 2004 and is based upon the physical (PHY) and media access control (MAC) layers of IEEE standard 802.15.4.

Since Zigbee and its underlying standard IEEE 802.15.4 are relatively new, there has been little research investigating the power consumption of the different network topologies and comparison between the operating modes. Most studies focus on the beacon enabled network mode because many applications require bi-directional data flow, and in star networks research has shown that it gives worthwhile energy savings. However in a real time monitoring application, data flow is mostly one way from the sensor node to a central storage device for processing and recording and therefore beacon enabled mode might not give the best energy savings.

This research investigates the power consumption of the transceiver and packet delivery ratio with a varying traffic rate for both Zigbee and pure IEEE 802.15.4 wireless networks under different system configurations. For IEEE 802.15.4 using the star network topology, the effect of two system parameters, Beacon Order (BO) and

Superframe Order (SO) was recorded on the power consumption and delivery ratio of packets sent from nodes to the coordinator using simulations with NS-2 and results obtained from real network experiments.

The Zigbee network was analysed on a real network using both the star topology and the much more popular mesh network topology with a varying traffic rate. Due to limitations with NS-2, the Zigbee network was not able to be simulated.

1.1 Related Research

Since Zigbee and its underlying standard IEEE 802.15.4 are relatively new, there has been little research investigating the power consumption and delivery ratio of the different network topologies and comparison between the operating modes.

Zheng and Lee [1] developed a computer simulation model of the media access control (MAC) and physical (PHY) layers of IEEE 802.15.4 for Network Simulator-2 (NS-2)[2] to quantify its operation. Their research shows that IEEE 802.15.4 is an excellent low power low data rate wireless standard upon which applications can be built. The simulation model has been the subject of a large number of research projects, several of which are reviewed in the following paragraphs.

Huang and Pang [3] and Kovakka et al [4] investigated the effect of BO and SO on the power consumption of a small star network, while [5] also investigated non beacon mode. Their work showed that these parameters can considerably increase the power consumption if the parameters are set incorrectly.

A beacon enabled star network was also researched by Ling-xi et al. [6]. They demonstrated that adjusting the transmission power does not have a large effect on overall power consumption because the time spent in transmit mode is small. They and [7] showed that increasing the number of network nodes increases node power consumption due to extra network contention.

Research by [8] is one of the few studies that focused on real Zigbee devices. Their research focused on investigating data rate and delivery ratio with varying BO, number of nodes and data packet sizes. Although they only focused on the situation where BO = SO (duty cycle 100%) they show that a higher throughput is possible in non beacon mode.

Singh et al [9] highlighted several inconsistencies between the NS-2 IEEE 802.15.4 simulator model and the IEEE 802.15.4 standard and demonstrated that the clear channel assessment functionality in IEEE 802.15.4 consumes a significant amount of energy, particularly at higher data rates.

Rao [10] performed extensive analysis on a beacon mode IEEE 802.15.4 star network using NS-2. His research discovered several issues with the NS-2 simulation model which were resolved. Battery life was investigated which showed that IEEE 802.15.4 nodes are likely to have a long battery life and good delivery ratio.

1.2 IEEE 802.15.4 and Zigbee Overview

IEEE 802.15.4 was conceived due to the unsuitability of current wireless standards such as Wi-Fi and Bluetooth for low data rate battery powered ad hoc networks. The standard specifies the MAC and PHY layers of the open standards interconnection (OSI) network model while leaving the development of the upper layers to the designer.

The standard defines three types of network nodes, PAN coordinators, Full Function Devices (FFDs) and Reduced Function Devices (RFDs). The PAN coordinator is responsible for creating the network and is often used as a gateway to other networks such as Ethernet. There must be only one PAN coordinator per network. FFDs are capable of communicating with all device types as well as creating sub networks and managing routing and addressing of RFDs. RFDs are intended to be extremely simple devices with minimal hardware and software resources. RFDs can only communicate with FFDs or the PAN coordinator, not with another RFD.

1.2.1 IEEE 802.15.4 Network Topologies

IEEE 802.15.4 networks can be divided into two main topologies, star and peer-to-peer. A third topology called cluster tree is a variation of the peer-to-peer topology.

The star network topology is more structured than peer-to-peer. The PAN coordinator of the network is always the central node. Communication between devices always occurs via the PAN coordinator which relays mes-

sages between devices. Direct messaging between end devices is not permitted.

In the peer-to-peer topology, an arbitrary array of connections can be created between full function devices and the PAN coordinator. Since a network layer is not defined in the standard, routing is not directly supported.

A cluster tree topology is also possible. This is a special case of the peer-to-peer topology. This exploits the fact that RFDs can only associate with FFDs and is used where most devices in the network are FFDs.

1.2.2 IEEE 802.15.4 Physical Layer

The Physical layer in IEEE 802.15.4 is designed to operate in unlicensed frequency bands world wide. Since not all the frequency bands are the same world wide, IEEE 802.15.4 provides three possible operating frequencies, 868 MHz, 915 MHz and 2.4 GHz. Each frequency except 868 MHz has several channels which can be selected by the user. Table 1 lists the available frequencies and channels that can be used by IEEE 802.15.4.

For this research, the 2.4 GHz physical layer has been chosen due to the world wide availability of this frequency band and the wider availability of wireless transceivers at this frequency.

Table 1: Available channels at different frequency bands and locations

Frequency	Available channels	Data rate (kbps)	Locale
2.4 GHz	16	250	World wide
915 MHz	10	40	USA
868 MHz	1	20	Europe

1.2.3 IEEE 802.15.4 MAC Layer

IEEE 802.15.4 networks are able to operate in two different modes of operation, beacon mode or non beacon mode. In non beacon mode, nodes contend for channel access using CSMA/CA. In beacon mode, the network is fully synchronised as the coordinator sends out periodic packets or beacons. This enables nodes to sleep between beacons thus conserving energy. In beacon mode, all transmissions use the superframe structure illustrated in Fig 1.

The superframe is divided into 16 slots with the first slot for the beacon and the rest of the slots for nodes to communicate. At the end of the superframe is an inactive period where devices can sleep. The structure of the superframe is determined by the coordinator and consists of three main sections, the contention access period (CAP), an optional contention free period (CFP, also known as guaranteed time slots or GTS) and the inactive period. The coordinator also determines the composition of the superframe with the Beacon Interval (BI) and Superframe Duration (SD) being adjusted by the Beacon Order (BO) and Superframe Order (SO) parameters.




Figure 1: IEEE 802.15.4 superframe structure[11]



Figure 2: CC2431BB development board

1.2.4 Zigbee

Zigbee is a wireless mesh networking standard which is based on the MAC and PHY layers of IEEE 802.15.4. The Zigbee stack layer sits above IEEE 802.15.4 and provides most of the network functionality for the network. This layer is responsible for network formation (if the device is a coordinator), assigning network addresses, routing, security and route discovery.

The network topologies possible with Zigbee are the same as with IEEE 802.15.4 but the peer-to-peer topology has been enhanced due to the routing mechanisms present in Zigbee and is now referred to as the mesh network topology.

With Zigbee the three device types are typically called Zigbee Coordinators (ZC), Zigbee Routers (ZR) and Zigbee End Devices (ZED or just end device).

2 SIMULATION & EXPERIMENTAL ENVIRONMENT

For this research, the following metrics were analysed:

Delivery ratio: This is the percentage of data packets which are successfully received versus the number of packets actually transmitted and is an indicator of how reliable the network link is.

Power consumption: This is the amount of power consumed by an end device in milliwatts.

Battery life: This is the node power consumption converted into a more human readable metric using a chosen battery capacity.

The performance of IEEE 802.15.4 was analysed using the star network topology using simulations and hardware experiments in both beacon and non beacon mode.

Experiments were also performed on hardware devices using the same parameters as the simulations. For this network topology, 10 end devices were arranged in a circle with a 10 metre radius. A single coordinator was located at the centre. This is illustrated in Fig 3.

The chosen hardware platform for this research was a Chipcon CC2431BB development board which can be seen in Fig 2. The simulator was configured to match the characteristics of this hardware.




Figure 3: IEEE 802.15.4 and Zigbee star network topology




Figure 4: Zigbee mesh network topology

Experiments were also performed for Zigbee networks in both the star and mesh network architectures, however they were not able to be simulated due to limitations with the simulator. Fig 3 and 4 show the network topologies investigated and how the results were obtained

from the nodes. Table 2 illustrates the parameters that were common to both the simulations and experiments.

2.1 Simulations

NS-2 is a discrete event simulator developed in a collaborative effort by many institutions and contains code contributions from many researchers. For this research, NS-2 version 2.33 [2] was used which was the most recent version at the time.

The Zigbee/IEEE 802.15.4 simulation model that has been used in this research and is now included in NS-2 was developed by Zheng and Lee [12] from the City College of New York. Ramachandran [13] also contributed many modifications to the original simulation model which have been used in this research.

To get accurate results from the simulations, the simulator needs to be configured to match the hardware characteristics of the Chipcon CC2431BB development board. The additional parameters that need to be configured which are specific to the simulator are listed in Table 3.

Several scripts were written to automate the simulation process. The scripts ran the simulation scenarios, automatically varying the required parameters over the required range as well as extracting the desired metrics from the NS-2 trace files and performing the require data analysis to determine the power consumption, battery life and packet delivery ratio.

2.1.1 Simulator Modifications

Initial simulator results were very different from the initial hardware experiments. After investigating the operation of the simulator the following modifications were made.

2.1.2 AODV and ARP Packets

During the simulations it was noticed that large numbers of Ad-Hoc On Demand Distance Vector (AODV) and Address Resolution Protocol (ARP) packets were being broadcast between nodes which did not match the experimental network traffic on the packet sniffer. Because the simulations involve a star network with the coordinator always being the destination, AODV routing broadcasts are not required. The AODV module in NS-2 was subsequently modified to prevent the route resolve procedures being called using modifications by Rao [10]. ARP was also deemed unnecessary as noted by [10] and [14] and was disabled using modifications provided by Rao [10].

2.1.3 State transitions

While the energy model in NS-2 took into account the energy consumed during state transitions i.e. from receive to transmit, it did not take into account the energy consumed during the power up phase. The Chipcon data sheet specifies a 320µS delay, which was added to the energy model.

2.2 Experiments

To validate the results of the simulator, the same experiments were performed on Chipcon CC2431BB development boards using the Texas Instruments IEEE 802.15.4 and Zigbee network software.

The power consumption was measured using a Measurement Computing data acquisition module connected to a PC running MATLAB which recorded the voltage drop across a precision resistor. A MATLAB script controlled the data acquisition module and calculated the power consumed of the device as well as calculating the battery life based on the battery capacity in Table 2. Due to the time taken to perform the experiments, the power consumption of a single node was measured, as shown in Fig 2 and 3 and it was assumed that it would be similar for the remaining 9 nodes. Research by [15] shows that this is likely to be a valid assumption.

The IEEE 802.15.4 energy model in NS-2 only models the power consumed by the transceiver, and does not include the microcontroller. Because of this, additional post processing was required to exclude the microcontroller power consumption during the wakeup phase from sleep mode before the packet is transmitted.

The coordinator was configured to output received packets to its serial port and an application on the computer analysed the received packets to calculate the delivery ratio for each node.

Zigbee was only able to be tested in non beacon mode due to limitations with the Texas Instruments Zigbee stack.

Table 2: Configuration parameters common to simulations and hardware experiments

Parameter	Range
BO	6-10 and 15(Non beacon mode)
SO	0-2 and 15 (Non beacon mode)
Data rate (s/packet)	30, 60, 100, 200 and 1000
Simulation/Experiment time	1000 seconds
Number of end devices	10
Distance between nodes	10 metres
Traffic direction	Node to coordinator
Packet size	64 bytes
Battery capacity	900 mAH

Table 3: Simulation specific configuration parameters

Parameter	Value
Transmit current consumption	30mA
Receive current consumption	33mA
Sleep current consumption	0.5µA
Random traffic jitter	Disabled
Energy capacity	9720 Joules
Network topology	Star

3 RESULTS

In this section the performance of IEEE 802.15.4 and Zigbee wireless networks are analysed to evaluate the power consumption and delivery ratio of both network types and to compare the simulation and experimental results. Results for the star network topology simulations in both beacon and non beacon mode are compared with results from experiments to determine the optimum network operating mode and also validate the accuracy of the simulator.

The graphs in the following sections illustrate the battery life and delivery ratio results. The power consumption figures have not been included as, for this evaluation, an estimate of battery life was more meaningful in evaluating the suitability of this technology for its use in a real time monitoring system. The aim is to maximise both battery life and packet delivery ratio.

3.1 Beacon mode results

For beacon mode networks, the modified simulator is able to estimate battery life to within 15 percent, often with much greater accuracy (Fig 5 and 7). The exception is beacon order 9 at 1000 seconds/packet, where the error is around 20 percent (due to a lack of space this figure was omitted).

Figures 6 and 8 show a considerable difference in delivery ratio between the simulated and experimental results. The experimental results usually had a 100 percent delivery ratio whereas the simulated delivery ratio began to decrease as soon as the beacon order increased. Analysing the NS-2 trace files shows that this is due to packet collisions between devices. It is suspected that this behaviour is due to a scheduling issue in the simulator as the nodes were started at slightly different times to avoid collisions.

These results show that when the network is operating in beacon mode, in order to maximise battery life (minimise power consumption) and delivery ratio, a good estimate of the traffic rate is required in order to determine the correct BO and SO. It is also noted that increasing SO usually results in a better delivery ratio but often results in a small increase in increase in power consumption due to the longer CAP.

3.2 Non beacon mode results

The delivery ratio graph for non beacon networks has been omitted as the delivery ratio was 100% for all results. Fig 9 shows that using Zigbee results in a significant decrease in battery life (increase in power consumption) compared to IEEE 802.15.4 networks due to the extra routing packets being broadcast. However the resulting battery life is so long that it would outlast the shelf life of most batteries.




Figure 5: Battery life at 30 s/pkt




Figure 6: Delivery ratio at 30 s/pkt




Figure 7: Battery life at 100 s/pkt




Figure 8: Delivery ratio at 100 s/pkt




Figure 9: Non beacon mode results

4 CONCLUSIONS

This research has shown that despite previous research focusing on beacon mode, in single sink IEEE 802.15.4/Zigbee wireless networks non beacon mode gives the longest battery life (lowest power consumption) and the best delivery ratio at all tested data rates in both computer simulations and the experimental investigation.

From the results presented, IEEE 802.15.4 in beacon mode in a single sink wireless network results in significantly lower battery life (higher power consumption) and a poorer delivery ratio. However it is likely that in applications where there is bidirectional traffic, the synchronisation provided by beacon mode could be an advantage and could result in fewer collisions than in non beacon mode. This could be the subject of future research.

As mentioned by previous research, and as is apparent in the above results, when operating in beacon mode an idea of the traffic rate is required in order to correctly set BO and SO to maximise delivery ratio and minimise power consumption.

After performing the modifications described to NS-2, results obtained from the simulator were usually within 15 percent of the experimental results. However the poor agreement of the delivery ratio results and apparent scheduling issue indicates that further work is required on the simulator to improve its performance in this area.

5 REFERENCES

1. Zheng Jianliang, Lee Myung J. "Will IEEE 802.15.4 Make Ubiquitous Networking a Reality?" *IEEE Communications Magazine*. 2004. p. 140-146.
2. "The Network Simulator - NS-2 - Version 2.33". Available from: <http://www.isi.edu/nsnam/ns/>
3. Yu-Kai Huang, Ai-Chun Pang. "A Comprehensive Study of Low-Power Operation in IEEE 802.15.4". *MSWiM '07: Proceedings of the 10th ACM Symposium on Modeling, analysis, and simulation of wireless and mobile systems*. ACM, 2007. p. 405-408.

4. Kovakka Mikko, Kuorilehto Mauri, Hannikainen Marko, Hamalainen Timo D. "Performance analysis of IEEE 802.15.4 and Zigbee for large scale wireless network sensor applications". *Proceedings of the 3rd ACM international workshop on Performance evaluation of wireless ad hoc, sensor and ubiquitous networks*. ACM, 2006. p. 48-57.
5. Hoffat Joe, Klues Kevin, Orjih Obi. "Configuring the IEEE 802.15.4 MAC layer for single sink wireless sensor network applications". Washington University, St. Louis, Missouri, 2005.
6. Ling-xi Wu, Jie Zhan, Wei Shi. "The Low Power Study Based on Zigbee Network". *4th International Conference on Wireless Communications, Networking and Mobile Computing, WiCOM 2008*. IEEE. p. 1-4.
7. Chen Zhijia, Lin Chuang, Wen Hao, Yin Hao. "An Analytical Model for Evaluating IEEE 802.15.4 CSMA/CA Protocol in Low-Rate Wireless Application". *21st International Conference on Advanced Information Networking and Applications Workshops*. IEEE, 2007. p. 899-904.
8. Lee Jin-Shyan. "An Experiment on Performance Study of IEEE 802.15.4 Wireless Networks". *10th IEEE Conference on Emerging Technologies and Factory Automation*, 2005. p. 451-458.
9. Singh Chandramani Kishore, Kumar Anurag, Ameer P.M. "Performance evaluation of an IEEE 802.15.4 sensor network with a star topology". *Wireless Networks*. 2007;vol. 14, no. 4. p. 543-568.
10. Rao Vaddina Prakash. The simulative Investigation of Zigbee/IEEE 802.15.4: Dresden University of Technology; 2005.
11. "IEEE Computer Society 802.15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs)". 2006.
12. Zheng Jianliang, Lee Myung J. "NS-2 simulator for 802.15.4". The City University of New York, Available from: <http://ees2cy.engr.cuny.cuny.edu/zheng/pub/file/wpan11.tar.gz>
13. Ramachandran Iyappan. "Changes made to the IEEE 802.15.4 NS-2 Implementation". University of Washington, 2006 Available from: http://www.ee.washington.edu/research/funlab/802_15_4/
14. Zheng Jianliang. "Low Rate Wireless Personal Area Networks (LR-WPANs) FAQ". The City University of New York, 2006.
15. Fang Miaoqi, Wan Jian, Xu Xianghua, Wu Guanrong. "System for temperature monitor in substation with ZigBee connectivity". *11th IEEE International Conference on Communication Technology*, 2008. p. 25-28.

Low-Cost Temperature Sensor on a Modern Submicron CMOS Process

Robert P. Fisk¹ and S. M. Rezaul Hasan²

Center for Research in Analog & VLSI Microsystem Design (CRAVE)
School of Engineering and Advanced Technology, Massey University
Albany Highway, Gate 4 Building 106, Albany, Auckland, New Zealand
¹rfisk@massey.ac.nz and ²hasanmic@massey.ac.nz

Abstract: The design of an integrated smart temperature sensor is presented that achieves high accuracy and low cost. The use of a modern 0.13 μm process allows high layout density, while circuit techniques are used to optimise area usage thus further reducing costs. A second-order delta-sigma analog-to-digital converter (ADC) is used with a third-order sinc decimation filter to achieve a 16-bit digital output resolution while remaining insensitive to analog imperfections. Active circuitry of the completed sensor occupies an area of 0.21mm².

Keywords: Smart temperature sensor, substrate bipolar transistor, delta-sigma conversion, symmetric decimation filter.

1 INTRODUCTION

So-called ‘smart’ temperature sensors integrate both a temperature sensor and an analog-to-digital converter (ADC) on the same integrated circuit. IC manufacturing processes excel at producing high product volumes at a low unit cost. Smart integrated temperature sensors leverage this advantage, and are ideally suited to high-volume applications requiring a rugged sensor with minimal support circuitry.

Smart IC temperature sensors were originally developed for bipolar processes, and thus had very simple ADCs [1]. The transition to high-density CMOS processes enabled increasing ADC complexity and a corresponding increase in sensor accuracy. Published accuracy figures for smart temperature sensors have improved from $\pm 1^\circ\text{C}$ [2] to $\pm 0.1^\circ\text{C}$ [3, 4], the latter representing the highest calibrated accuracy for a smart integrated sensor reported to date.

The above smart sensors were manufactured on IC processes that are considered rather elderly today; [4] used a 0.7 μm process with 2 metal layers and occupied a chip area of around 2.5mm² excluding bondpads and delta-sigma decimation filter. Using a modern process with reduced linewidth allows a higher functional density, thus reducing die area and manufacturing costs. Circuit techniques are described in this paper to further reduce layout area.

This paper presents a design that targets the same $\pm 0.1^\circ\text{C}$ accuracy level as [4], yet occupies significantly less chip area. Section 2 of this paper introduces the fundamentals of integrated temperature measurement and presents the block structure of the smart temperature system. Section 3 describes the factors influencing the choice of target manufacturing process. Sections 4 and 5 present the de-

sign of the analog and digital sections of the on-chip ADC. The paper concludes with Section 6.

2 SYSTEM STRUCTURE

2.1 On-chip sensing techniques

Smart sensors produce an output in the digital domain by integrating an on-chip ADC alongside the sensing element. This digital output value is a ratiometric quantity, dependent on the ratio between a temperature-dependent signal and an invariant reference. A smart sensor therefore requires both a temperature sensor and a reference generator to be integrated on-chip, and both of these are limiting factors to the overall system performance.

Fortunately, techniques for creating a temperature-invariant voltage reference have been well-understood since the 1970s [5, 6]. A stable voltage reference cannot be created using any single on-chip device, as all significant device parameters exhibit some degree of temperature dependence. The well-known bandgap reference technique uses bipolar transistors to create two voltages with predictable and opposing temperature coefficients. An appropriate linear combination of these voltages produces the desired temperature-invariant behaviour.

The voltage across the base-emitter junction of a bipolar transistor is related to its collector current by:

$$V_{BE} = V_T \ln\left(\frac{I_C}{I_S}\right), \quad (1)$$

where the thermal voltage V_T is linearly proportional to absolute temperature. Eq. (1) can be manipulated to reveal the temperature dependency of V_{BE} [7]:

$$V_{BE} = V_{G0} + V_T \ln\left(\frac{I_C T^{-\gamma}}{CA}\right), \quad (2)$$

where V_{G0} is the bandgap voltage of silicon, A is the area of the base-emitter junction, and C is a process-dependent constant. This forward diode junction voltage exhibits a nearly linear temperature characteristic with a slope of around -2mV/K (Fig. 1). The gradient of this temperature dependence is influenced by the bias current density (I_C / A) and device characteristics C of (2). Higher-order curvature is introduced through the factor $T^{-\gamma}$.

In order to produce a reference voltage, the negative temperature coefficient of V_{BE} must be cancelled by a voltage with a positive temperature coefficient. This can be produced as follows. Two bipolar transistors carefully laid out on the same IC will have virtually identical device characteristics. Therefore the difference between their base-emitter voltages can be expressed as:

$$\begin{aligned} \Delta V_{BE} &= V_{BE1} - V_{BE2} \\ &= V_T \ln\left(\frac{I_{C1} \cdot A_2}{I_{C2} \cdot A_1}\right). \end{aligned} \quad (3)$$

This neatly eliminates the various sources of non-linearity and process dependence present in (2), and leaves a voltage dependent on a few easily-controlled parameters and linearly proportional to absolute temperature (PTAT) (Fig. 1). A current density ratio of 8:1 produces a PTAT voltage of approximately $180\mu\text{V/K}$, which must be amplified in order to completely cancel the temperature dependence of V_{BE} and obtain a constant reference (Fig. 1):

$$V_{REF} = V_{BE} + \alpha \Delta V_{BE}. \quad (4)$$

This reference V_{REF} is required by the sensor's on-chip ADC to convert the temperature signal to a digital value. Either of the two temperature-dependent voltages in (4) may be used as the temperature signal itself, although in practice the PTAT voltage ΔV_{BE} is always used due to its predictable linear temperature relationship and insensitivity to process-induced variations.




Fig. 1. The temperature characteristics of the base-emitter junction voltage V_{BE} , and the derived voltages ΔV_{BE} and V_{REF} .

2.2 Block Diagram

At a conceptual level, the smart temperature sensor system can be divided into two distinct regions – an analog sensor / reference, and an ADC (Fig. 2). The analog front-end is further divided into a bias current generator, process compensation circuit, and bipolar sensor core. The raw analog values V_{BE} and ΔV_{BE} produced in the bipolar core are connected directly to the input of the ADC.




Fig. 2. Block diagram of the smart temperature sensor.

To allow the complete system to achieve the highest possible accuracy, it is necessary for the ADC to surpass the system accuracy target of $\pm 0.1^\circ\text{C}$. An output resolution of 16 bits was thus chosen, which is not achievable using nyquist-rate converters and poorly-controlled passive on-chip components. Delta-sigma ($\Delta\Sigma$) ADCs are an ideal choice for this application as they perform many coarse quantisations of the input value to produce one digital output value. This technique allows a trade-off between speed and accuracy, allowing the converter to achieve high resolution with a simple 1-bit quantiser [8].

The on-chip $\Delta\Sigma$ ADC occupies the largest proportion of chip area, so it presents more opportunities for area savings than the analog front-end. Area-saving circuit techniques will be presented in Sections 4 & 5.

3 MANUFACTURING PROCESS

The primary goal for this work is to minimise costs by producing a smart sensor featuring a high level of integration and reduced die area compared to previous designs. The IC manufacturing process significantly influences the layout area of a design, and should be chosen with care.

A smart temperature sensor requires both analog and digital circuitry to perform its intended function. Advanced IC processes with reduced minimum line width implement digital logic functions at higher speed while consuming less power and die area than previously possible. Analog circuitry however, does not physically scale at the same rate as digital logic, while secondary considerations such as reduced supply voltage and short-channel effects complicate the migration path to a newer process. The optimal process choice for mixed-signal

designs therefore requires the considered balancing of opposing criteria.

Recent smart sensor designs in the literature [2, 4] were implemented on 0.5 μm and 0.7 μm processes respectively. While this allowed the analog sections to achieve good performance with reasonable voltage headroom, the same functionality could be achieved with reduced die area by moving to a process with a lower minimum feature size. On the other hand, analog circuits typically require accurate passive components – resistors, capacitors, inductors – that are not available on the newest processes targeted at digital applications. Process extensions that provide such passive components take time to develop, thus analog designs requiring these passive devices are only possible on processes a generation or two behind the leading edge.

An example of such a modern process offering analog extensions is a 0.13 μm IBM process targeted at RF applications. This offers precision polysilicon resistors and linear metal-insulator-metal (MIM) capacitors to allow the implementation of accurate analog circuits, while the 0.13 μm minimum feature size reduces the layout area required by digital logic. This 0.13 μm process was chosen as the target for the smart sensor design presented in this paper.

As described in Section 2.1, the temperature sensor uses the characteristics of bipolar transistors to produce the temperature signal and reference. Bipolar transistors are not used in modern CMOS designs, so it is necessary to verify the ability to manufacture such devices on the target process. Fortunately many bandgap reference circuits are still commonly used in mixed-signal designs, so a substrate PNP bipolar device is supported and modelled in the target IBM process. However the diffusions in modern manufacturing processes are not optimised for creating substrate transistors, so these devices perform relatively poorly in advanced submicron processes [9, Chapter 8]. As a substrate bipolar’s collector is tied to substrate (usually the most negative potential), such devices must be biased through their emitter. And because the bias current is defined at the collector (1), the bipolar’s current gain β becomes a quantity of interest. Shallow wells and extremely thin source / drain regions contribute to low β in submicron processes, and the IBM-provided models predict a substrate β of around 2. Such a low value requires the use of circuit techniques to allow the value of I_C to be accurately controlled [4].

4 $\Delta\Sigma$ MODULATOR

A $\Delta\Sigma$ ADC consists of an analog modulator and a digital filter. The modulator consists of a coarse quantiser and one or more integrators arranged in a feedback loop. Many coarse samples of the input signal are taken, and quantization noise introduced by the sampling process is modulated to frequencies above the signal band by the noise transfer function (NTF).

$\Delta\Sigma$ modulators designed for modern mixed-signal ICs are invariably implemented using discrete-time differential switched-capacitor circuitry. Fully differential circuits offer increased signal swings and greater immunity to external noise sources than single-ended designs, both of which are significant advantages in a low-voltage mixed-signal IC [10, Chapter 6]. Switched-capacitor circuits exhibit good linearity and well-defined gain due to the excellent matching of on-chip capacitors. The discrete-time nature of their operation also eliminates sensitivity to clock jitter [8, Chapter 6].

A second-order modulator was chosen as a compromise between circuit complexity and conversion time, as it can achieve the desired 16-bit resolution with several hundred clock cycles. The loop can be implemented with several different architectures [8]. A Cascade of Integrators Feedforward (CIFF) structure was chosen as it can implement more complex NTFs without placing additional load on the sensitive input signals (Fig. 3).




Fig. 3. Second-order CIFF loop structure used as the $\Delta\Sigma$ modulator.

4.1 Selection of NTF

The values of the coefficients a_1 and a_2 in Fig. 3 determine the characteristics of the noise transfer function. The NTF affects the magnitude of quantization noise within the loop, which in turn affects loop stability and the number of cycles required to achieve a given resolution. The simplest NTF for a second-order loop is the differential function:

$$NTF(z) = (1 - z^{-1})^2. \quad (5)$$

The high-frequency peak in (5) can be limited by introducing appropriately-chosen poles. A commonly-used rule of thumb [8, Chapter 4] is to limit the maximum gain of the NTF to a suitable value, such as 1.5, notated as $\|NTF\|_\infty = 1.5$. The NTF is therefore implemented in the form:

$$NTF(z) = \frac{(1 - z^{-1})^2}{D(z)}, \quad (6)$$

where $D(z)$ can implement the poles of any high-pass filter, such as Butterworth.

Limiting the peak NTF value in this way has the effect of reducing signal swings within the modulator. This is desirable because each integrator’s gain must be scaled to ensure the integrator output remains within the range determined by the limited supply voltage. The capacitors used at the input to the first integrator are often the largest in the modulator to minimize the effect of sampled thermal noise, so reducing integrator swings reduces the

need for large integrating capacitors, thereby saving layout area. On the other hand, limiting the peak NTF magnitude also increases the proportion of noise present in the signal band, reducing the achievable conversion resolution.

Because a primary design goal is to minimize chip area and because a temperature signal is effectively a DC value, the NTF of (6) was implemented to reduce capacitor layout area. With a peak NTF gain of 1.5, (6) resolves to [8, Chapter 8]:

$$NTF(z) = \frac{(1 - z^{-1})^2}{1 - 1.225z^{-1} + 0.4415z^{-2}}. \quad (7)$$

4.2 Area-Efficient Integrator Gain

To create the ADC reference voltage V_{REF} of (4), it is necessary to multiply ΔV_{BE} by the factor α . The characteristics of the substrate bipolars and the bias currents used in the analog front-end dictate that a gain of $\alpha = 16$ is required. To minimize errors introduced by additional circuitry, it should be implemented in the first switched capacitor of the analog modulator (Fig. 3). This gain can be implemented either by switching in larger sampling capacitors, or by integrating ΔV_{BE} multiple times. In the interests of saving layout area the latter option was chosen. As the ADC reference depends on this gain, it must be implemented as precisely as possible.

Fig. 4 shows the switched-capacitor integrator used in [4] as the first loop integrator. This integrator uses correlated double-sampling (CDS) to eliminate offsets and low-frequency noise at the inputs to the amplifier [11]. CDS operation is achieved by inverting V_{IN} between sampling phase ΦA and integrating phase ΦB . The integrator's differential output toggles between zero and the integrated value V_{INT} between switch phases.

The switches in Fig 4 are implemented using NMOS devices, which invariably introduce small parasitic capacitances between the switched terminals and chip substrate (Fig. 5). This is a concern in a high-accuracy, high-cycle integrator as these parasitic capacitances will cause a small amount of charge to leak off C_{INT} every time the integrator toggles between ΦA and ΦB .




Fig. 4. Switched-capacitor integrator using correlated double-sampling.




Fig. 5. Parasitic capacitances of a MOS switch.




Fig. 6. Improved CDS integrator with minimal charge leakage.

An improved CDS integrator is shown in Fig. 6. The input signal V_{IN} is still inverted between phases ΦA and ΦB . But unlike the integrator of Fig. 4 the output voltage remains at V_{INT} during ΦA . This eliminates leakage due to parasitic capacitances in the switches, and allows the integrator to sample ΔV_{BE} 16 times with high precision.

5 DECIMATION FILTER

Conventional $\Delta\Sigma$ modulators designed for telecommunications and audio applications operate on a continuous waveform. The spectral characteristic of these converters is a primary performance measure, and less attention is paid to DC characteristics such as linearity and offset. The ADCs used in low-frequency instrumentation and measurement applications, however, do not process an uninterrupted waveform. Converters in these applications are required to accurately reproduce a near-DC signal in the digital domain. Each sample is independent from previous values, and performance measures of interest are linearity, gain, and offset.

To eliminate the influence of previous values on the current conversion, the integrators and digital filter of a $\Delta\Sigma$ ADC must be reset prior to the commencement of conversion. This is known as *incremental*, or *one-shot* operation. Because an incremental $\Delta\Sigma$ ADC must operate for the total impulse length of the filter before a digital output value can be obtained, the order of the digital filter is an important design variable that affects conversion time as well as output resolution.

Earlier research into incremental $\Delta\Sigma$ ADCs [12] used time-domain analysis to determine that the optimal filter is non-symmetric, and of the same order as the modulator. More recent research [13, 14] confirms this observation. A non-symmetric filter is known as a *Cascade of*

Integrators (CoI); the transfer function of a second-order CoI filter is:

$$H(z) = \frac{1}{N^2} (1 - z^{-1})^{-2}, \quad (8)$$

where N is the oversampling ratio. Due to their non-symmetric impulse response, filters of this type cannot suppress periodic noise originating from the input signal or within the modulator itself.

A symmetric impulse response is required for this smart sensor design for two reasons. Firstly, a symmetric response will provide an average of the chip temperature during the entire conversion. And secondly, chopping [11] can be used within the $\Delta\Sigma$ modulator to remove residual switch-induced offsets [4]; a symmetric filter is necessary to average these chopped offsets present at the modulator output.

An efficient and popular filter for $\Delta\Sigma$ applications requiring a symmetric impulse response is the *sinc* filter [15]. The analysis presented in [13] showed that a second-order modulator combined with a third-order sinc filter produced a 16-bit result with around half the number of cycles required by the same modulator with a second-order sinc filter. The transfer function of this third-order sinc filter is:

$$H(z) = \frac{1}{N^3} \left(\frac{1 - z^{-N}}{1 - z^{-1}} \right)^3. \quad (9)$$

To confirm the results in [12], the output of a second-order modulator implementing the NTF of (7) was filtered with both second- and third-order sinc filters of various lengths. The resulting conversion error was converted to an effective number of bits (ENOB), and plotted against the number of cycles required to fill the decimation filter (Fig. 7). The symmetric third-order filter achieves the target resolution of 16 bits with around 500 cycles, while the second-order filter fails to reach this accuracy at 800 cycles. Thus a symmetric third-order filter is clearly the better choice for decimating the output of a second-order $\Delta\Sigma$ modulator where a symmetric impulse response is required.




Fig. 7. Plot of output resolution vs. total modulator cycles for symmetric second-order and third-order decimation filters.

6 LAYOUT RESULTS

The analog sensor and reference generator were laid out on the target 0.13 μ m IBM process, along with the second-order switched-capacitor $\Delta\Sigma$ modulator and third-order sinc filter. The simple structure of the sinc filter combined with the high digital density possible on the target process enabled the filter to occupy minimal area. High-density metal-insulator-metal capacitors combined with the accurate multi-counting integrator of Section 4.2 significantly reduced the area requirements of the $\Delta\Sigma$ modulator. Performing the required 16x multiplication with increased capacitor values alone would require a layout area approximately 16x greater for the capacitors of integrator 1; it is apparent from Fig. 8 that this would lead to a vastly larger IC.

The total area used by the sensor's active circuitry is 0.21mm². Prototype devices are being manufactured at the time of writing.




Fig. 8. Layout of the complete smart temperature sensor.

7 REFERENCES

- [1] G. Meijer, A. Boomkamp, and R. J. Dugesnoy, "An accurate biomedical temperature transducer with on-chip microcomputer interfacing," *IEEE J. Solid-State Circuits*, vol. 23, no. 6, pp. 1405-1410, Dec. 1988.
- [2] A. Bakker and J. H. Huijsing, "Micropower CMOS temperature sensor with digital output," *IEEE J. Solid-State Circuits*, vol. 31, no. 7, pp. 933-937, July 1996.
- [3] A. L. Aita, M. Pertijs, K. Makinwa, and J. H. Huijsing, "A CMOS smart temperature sensor with a batch-calibrated inaccuracy of $\pm 0.25^\circ\text{C}$ (3σ) from -70°C to 130°C ," in *IEEE Int. Solid-State Circuits Conf.*, San Francisco, 2009, pp. 342 - 343.
- [4] M. A. P. Pertijs, K. A. A. Makinwa and J. H. Huijsing, "A CMOS Smart Temperature Sensor With a 3σ Inaccuracy of $\pm 0.1^\circ\text{C}$ From -50°C to 120°C ," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, Dec. 2005.
- [5] R. J. Widlar, "New developments in IC voltage regulators," *IEEE J. Solid-State Circuits*, vol. 6, no. 1, pp 2-7, Feb. 1971.

- [6] A. P. Brokaw, "A simple three-terminal IC bandgap reference," *IEEE J. Solid-State Circuits*, vol. 9, no. 6, pp. 388-393, Dec. 1974.
- [7] W. Rasmussen, "Thermal sensors," in *Sensor Technology and Devices*, L. Ristic. MA, USA: Artech House, 1994, pp. 287-291.
- [8] R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*. Piscataway, NJ, USA: IEEE Press, 2005.
- [9] A. Hastings, *The Art of Analog Layout*. NJ, USA: Prentice Hall, 2006.
- [10] D. Johns and K. Martin, *Analog Integrated Circuit Design*. USA: John Wiley & Sons, 1997.
- [11] C. C. Enz and G. C. Temes, "Circuit Techniques for Reducing the Effects of Op-Amp Imperfections: Autozeroing, Correlated Double Sampling, and Chopper Stabilization," *Proc. IEEE*, vol. 84, no. 11, pp. 1584-1614, Nov. 1996.
- [12] J. Robert and P. Deval, "A second-order high-resolution incremental A/D converter with offset and charge injection compensation," *IEEE J. Solid-State Circuits*, vol. 23, no. 3, pp. 736-741, June 1988.
- [13] J. Márkus, J. Silva and G. C. Temes, "Theory and applications of incremental $\Delta\Sigma$ converters," *IEEE Trans. Circuits Systems I*, vol. 51, no. 4, pp.678-690, April 2004.
- [14] J. Márkus, P. Deval, V. Quiquempoix, J. Silva and G. C. Temes, "Incremental Delta-Sigma Structures for DC Measurement: an Overview," in *IEEE Custom Integrated Circuits Conf.*, San Jose, 2006, pp. 41-48.
- [15] E. B. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Trans. Acoustics, Speech and Signal Processing*, vol. 29, no. 2, pp. 155-162, April 1981.

A toy example of electrical impedance imaging using computational inference

Erfang Ma and Colin Fox

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

Abstract: We present an example of reconstruction in a toy example of electrical impedance tomography. Reconstructed images are computed using Bayesian inference by Markov chain Monte Carlo (MCMC) sampling implemented using a single-site Metropolis-Hasting algorithm.

Keywords: Inverse problem, Bayesian inference, EIT, MCMC

1 INTRODUCTION

Figure 1 shows a phantom conductivity that we use to generate a toy example of imaging. The imaging problem is to recover the spatially varying electrical conductivity from boundary measurements of current and voltage, often called electrical impedance tomography (EIT). The phantom occupies a square region, and has pixel-wise constant values of 4 (black) and 3 (white) in arbitrary units of conductivity. Measurements are made at 16 point electrodes (i.e. of no width) on the boundary (see Figure 2). This is same setup as used in [8].

In EIT, we normally inject a set of currents $\{\mathbf{j}_E^i, i = 1, 2, \dots, k\}$ in different pattern through these electrodes and measure the resulting potentials at electrodes $\{\mathbf{u}_E^i, i = 1, 2, \dots, k\}$. The set of measurements is the current-voltage pairs $\{(\mathbf{j}_E^i, \mathbf{u}_E^i), i = 1, 2, \dots, k\}$. The imaging problem is to estimate the conductivities from these measurements [5].




Figure 1: The phantom conductivity used in the toy example of EIT




Figure 2: Diagram showing a point electrode

2 MCMC

In the Bayesian approach to EIT, the conductivity σ is treated as a random variable. The *posterior probability distribution* over σ , given measurements of current and potential $f(\sigma | \mathbf{u}_E, \mathbf{j}_E)$, is the focus of subsequent inference [5] [3].

Here, we employ a single-component Metropolis-Hastings algorithm [7] [9] to explore the posterior distribution and to obtain a set of samples $\sigma^1, \sigma^2, \dots, \sigma^n$ from this multivariate distribution. The algorithm is comprised of iterations $\sigma^k = (\sigma_1^{(k)}, \sigma_2^{(k)}, \dots, \sigma_m^{(k)}) \rightarrow \sigma^{k+1}$ each of which involves the following three steps

$$\mathbf{1} \quad \sigma^k \rightarrow \sigma' = (\sigma_1^{(k)}, \sigma_2^{(k)}, \dots, \sigma_{i-1}^{(k)}, \sigma'_i, \sigma_{i+1}^{(k)}, \dots, \sigma_m^{(k)})$$

Here $\sigma'_i = \sigma_i^k + z_i$, with $z_i \sim N(0, \sigma_{z_i})$

2 compute acceptance probability $a(\boldsymbol{\sigma}, \boldsymbol{\sigma}') = \frac{f(\boldsymbol{\sigma}' | \mathbf{u}_E, \mathbf{j}_E)}{f(\boldsymbol{\sigma} | \mathbf{u}_E, \mathbf{j}_E)}$

3 accept update $\boldsymbol{\sigma}^{(k+1)} = \boldsymbol{\sigma}'$ with probability $\min\{1, a(\boldsymbol{\sigma}, \boldsymbol{\sigma}')\}$.

Note that the normalization constant in the posterior distribution is not needed for the computation of acceptance probability, since

$$f(\boldsymbol{\sigma} | \mathbf{u}_E, \mathbf{j}_E) = \frac{f(\mathbf{u}_E | \boldsymbol{\sigma}, \mathbf{j}_E) \cdot f(\boldsymbol{\sigma})}{\int f(\mathbf{u}_E | \boldsymbol{\sigma}, \mathbf{j}_E) \cdot f(\boldsymbol{\sigma}) d\boldsymbol{\sigma}}$$

hence,

$$a(\boldsymbol{\sigma}, \boldsymbol{\sigma}') = \frac{f(\mathbf{u}_E | \boldsymbol{\sigma}', \mathbf{j}_E) \cdot f(\boldsymbol{\sigma}')}{f(\mathbf{u}_E | \boldsymbol{\sigma}, \mathbf{j}_E) \cdot f(\boldsymbol{\sigma})} \quad (1)$$

The prior distribution $f(\boldsymbol{\sigma})$ here [6] is

$$f(\boldsymbol{\sigma}) \propto \exp \left\{ \beta \sum_{i \sim j} s(\sigma_i - \sigma_j) \right\}, \quad \boldsymbol{\sigma} \in [2.5, 4.5]^m \quad (2)$$

where $s(\cdot)$ is the tricube function [2]

$$s(d) = \begin{cases} \frac{1}{d_0} (1 - |d/d_0|^3)^3 & \text{if } -d_0 < d < d_0 \\ 0 & \text{if } |d| \geq d_0 \end{cases}$$

The notation $i \sim j$ in the summation in equation 2 indicates an edge in the graph of Markov random field, as depicted in Figure 3. Figure 4 shows a realization from this prior distribution. We can see that this model allows abrupt changes of intensity in the image, which is an important feature for this application.




Figure 3: Graph of first order Markov random field corresponding to the pixel lattice

Evaluating the likelihood term $f(\mathbf{u}_E | \boldsymbol{\sigma}', \mathbf{j}_E)$ involves solving a boundary value problem (BVP) in order to compute the forward map.

3 FORWARD MAP

If noise in the forward map is modeled as

$$\mathbf{u}_E = \mathbf{u}_E^0(\boldsymbol{\sigma}, \mathbf{j}_E) + \mathbf{n} \quad (3)$$




Figure 4: A realization from the prior showing typical properties

where $\mathbf{n} \sim N(0, \sigma_n \mathbf{I})$, then the likelihood function is

$$f(\mathbf{u}_E | \boldsymbol{\sigma}, \mathbf{j}_E) = \exp \left\{ -\frac{1}{2\sigma_n^2} \|\mathbf{u}_E - \mathbf{u}_E^0(\boldsymbol{\sigma}, \mathbf{j}_E)\|_2^2 \right\} \quad (4)$$

up to a constant of proportionality that does not depend on $\boldsymbol{\sigma}$. Here $\mathbf{u}_E^0(\boldsymbol{\sigma}, \mathbf{j}_E)$ is potential at electrode $u|_E$, in which the potential u is the solution of the BVP

$$\begin{aligned} -\nabla \cdot \boldsymbol{\sigma} \nabla u &= 0 & \Omega \\ -\boldsymbol{\sigma} \frac{\partial u}{\partial n} &= j_E & \partial\Omega \end{aligned}$$

We use the finite element method (FEM) to solve this BVP. The weak form of the BVP is

$$\int_{\Omega} (\nabla \cdot \boldsymbol{\sigma} \nabla u) \cdot v = 0 \quad (5)$$

where v is the test function. Applying Green's identity we have

$$\int_{\Omega} \boldsymbol{\sigma} \nabla u \cdot \nabla v = \int_{\partial\Omega} \boldsymbol{\sigma} \frac{\partial u}{\partial n} \cdot v. \quad (6)$$

The associated quadratic form gives the forms for the components in the stiffness matrix \mathbf{K} and load vector \mathbf{f} in the FEM system

$$\mathbf{K} \mathbf{u} = \mathbf{f}. \quad (7)$$

They are

$$K_{ij} = \int_{\Omega} \boldsymbol{\sigma} \nabla \phi_i \nabla \phi_j = \sum_{e=1}^E \int_{\Omega_e} \boldsymbol{\sigma} \nabla \psi_i \nabla \psi_j$$

and

$$f_i = \int_{\partial\Omega} \boldsymbol{\sigma} \frac{\partial u}{\partial n} \phi_i = \sum_b \int_{\Gamma_b} \boldsymbol{\sigma} \frac{\partial u}{\partial n} \psi_i.$$

We use bilinear elements to interpolate within each (square) pixel. The 4×4 local stiff matrix is

$$\left[\int_{\Omega_e} \nabla \psi_i \psi_j \right] = \frac{1}{6} \begin{pmatrix} 4 & -1 & -2 & -1 \\ -1 & 4 & -1 & -2 \\ -2 & -1 & 4 & -2 \\ -1 & -2 & -1 & 4 \end{pmatrix}.$$

The force vector f depends on the current pattern injected. In this paper, we inject a current at one of the 16 electrodes and extract the current uniformly around the boundary. This procedure is repeated 16 times with each electrode taking turn as being an injector. So there is 16 sets of measurements overall, so we assemble the force vectors into the matrix F having 16 columns, with each column corresponding to a set of measurements. If I amount of current is injected and there are m number of nodes on the boundary, the nonzero element of the i th column of F , are at the entry of boundary nodes, and of value $1/m$, except for the entry of injector being $\frac{m-1}{m}$.

The steps required for a computation of the acceptance ratio $a(\sigma, \sigma')$ are illustrated in the diagram in Figure 5.




Figure 5: Diagram of computational steps for the acceptance ratio

4 NUMERICAL RESULTS

We started the chain at $\sigma_0 = (3, 3, \dots, 3)$, and didn't stop until $80,000 \times m$ simulations had been completed. During the time, we recorded the state of σ every 10 scans ($10 \times m$ iterations).

We also adjusted the step size σ_{z_i} in the proposal step to make sure that the acceptance rate for each pixel is within [30%, 70%]. Specifically, σ_{z_i} for the pixels near electrodes are set to be relatively small and get larger for the pixels towards the centre. The exception is those pixels at the

four corners, which have the biggest step size. An image of step size is show in Figure 6.




Figure 6: Image of step size as a function of pixel. Brightness is proportional to σ_{z_i}

It is easy for us to provide a rough estimate of burn-in from the trace plot of log-prior (Figure 7) and that of log-likelihood (Figure 8), which is about first 40,000 samples.




Figure 7: Trace of log-prior




Figure 8: Trace of log-likelihood

Four realizations after the burn-in, which are separated by 1,000 scans ($1,000 \times m$ iterations), are shown in Figure 9.




Figure 9: Four samples from the posterior distribution

Traces of the conductivity at three pixels in the image are shown in Figure 10. The location of the three pixels are depicted as red, blue, and green circles in Figure 14. We see that the sampler moves closely around desired low conductivity during the course, while it is jumping up and down from time to time in the dimension of the other two pixels. This is in accordance with the marginal distributions for these pixels (Figure 11 12 13). The marginal distribution for pixel red has only one mode, while the other two marginal distributions are bimodal.

All these results show that single-site Metropolis-Hasting algorithm did not do a bad job in sampling this high-dimensional distribution.

If any single estimator of conductivity is of interest, posterior mean (Figure 14) and marginal posterior mode (Fig-




Figure 10: traces of three pixels during the MCMC run




Figure 11: Sample marginal distribution for pixel red




Figure 12: Sample marginal distribution for pixel blue

ure 15) are two popular choices [4]. As we see in Figure 14, most of the uncertainty is about boundary of high conductivity region. The discrepancy between marginal posterior mode (Figure 15) and original one (Figure 1) happens to be around the same place.

However, note that 80,000 iterations had been done to achieve the above result, and it took quite a few days to run




Figure 13: Sample marginal distribution for pixel green




Figure 14: Posterior mean conductivity




Figure 15: Marginal posterior mode

this long simulation in Matlab. We also mention that about 99% of CPU time had been spent on solving the system of linear equation in FEM stage for the forward map.

5 CONCLUSION AND FURTHER WORK

Single-site Metropolis-Hasting algorithm can be used to sample the posterior distribution arising in Bayesian infer-

ence in EIT, as long as speed is not a top priority. Future work could focus on alternative MCMC algorithms which are more efficient and fast in this application, e.g. delayed acceptance MCMC [1].

REFERENCES

- [1] J.A. Christen and C.Fox, *MCMC using an approximation* J.Comput. Graph. Stat. 14 (2005), pp. 795-810.
- [2] W.S. Cleveland *Robust Locally Weighted Regression and Smoothing Scatterplots* Journal of the American Statistical Association, 74 (1979), 829-836.
- [3] C. Fox *Recent advances in inferential solutions to inverse problems* Inverse Problems in Science and Engineering Vol. 16, No. 6, September 2008, 797-810.
- [4] C. Fox and G. Nicholls *Exact map states and expectations from perfect sampling: Greig, Porteous and Seheult revisited*, in *Twentieth International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, A.M. Djafari, ed., AIP, Melville, New York, USA, 2000, pp. 252-263.
- [5] C. Fox and G. Nicholls *Sampling Conductivity images via MCMC* The art and science of Bayesian image analysis. Proceedings of the Fifth Valencia International Meeting, June 5-9, 1997
- [6] S. Geman. and D. McClure *Statistical methods for tomographic image reconstruction* Bulliten of the International Statistical Institute, 52, no.4 (1987) 5-21.
- [7] N. Metropolis, M. Rosenbluth, A. Teller, and E. Teller "Equations of state calculations by fast computing machines" *Journal of Chemical Physics*, 21, (1953) 1087-1091.
- [8] J. D. Moulton, C. Fox, and D. Svyatskiy *Multi-level Approximations in Sample-Based Inversion from the Dirchlet-to-Neuman Map*. Tech. Rep. LA-UR 07-7958, Mathematical Modeling and Analysis Group, Los Alamos National Laboratory.(2007)
- [9] A. F. M. Smith, G. O. Roberts *Bayesian Computation via the Gibbs Sampler and Related Markov Chain Monte Carlo Method* J. R. Statist. Soc. B 55, No.1 (1993) pp. 3-23

Low-cost contour and groundwater mapping

Colin Fox

Electronics Group, Department of Physics, University of Otago, PO Box 56, Dunedin, New Zealand

Email: fox@physics.otago.ac.nz

Abstract: Faced with the need to establish contours for a sloping quarter acre section, I developed a simple and low-cost way of making a survey from which accurate contours can be produced. The solution requires measuring the distance between suitably-placed survey pegs, and also their relative heights, followed by computation of a sizable inverse problem. A groundwater problem on a neighbouring property prompted me to also develop a low-cost piezometer for measuring the water level in bore holes. The surveying method is described with the resulting plan of contours, and the piezometer circuit and details are given.

Keywords: Contour mapping, surveying, inverse problem, groundwater, fluid sensor, piezometer

1. INTRODUCTION

In 2008 we¹ bought a quarter-acre section in North Dunedin. Being warm-blooded immigrants from the north we chose a section that catches lots of sun, which in Dunedin means that the section is on a hill side. (In Dunedin land on the flat is either in the valley, which gets the frost, or on the ridge, where it can be very windy.) A picture of the land, and slope is in Figure 1.



Figure 1: A sunny north-facing slope in Dunedin. The slope can be seen to be roughly 1 in 3.

With an intention to build, I decided to survey the land to provide contours to our architect. I first went the high-tech route using GPS that allows

post processing to sub-metre accuracy. That allowed absolute placement of the section pegs, but I found the process too slow and imprecise for the contouring job. In the end I settled on a low-tech approach using a spirit level made from a garden hose, a tape measure, and string lines to fill in detail as needed. Measurements are then quite quick and accurate, but leave a quite large implicit problem to determine the position of pegs and hence contours. The measurement equipment, procedure, data protocols, and structure of the inverse problem are covered in section 2.

Our house design would require excavation to a depth of 3 metres, so we became concerned about the possibility of groundwater flow to the house site, both during construction and after the house was built. We are aiming at building a well-insulated house; contrary to popular belief, Dunedin actually gets plenty of regular sun. So it is feasible to build a house that maintains an internal temperature around 20 degrees C, without any heating at all [1]². The key is thermal mass and insulation – the combination stores the heat from the sun for the days when there is none.

We had experienced this system when living in Guanajuato in México, which (contrary to popular belief) is often a cold place. There, the houses are built against the bedrock which stores the heat from the day, keeping the house warm during the very cold nights. Only after three days of overcast weather did our house actually need heating.

¹Andrea and I

²We will put in a log burner for those times when theory and practice don't match.

Guanajuato has the advantage of being incredibly dry, which aids thermal insulation.

Water flow against a building is not conducive to maintaining warmth (or dryness). To measure the groundwater on the site, I enlisted the help of a soil geologist friend. We bored three 3-5 metre deep bore holes and measured the level of water. The first measurements were made using an industrial water-level sensor that we borrowed. For further measurements I built a fluid-level sensor that give the same functionality, using bits out of my electronic junk box. The electronics and sensor are described in section 3.

2. LOW-COST SURVEYING

2.1 Low cost GPS surveying

My first intention was to survey the land solely using some cheap global positioning system (GPS) gear. Standard GPS equipment has an accuracy of about 10m, which is obviously not good enough for determining contours. Professional surveying GPS gear uses a differential system that requires a local base station to transmit it's own location, so that a moving station can subtract out most of the errors in the system. This still leaves something like 1 metre accuracy, but can be fast. However, these systems cost something like \$100k, so are well outside my budget.

Instead I opted for some affordable gear from DeLorme [2] that allows post-processing of the data to give sub-metre accuracy. The units I bought were an Earthmate Blue Logger GPS, and an Earthmate USB GPS, costing about NZ\$300 all up. That allowed me to log one GPS directly to my laptop, while the other logged itself, for later download via Bluetooth. The gear worked well, though with a few teething and software problems and the usual Bluetooth issues, and did indeed allow sub-metre measurements. I found that very useful for accurately locating boundary pegs with respect to the survey reference points in Dunedin. However, the post-processing essentially works by *averaging* a long record to reduce errors. In practice that means recording for an hour or two, at a time when many satellites are overhead. DeLorme provided great software for finding out in advance when a suitable satellite configuration was overhead, but that did mean planning measurements in advance to fit them around other commitments. In practice I could manage a few measurements per day which was simply too slow for the surveying job. On top of that, accuracy around 1 metre was not really good enough for the half-metre contours that the architect wanted. So I turned to the low-tech route, described next.

2.2 Low-tech surveying, in theory

The GPS route had created the hope that I could determine the absolute location of any point. With that easy route gone, I decided to measure the relative location of selected points by the tried-and-true method of triangulation. I fixed a set of points by putting in pegs at regular locations, with a clout-head nail in the top of each peg to give a precise measurement point.

To locate the pegs in three dimensions requires measuring the relative heights, as well as relative distances. Consider three points in space $r_1 = (x_1, y_1, z_1)$, $r_2 = (x_2, y_2, z_2)$, and $r_3 = (x_3, y_3, z_3)$. The relative heights $h_{1,2} = z_1 - z_2$, $h_{2,3} = z_2 - z_3$, and $h_{3,1} = z_3 - z_1$ and distances $d_{1,2} = \|r_1 - r_2\| = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2}$, $d_{2,3}$, $d_{3,1}$ are all that is needed. For these three points, the relative z -values are determined by $h_{1,2}$, $h_{2,3}$, and $h_{3,1}$ (note there is a redundancy since $h_{1,2} + h_{2,3} + h_{3,1} = 0$). The x and y values can be found by writing $(x_1 - x_2)^2 + (y_1 - y_2)^2 = d_{1,2}^2 - h_{1,2}^2$, etc, and solving for three of the six unknowns. The position of the three points is then determined up to an arbitrary translation in 3-dimensions, a rotation in the horizontal plane, and a mirror symmetry of the horizontal plane³.

The location of further pegs can be found by further triangulation. For example, if we include a peg at r_4 and measure distances to two previous pegs and one relative height, the three measurements determine the three coordinates (x_4, y_4, z_4) , up to the mirror symmetry.

The translation and rotation ambiguities may be resolved by choosing one peg to be at the coordinate origin (determining three values) and by determining the orientation of one line with respect to north. I did this at the final step, by determining the absolute location of two of the boundary pegs. The mirror ambiguity is resolved by knowing the approximate location of pegs, and picking the closest solution.

2.3 Measuring distance and height

Measuring distance is easy with a tape measure, but measuring relative heights requires a bit more ingenuity.

My first thought was to use a hose filled with water and measure the pressure drop across the hose, and hence determine relative heights of the two ends. I soon found that pressure transducers with

³There are *three* mirror symmetries, but the result depends only on whether an odd or even number of mirror operations are applied.

the required accuracy are a bit pricey, so I opted for a simpler approach. I made a simple spirit-level using a garden hose with about a metre of clear plastic hose attached at each end. A ball tap at each end allowed the ends to be closed so the water did not leak out while moving the hose, and could be opened to the atmosphere when making measurements – so the water at both ends have the same (atmospheric) pressure and hence are at the same height. I measured relative heights by mounting a pole of known height vertically over each peg, and measuring the height from the water level to the top of the pole. I squirted some dish-washing liquid into the water in the pipe to reduce any affects of surface tension. That had the unexpected benefit of creating foam on top of the water, making it easy to see the water level. The hose level is shown in Figure 2, with both ends clamped to a pole placed over a peg. The pole has a short section of pipe attached to the bottom, to allow easy placement over the peg, and a post-level attached to make it easy to ensure the pole is vertical.



Figure 2: The hose level made up of a water-filled garden hose with clear tube ends, with both ends clamped to a post. A red post-level is visible, while foam on top of the water is just visible.

I also used a string line between pegs to fill in surface levels, as needed. I used thick fishing line as the string line; being light and strong it can be pulled tight and does not droop much. Measurement of the distance from one peg deter-

mines the 3-dimensional location of a point on the string line, and the height above the surface was recorded.

2.4 Measurement protocol, parsing, and data structures

Measurements were initially written (in pencil) in a notebook, as they were made. In total, several hundred measurement were made and by the end I wished I had an electronic way of recording measurements without the need to write. I came up with a simple readable space delimited text file format for entering this data into the computer, that kept the process simple and easy to check.

All data was entered into one text file, with one measurement per line. The lines have the format

- H peg1 peg2 hhh – height hhh cm of peg1 above peg2
- D peg1 peg2 ddd – distance ddd cm between peg1 and peg2
- SL peg1 peg2 ddd hhh — the height hhh (cm) above ground of a string line a the point distance ddd cm from peg1

The terms `peg1` and `peg2` were the names given to the pegs. At the time of putting them in, each peg was marked with a short name such as ‘L1’ for pegs establishing basic levels, ‘BP4’ for pegs on the boundary, and ‘C3’ for fill-in pegs needed for more accurate contours.

The file was first parsed to build a table of peg names, and a mapping from names to numbers. Two arrays were constructed for each of the height and distance data. Array H was a measurement mask for heights with $H(i, j) = 1$ if the height between the i^{th} and j^{th} pegs was measured, and is otherwise zero, with the height array h storing the value of $h_{i,j}$. Note that H has even symmetry while h has odd symmetry. Similarly array D was a measurement mask for distances with $D(1, j) = 1$ if the distance between the i^{th} and j^{th} pegs was measured, and is otherwise zero, with the distance data array d storing the value of $d_{i,j}$. Note that D and d have even symmetry.

2.5 An inverse problem

Each of the raw distance or height measurements has an accuracy of about $\pm 0.5\text{cm}$, from reading the tape measure. Once geometry inaccuracies are accounted for, such as sag in a tape measure line or wobbly pegs, the errors are about 2 or 3 times

that, and maybe more in long distance measures where tape measure stretch is an issue.

These errors mean that the positions of pegs cannot be determined by exact solution of the equations, above. Attempting that route would probably end up with massive errors by the time the last measurements were being analyzed. Further, redundancy in the measurements (which ought to improve results) in the presence of errors means there is no set of peg positions that exactly satisfy the distance and height equations. So what to do?

This is a common situation in inverse problems. A standard solution is to solve *all* of the equations at once, but only approximately. And how to do that? The late great Murray D. Johns once said that solving any problem is easy ... just pick a criterion of optimality and optimize with respect to it⁴. So that is what we do here, by minimizing a sum of squares of the data residuals, as follows.

Using the arrays defined in section 2.4, define the square error in heights

$$Q_h = \sum_{H_{i,j}=1} (z_i - z_j - h_{i,j})^2$$

and the square error in distances

$$Q_d = \sum_{H_{i,j}=1} (\|r_i - r_j\| - d_{i,j})^2$$

which are functions of the peg positions. The equation in heights decouples so it is easy to solve for heights first. Since Q_h is quadratic in the heights, the set of z -values that minimize Q_h is given by solving a linear system involving the Hessian of Q_h which is (in MatLab notation) $-H + \text{diag}(\text{sum}(H))$ – it is no coincidence that this has the same form as the admittance matrix for a network of 1 Ohm resistors. Fixing these z values, I then solved for the x - and y -positions by minimizing Q_d . This functional is not quadratic in peg positions so requires a bit more work to minimize. I took the lazy route of using MatLab's `fmins` function, which takes ages to converge but does get there. It is important to resolve the translational and rotational symmetries to avoid unnecessary numerical work, and to choose a starting position that resolves the mirror symmetries.

Once peg positions were determined, I used a radial basis function interpolation to evaluate the ground level at any point. Figure 3 shows the contours determined this way, with the location of pegs shown as blue crosses, plotted on an aerial photograph downloaded from the Dunedin City web site.

⁴It turns out this suggestion of making a ‘point’ estimate can sometimes give terrible predictions. See the other paper co-authored by me in this proceedings for some examples that demonstrate this.

3. MEASURING GROUNDWATER

Underground water was probed by drilling bore holes, and measuring the water level in each bore hole over time.

With the help of a friend who is a soil geologist, we borrowed a 60mm diameter hand auger that had extensions to allow bore holes of several metres depth. In the end we bored three holes, two at the location that will be the back wall of the house (and the deepest excavation) and one near an area where a neighbour had experienced problems from excessive groundwater. The bore hole varied from 3.5 to 4.5 metres – basically we kept going until we hit a rock and could not bore further.

I put 30 mm diameter stand pipes in each bore-hole, with horizontal hacksaw cuts each 50mm for 1.5 m from the bottom to allow water flow, then covered in filter material to stop silting. The bore hole around the stand pipe was back-filled with coarse sand to stop the hole collapsing, which can also block flow. Figure 4 shows one stand pipe.



Figure 4: All that can be seen from the surface of one stand pipe

3.1 A junk-box piezometer

An easy way to measure water height is to use a ‘piezometer’. A sensor is lowered into the stand-pipe, and the piezometer indicates when the sensor touches water. The sensor head is essentially a pair of electrical contacts with the water closing an electrical circuit.




Figure 3: Contours, in green, interpolated from 3-dimensional peg positions, shown as blue crosses.

On the day we bored holes, we used a professional piezometer that we borrowed to measure the height of the groundwater. For subsequent measurements I built a unit that mimicked the operation of the professional unit. Figure 5 shows



Figure 5: The sensor head connected to a tape measure and wiring, sitting on the box that houses the electronics.

the sensor head I built, made from a stainless steel center rod with a 10mm long 3mm diameter probe turned at the bottom end, wrapped in insulating tape and shoved into a outer tube, with 10mm cross hole at the probe. The outer tube turned out not to be stainless, as can be seen from the surface rust.

Measurement of electrical connectivity at the sensor is best made at frequency, as opposed to d.c., to avoid electrolysis effects that can build up an insulating layer on the sensor. I chose a frequency of 1.2 kHz.

The electrical equivalent of the sensor at 1kHz is shown in Figure 6, when dry and when dipped in tap water. Not surprisingly, when dry the sensor




Figure 6: The electrical equivalent of the sensor head, wet and dry, measured on a RCL bridge at 1kHz.

is essentially an open circuit with the leads contributing about 430pF, so the impedance is primarily capacitive. When wet, the impedance at 1kHz is primarily resistive, looking like a 3kΩ resistor with some capacitance.

The water sensing circuit is shown in Figure 7, adapted from a circuit I found on the web [3]. This gives a device that gives an audible beep when




Figure 7: Circuit diagram for the piezometer.

the sensor is wet. It uses half of a CMOS 4093 quad two-input NAND schmitt trigger. I have to say that I'm a fan of this very versatile chip, and recommend the Fairchild application note [4]. I added the anti-static protection [5] after zapping the first version within a day. All the parts required were retrieved from my junk box. The speaker is a piezo transducer (probably scavenged from a musical Christmas card that had become annoying) glued to the inside of the box. The whole thing runs from a 9V battery and requires no on/off switch as it draws virtually no power when not beeping. With only the connecting terminals going through the box, I could seal the box completely making the whole unit weather proof. The unit I built is six months old now, and still beeps happily when the sensor is wet.

4. DISCUSSION AND CONCLUSION

I have now surveyed two plots of land using the method described here. One house is built using the resulting contours; unfortunately it's not ours. However, as we have shown, it is perfectly feasible to make contours of a quarter-acre plot with some \$100 worth of bits from the local hardware store, followed by several hours of high-tech computing courtesy of MatLab.

I was also surprised how easy it is to measure groundwater, and to build a piezometer. A hydrologist friend is now interested in replicating and using the design I have given here to save the thousands it costs to buy the professional unit.

As part of the groundwater measurements I also built a simple water pump for performing 'pump tests'. Electrical pumps that can suck muddy water up 4 metres are rather expensive, and rely on

a purpose-made non-return valve at the bottom of the bore. In keeping with the cheapo aspect of this project, I built a simple pump following a design I seen used in poor regions of Africa. It consists of a ball valve at the bottom of a long pipe; oscillating the pipe up and down pumps the water up. I made one using some 25 mm water pipe from the local hardware store, a few adapted fittings, and a marble. I was amazed how well it worked, and how good the valve turned out to be. This allowed me to perform a 'draw down' test to determine water availability – again requiring much processing in MatLab.

My conclusion is that a lot can be done with a few parts from the hardware store, an electronic junk box, and lots of MatLab processing.

5. REFERENCES

- [1] American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. "2005 ASHRAE Handbook Fundamentals", 2005.
- [2] www.delorme.com
- [3] Tony van Roon "Fluid-Level Sensor", <http://www.sentex.ca/mec1995/circ/sensor3.htm>, 1995.
- [4] Fairchild Semiconductor Corporation "CMOS Schmitt trigger – a uniquely versatile design component", Application Note AN-140, June 1975.
- [5] Fairchild Semiconductor Corporation "Electrostatic discharge prevention – input protection coircuits and handling guide for CMOS devices", Application Note AN-248, March 2003.

Precise and Highly Selective Tonal Frequency Tracking

Alex Opie^{*}, Dr. Michael Hayes^{*}, and Dr. Jim Cavers[†]

^{*}Electrical and Computer Engineering, University of Canterbury

[†]NZi3 Wireless Research Centre

Abstract: An algorithm for tracking the frequency of an interference tone is described and analysed. The method involves detection of the phase shift of a signal through a bandpass filter. The VHDL implementation of this method is described and simulation results are presented to show the operation of the system.

Keywords: EMI rejection, frequency tracking, VHDL

1 INTRODUCTION

Electromagnetic interference (EMI) can have a large impact on the operation of a circuit. As such, minimising susceptibility to EMI is a major design task in most electronic devices [1]. EMI can originate from both outside a circuit and within it, the latter being the focus of this paper. While it is common to solve the issue by placing shielding over the component(s) causing the interference [2] (usually a clock source), this paper looks at using digital filtering to remove the interference.

An algorithm for detecting the frequency of the interference and a notch filter for removing it have been developed. This paper focuses on tracking the interference frequency. This is important because a small amount of clock drift can result in the interference tone shifting away from the notch and escaping attenuation. The notch is narrowband to minimise damage inflicted on the wanted signal, so the frequency needs to be accurately tracked. Specifically, the notch has a bandwidth of ~ 7.5 Hz in a system with a 48 kHz sample rate.

Fig. 1 shows the basic structure of the system, comprised of a notch filter with adjustable centre frequency and a module to track the frequency of a tone in the incoming signal x . The notch filter is described in Section 2 and the tracker is the focus of this paper. An operational system also requires a means of determining the interference frequency initially, but this diagram only shows the components operating in steady state.

The notch filter and how the tracking unit fits into the system are described in Section 2, as well as a mention of alternative solutions. Section 3 describes the tracking algorithm, and Section 4 describes the implementation. Further applications of the tracker and recommended future research are described in Section 5.




Figure 1: Base level block diagram of the application.

2 BACKGROUND

Tonal interference can be reduced using a narrowband notch filter. A single-tap complex infinite impulse response (IIR) filter with adjustable centre frequency has been found to give good results. The structure of the notch is shown in Fig. 2. Its transfer function is given by

$$H(z) = \frac{z - \alpha}{z - \alpha(1 - \mu)}, \quad (1)$$

where α is a unit magnitude complex number whose angle determines the notch centre frequency and μ is a small real number proportional to the notch bandwidth.

The notch is supplemented by a differential detection algorithm to automatically determine the frequency of the interference at startup [3]. The notch and the detection algorithm were implemented in both Matlab and VHDL and were successfully tested on Altera Cyclone FPGA development boards. To complete the interference cancellation system, an algorithm to track the frequency of the interference in real time was required.




Figure 2: Notch filter structure.

Several solutions exist to the problem of frequency tracking. Tashev and Malvar [4] suggest using frequency domain techniques to predict the behaviour of interfering sinusoids in speech signals in order to perform cancellation. In [5], image processing techniques are applied to the spectrogram of a signal to restore a speech signal contaminated by tone(s). This technique is not appropriate for real time processing. Both of the above approaches have more complexity and require more computation than is desired. Also, neither of the above techniques allow for wanted tones to exist in a signal, rather, they treat all tones as interference to be removed.

In [6] and [7], features of orthogonal frequency division multiplexing are used to track the carrier frequency. Whilst they offer high-performing systems, the methods are not applicable to a general frequency tracking system. The method proposed in [8] based on dynamic signal delays is fast but does not offer enough accuracy. The methods proposed in [9] and [10] are too complicated for this application, with Markov models and artificial intelligence, and they would thus use far more resources than is justified for a small tone cancellation routine.

3 INTERFERENCE TRACKING

As shown in Fig. 3, the tracker operates by passing the incoming signal through a bandpass filter, then finding the phase shift that resulted. If the interference is located in the centre of the notch, the filter will apply no phase shift. But if the interference drifts away from the centre, a phase shift will be applied to the signal. This phase can be detected, and its sign indicates the direction of drift.




Figure 3: Tracker block diagram.

For a single tone input signal of amplitude A_x and digital frequency F_x , γ is given by

$$\gamma = |A_x|^2 H(F_x), \quad (2)$$

and so

$$\arg[\gamma] = \arg[H(F_x)]. \quad (3)$$

If the input signal contains both a wanted tone and an interference tone with amplitudes A_w and A_i and digital frequencies F_w and F_i respectively, the output becomes

$$\arg[\gamma] \approx \arg[H(F_i)|A_i|^2 + H(F_w)|A_w|^2]. \quad (4)$$

Comparing (4) and (3) it is clear that the effect on the tracker output of the wanted tone depends on the filter response at the frequency of the wanted tone as well as the relative magnitudes of the two tones.

For the general case, it can be shown that

$$\gamma = \langle v[n] \rangle \approx \int_{-0.5}^{0.5} H(F) S_x(F) dF, \quad (5)$$

where $\langle \cdot \rangle$ denotes a long-term average and S_x is the power spectral density of the signal x . This is equivalent to a weighted centroid operation.

3.0.1 First order vs second order

Two bandpass filters were tested, of first order and second order. Both filters have a phase response that diverges either side of its centre frequency and a magnitude which drops away with distance from the centre. These two features result in the filter producing a non-zero output when the interference tone is offset from the notch centre and reduce the effect of a wanted signal on the detection, respectively.




Figure 4: Frequency response of first order filter.

The first order filter has the transfer function

$$H(z) = \frac{z}{z - (1 - \mu)\alpha} \quad (6)$$




Figure 5: Frequency response of second order filter.

while the second order filter has the transfer function

$$H(z) = \frac{z^2}{(z - (1 - \mu)\alpha)^2}, \quad (7)$$

with α and μ as described in (1). The response of each of these filters is shown in Figs. 4 and 5.


The second order response has two advantages over the first order response. First, it has much greater attenuation and second, the response phase tends to $\pm\pi$ as it moves away from the centre¹, where the first order system tends to $\pm\frac{\pi}{2}$ (see Fig. 4). The greater attenuation is advantageous as it results in the tone nearest to the filter centre having more dominance over the detector output, reducing the effect of a wanted tone. The phase response is an advantage as it means that the contribution from the wanted tone on the detector output will lie approximately anti-phase with the contribution from the interference tone, minimising the error introduced. Both of these features can be seen in Fig. 6, where the error angle introduced by the wanted tone is significantly smaller with a second order filter.

Fig. 7 shows the error introduced to the phase of γ by the presence of a second tone in the signal. This shows that a wanted tone needs to be much closer to the notch centre frequency to have an effect on the tracking in the second order system than it does in the first order system.


While the above observations suggest that a second order filter offers better rejection of other energy in the signal, the first order system has the advantage that it can be implemented more simply and cheaply. However, the greater rejection offered by the second order filter outweighs this advantage. A third order filter was also tested, however, this offered minimal extra rejection. The reason for this is that with small separation, there are cross terms

¹Actually, the phase tends quickly to $\pm\pi$, then slowly falls back to reach zero again at a frequency $S/2$ from the filter centre. However, frequencies this far from the centre are neglected since their magnitude response is essentially zero.

which can no longer be neglected. For these reasons, a second order filter is considered to be the best for this tracking system.



(a) First order filter.



(b) Second order filter.

Figure 6: Argand plane representations of example values of γ produced by the two filter types. Note the smaller error with the second order filter (b).

4 VHDL IMPLEMENTATION


Once the second order filter phase detection method was chosen to be the best, it was implemented in VHDL. This required implementation of a first order narrowband band-pass filter (cascaded to produce a second order filter), an accumulator to approximate the expectation operation, and a unit to adjust α . Divisions and multiplication were implemented with Altera Megafunctions² to reduce development time. The implemented design is shown in Fig. 8.

The algorithm of the tracker for each input sample x_n is:


1. $w_n = x_n + 2(1 - \mu)\alpha w_{n-1} - (1 - \mu)^2 \alpha^2 w_{n-2}$ (2nd order bandpass filter)

2. $\gamma_n = \gamma_{n-1} + w_n \times x_n^*$

²Megafunctions are pre-designed modules provided by Altera for use on their FPGAs.



(a) First order filter.



(b) Second order filter.

Figure 7: Error induced by a wanted tone.




Figure 8: System diagram of the filter and tracking unit.

3. if accumulation time is up:

- if $\arg[\gamma] > \text{threshold}$:
 - $\alpha = \alpha \exp(-j\phi)$
 - normalise α using (8)
- else if $\arg[\gamma] < -\text{threshold}$:
 - $\alpha = \alpha \exp(j\phi)$
 - normalise α using (8)

Normalisation of α is required to ensure a unit magnitude despite quantisation of the real and imaginary components. This is accomplished by [11]

$$\alpha' = \frac{\alpha(3 - |\alpha|^2)}{2}. \quad (8)$$

The bandpass filter implementation uses a complex multiply, a bit shift, and several adds. The bit shift is used to effect the multiply by μ . The unit to perform the correlation step multiplies its two incoming signal samples together after conjugating one of them, and accumulates this product. After a set number of samples, the sum is latched and its angle is approximated using a division. This small-angle approximation holds as long as the interference frequency is near the notch centre frequency. If it is not, but the tone is still the main energy being passed through the filter, the quotient will have an inaccurate magnitude, but the correct sign, allowing the tracker to still move the notch in the right direction. This situation works well with the fixed step size α controller, since it can be difficult to determine how far α should be moved using this approximation to $\arg[]$.

The α adjusting module keeps a register containing the current value of α , used both internally and by the tracker and notch. When the module is signalled to adjust α , it rotates it in the direction signalled, then normalises the result using (8). This new result is then clocked into the register, updating the whole system.

4.1 Simulation

Functional simulation in Modelsim was performed for each module separately as it was developed, then the system as a whole was simulated. The filter simulation confirmed that the filter was operational by passing two tones through it, one inside and one outside the passband, and observing that one was attenuated.

Simulation of the α controller showed that it correctly responded to signals by rotating α , but downward movement occurred faster than upwards. This was due to the new value of α always being quantised downwards. This prompted a change to make the downward rotator smaller than the upward one. While this did not completely fix the issue, it is not a large problem since if the notch only needs to move down one step, then it will fall by two, then will move up one in the next iteration.

Having successfully simulated each component separately, the complete tracking system was assembled and simulated. For this simulation, α was initialised to a frequency of 1 kHz, and a 1 kHz tone was passed through the system. This was repeated for various other initial values of α around 1kHz. These runs showed the tracking system behaving exactly as expected, adjusting α until it matched the frequency of the incoming tone, then keeping it there. The same test was repeated for some other frequencies, such as 3 kHz and -4 kHz, to confirm operation.

Fig. 9 shows the tracker successfully adjusting the frequency to follow the frequency of the interference when the incoming signal contains wanted signals and noise. For this simulation, the signal had an interference-to-noise ratio of -30 dB, and the wanted tone was at 2970 Hz, with an amplitude about 3 dB greater than that of the interference. For the simulation, the interference tone started at 3000 Hz and drifted up at a rate of 0.2 Hz per second.

The step response of the tracking system is shown in Fig. 10. It can be seen that the system responds by changing the frequency at an approximately constant³ rate until it reaches the new reference frequency. The tracker could be described as a bang-bang controller with hysteresis. It moves up at a constant rate, down at a constant rate, or not at all. There is no variable rate control. The error present on the right hand side of the plot is not an offset error as would be found in proportional control; it is a quantisation error.

4.2 FPGA testing

To test the operation of the system on the FPGA, two more modules were required. The development board has an audio codec connected to line in and out jacks, used to take in and pass out the signals. Thus a driver was implemented to initialise and operate the codec. Since the system is designed to operate on I/Q signals from a radio channel, a

³The rate is nominally constant, but changes between steps due to quantisation effects.




Figure 9: Operation of tracker with interference tone and noise.




Figure 10: Step response of the tracker in simulation.

Hilbert transformer [12] was implemented to emulate this situation. For simplicity, the Hilbert transformer was implemented as a 15 tap FIR filter, which simplifies down to four coefficients. This offered good phase response, and the magnitude response was flat enough to observe the system working.

Two signal sources were used for testing, a signal generator and a laptop computer. The signal generator was used to represent an interfering tone, and the laptop could provide arbitrary signals such as music or a second tone to represent a wanted signal. Whilst the system was running, the value of α on the FPGA was monitored using SignalTap⁴, so that all changes could be seen. The system was operated with 48 kHz sampling. A hold-off period of approximately 350 ms was implemented at system startup to allow the filters in the ADC to stabilise. This length of time was determined through experimentation.

Tests began with just the interference tone as input to the system, to allow the detection routine to find the initial interference frequency. Then the ‘wanted signal’ was enabled, and the interference frequency was varied. For all tests undertaken, the tracker successfully tracked the interference frequency, adjusting the notch frequency to lie as close as possible to the interference within the bounds of quantisation. Some testing was done to see how small a separation between tones could be handled, and on average the system was unaffected by the presence of the second tone down to separations of 10 Hz, surpassing expectations.

5 CONCLUSIONS

This research has found that the best compromise between complexity and effectiveness is achieved by the use of a second order bandpass filter. This design offers good selectivity and is capable of tracking to a high level of precision.

⁴SignalTap is an Altera technology that allows signal values on an FPGA to be recorded and uploaded to a host PC.

The tracking system developed in this paper could easily be used in other applications. Instead of adjusting a notch filter, the system could be used to adjust a bandpass filter, for example to track CTCSS tones [13] or a carrier signal. Because the system is unaffected by out of band energy, several instances of this system can be run closely in parallel for tracking several tones within one signal. This would require a more sophisticated detection routine than the one used in this application. Future work is also required to quantify the effect of word lengths on the performance of the algorithm.

6 REFERENCES

- [1] X. C. Tong, *Advanced Materials and Design for Electromagnetic Interference Shielding*. CRC Press, 2009.
- [2] K. L. Kaiser, *Electromagnetic Compatibility Handbook*, p. 22.1. CRC Press, 2005.
- [3] Y. H. Kho and J. Cavers, "Killing the birdie: Suppression of tonal interference in radio receivers v1.0," Tech. Rep. WRC002-003, NZi3:WRC, January 2009.
- [4] I. Tashev and H. S. Malvar, "Stationary-tones interference cancellation using adaptive tracking," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 201–204, 2007.
- [5] B. I. Andia, "Restoration of speech signals contaminated by stationary tones using an image perspective," in *ICASSP*, pp. 61–64, 2006.
- [6] M. R. Dacca, G. Levin, and D. Wulich, "Frequency offset tracking in OFDM based on multicarrier PLL," in *21st Century Military Communications Conference*, pp. 912–916, October 2000.
- [7] M. Luise, "Carrier frequency acquisition and tracking for OFDM systems," *IEEE Transactions on Communications*, vol. 44, November 1996.
- [8] Z. Cheng and T. T. Tjhung, "Explicit frequency estimation and tracking using modulated Lagrange delay filter," in *Asia-Pacific Conference on Communications*, pp. 916–920, October 2005.
- [9] D. Van Cappel and P. Alinat, "Frequency line extractor using multiple hidden Markov models," in *OCEANS*, pp. 1481–1485, September 1998.
- [10] B. C. Phan and Y.-H. Hu, "New method for time-varying harmonic frequency tracking," in *ICASSP*, pp. 2228–2231, April 1988.
- [11] J. Cavers, "A low computation method for normalising the magnitude of a complex number," Tech. Rep. WRC027-002, NZi3:WRC, May 2009.
- [12] R. G. Lyons, *Understanding Digital Signal Processing*, pp. 361–380. Prentice-Hall, 2004.
- [13] D. Biddulph, ed., *Radio Communication Handbook*, p. 5.55. Radio Society of Great Britain, 1994.

OGRE Compute Client Architectures

Patrick Suggate, Dr T. C. A. Molteno*

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

Abstract: OGRE is grid middleware that connects desktops, workstations, clusters, and even other grids together as one large, heterogeneous grid. OGRE is designed for embarrassingly-parallel workloads and cycle scavenging from idle computers and clusters. Its network client/server model uses just the HTTP and HTTPS protocols and is designed to pass through firewalls with a minimum of open ports since the clients initiate all network connections and HTTP proxies can be used.

Keywords: Grid, High Throughput Computing, High Performance Computing, Gridware, Distributed Computing

1 INTRODUCTION

Many large institutions, the University of Otago being one of them, have a large number of desktop PCs, some workstations, and even a few clusters. These are often under-utilised, that is the average CPU usage is well below 100%, often with large periods where the computer is completely idle – like overnight. The combined computational power of all of these computers can be extremely large. Distributed and grid computing are attempts to combine the computational power of many smaller compute resources into a large combined compute resource.

1.1 Early Distributed Computing Projects

Several attempts to harness all of these available idle processors have been made before, with varying levels of success. The first popular distributed computing project was SETI@home. This project uses software running on thousands of PCs to analyse radio telescope signals for evidence of extra-terrestrial intelligence. Private users volunteer the idle cycles of their PCs to the project by downloading the SETI@home client software. This software runs as a screen-saver so only uses the computer while it is idle.

Folding@home[1], another example of volunteer computing, is the most powerful ‘supercomputer’ in the world. This is also distributed computing application where those willing to donate their spare CPU cycles can download and install the Folding@home client and it runs whenever the computer is idle. Estimated total computational performance of all of the compute clients that are part of the project is more than 8 petaFLOPS.

CPU's	Machine	Performance (GFLOPS)	Cost (2009 USD)
289,000	SETI@Home	724,000	0
329,000	Folding@Home	8,255,000	0
1792	IBM p575	34,000	9,000,000
150,152	Cray XT-5	1,059,000	100,000,000

Table 1: Volunteer computing processing power compared with more traditional supercomputers.

Table 1 compares the computing power of SETI@home and Folding@home with two more traditional supercomputers. As the table clearly shows, the cost to the volunteer computing projects for the compute resources is zero. The projects do have to supply all of the server hardware, network connections, manpower, and other infrastructure for the projects. A spin-off from SETI@home is BOINC[2], a C/C++ API that implements the necessary functionality for other projects wishing to develop distributed computing applications.

1.2 Grid Computing

Grid computing is combining the computing resources from multiple domains to form a much larger computing resource. Two prominent grid software projects are Sun's Grid Engine[3] and the University of Wisconsin's Condor[4]. These projects also include cycle-scavenging features so that the computing power of idle desktop machines can be added to the Grid.

Condor and Grid Engine have serious problems in some situations though. The master nodes have to be able to

connect to the compute clients to give them jobs. This requires that the master node must be able to initiate network connections to each of the clients. For such a connection to be established the client needs to have many network ports exposed and made available. This often requires opening up a lot of ports within any existing firewalls and this can be insecure, and relies on system administrators being willing to make these ports available. Condor provides tools to help configure its networking requirements but these are complicated and do not solve these problems as well as OGRE does.

1.3 OGRE Overview

OGRE is a grid-computing platform that is designed to excel at embarrassingly parallel workloads. It is primarily designed for cycle scavenging on heterogeneous clients, with all clients initiating network transactions, and huge job queues. OGRE is written in the Ruby programming language and uses standard interfaces for communication and storage. Amongst these standards are SOAP, XML-RPC, YAML, SQL, XML, SSH, HTTP, HTTPS and OpenSSL. The jobs are stored on the MasterNode in a SQLite3 database, job data files are stored on the DataServer in the native filesystem, though accessible via HTTP. Since all compute clients initiate connections, and can use standard SSH or HTTP ports, firewalls are not a problem with OGRE.

2 EXAMPLE OGRE APPLICATIONS

2.1 Bird Tags

A new GPS algorithm (currently in development a patent is pending) can calculate a GPS location fix with only a few milliseconds of GPS satellite data. An off-shoot of this technology is a new generation of bird-tags that are smaller, lighter, and with better battery life (see Figure 1 for a picture of a tag currently in development). Now, months or years of bird-position data can be recorded by each bird-tag.

A tradeoff with this algorithm is that the computational costs are far greater than with the conventional GPS algorithm. But by submitting jobs containing the data recorded by the tags to OGRE, the workload is spread over many, many processors so that these positions can be calculated in a reasonable time-frame.

2.2 Bioinformatics

OGRE has support for using some pre-installed software packages for computation. An example of this is the package R. This is a programming platform (interpreter + libraries) designed primarily for statistical computing and subsequent visualisation. If OGRE compute clients have an R package installed then they can run R jobs retrieved



Figure 1: Bird-tag that uses the GPS FastFix algorithm to improve battery life and reduce size and mass.

from the MasterNode. Examples of R jobs OGRE has been used to run is bioinformatics, gene-survival jobs.

2.3 Search for a Faster Counter

Conventional base-2 counters typically have a large combinatorial delay (due to the ripple-carry operation) or require a large number of logic gates[5]. By using non-standard number systems faster counters can be developed. A project that OGRE has been used for, requiring millions of CPU hours, is looking at gate combinations for faster binary counters. The brute force search had to search for gate combinations in highly-connected networks of gates and D-type flip-flops (see Figure 2).




Figure 2: Brute-force search for faster binary counters.

3 OGRE ARCHITECTURE

OGRE was designed to run embarrassingly parallel computing tasks distributed across a large, institute or nationwide grid. OGRE supports heterogeneous compute nodes,

i.e. compute nodes can have different architectures, operating systems, and installed software packages. Additionally, OGRE's network model allows it to communicate through firewalls far better than existing tools like Sun's Grid Engine or Condor.

The OGRE architecture uses a client-server model. Key features are that OGRE has a separate job server (called the MasterNode) and data server (called the DataServer). Each of these can be run on a separate computer. Currently the OGRE project is written almost entirely in the program language Ruby. This language is very suitable for Rapid Application Development and yet flexible and powerful enough to not to have to re-write sections of code in a lower-level, compiled language like C/C++.

OGRE has compute clients for GNU/Linux, PBS Pro, Torque, Apple Mac OS (via Xgrid), and Microsoft Windows¹. OGRE also allows jobs submitted to the MasterNode (see Section 3.2) to be run on any compatible platform without the user specifically choosing compute resources.




Figure 3: OGRE grid hierarchy showing firewalls, clients and clusters.

3.1 OGRE's Networking Code

One of OGRE's biggest advantages is its network model, OGRE does not lose functionality even when having to communicate through very restrictive firewalls. All clients must initiate network transactions to the MasterNode and DataServer. Since the servers do not initiate transactions the clients need not be directly accessible by the server. The clients can also communicate with the servers using

¹Some of these platforms are still in the early stages of development.

standard SSH and/or HTTP. These network services are normally enabled on many networks, and therefore the ports they require are normally open.

Condor and Grid Engine are also designed for grid computing but have features that are not very firewall friendly. They are built to support MPI and other parallel jobs as well as standard serial (or embarrassingly parallel) jobs. All compute nodes running parallel jobs need to be able to initiate and receive connections with the other compute nodes running the same job. This can be very difficult to setup the firewalls for with large distributed computing projects.

Embarrassingly parallel jobs do not have this problem though. OGRE, by focusing on this class of job, can form a huge grid spanning a large geographical areas and many network firewalls. The only networking requirement for an OGRE compute client is that it has either SSH or HTTP access to a OGRE MasterNode and DataServer.

3.2 MasterNode

An OGRE MasterNode can currently manage > 100,000 jobs on a single-core server using the SQLite3 database. The SQL code used within the OGRE MasterNode is generic SQL so the data-base back-end could be easily changed from SQLite3 to PostgreSQL or MySQL or any other SQL database².

3.3 DataServer

The DataServer stores data on its local file-system and a simple web-server is used for handling file operations across the network. Files are transferred using just the standard HTTP GET and PUT commands. The data server uses Ruby's WEBrick libraries and has built-in support for MD5 check-sums to ensure data integrity.

3.4 Compute Clients

OGRE compute clients either retrieve and run jobs locally or retrieve jobs and then pass them onto another job scheduler like PBS or Xgrid. The clients have to monitor the state of the compute resources to determine if current jobs have been completed, so the results can be uploaded, or if more jobs have to be fetched. A more thorough discussion of compute clients is presented in Section 4.

3.5 The Job Manager

Users can view the status of OGRE as well as submit and manage their jobs using the OGRE Job Manager. The Job Manager uses the Qt4 API (Application Programming Interface) as its GUI (Graphical User Interface) toolkit.

²This feature could be used to attach a PostgreSQL cluster for handling an even larger number of jobs.



Figure 4: The OGRE Job Manager presents a simple interface for an end-user to submit and manage their OGRE jobs.

4 COMPUTE CLIENT ARCHITECTURE

OGRE clients have been written for multiple platforms. The current list of supported client platforms are shown in Table 4³.

OGRE Client Platform	Status
Linux and other UNIX-like operating systems.	Stable
Torque an improved and FOSS version of PBS.	Stable
Xgrid (Mac OS-based grid-computing platform).	Beta
PBS Professional.	Stable
Sun Grid Engine.	Beta
Microsoft Windows.	Obsolete Planning

Table 2: OGRE compute clients listed by platform.

4.1 Cycle Scavenging

An OGRE grid is designed to scavenge processor cycles for idle computers. The GNU/Linux client achieves this by running at a very low priority and limiting the

³A new Microsoft Windows client is planned that will be written in C# but the Ruby version was working until it was discontinued due to a couple of issues.

OGRE client to using only a small percentage of the available RAM. This causes minimal perceivable performance penalty to the end-user. The Windows client is designed to run as a screen-saver but this is a work in progress as the old client has been obsoleted and the new client is still in the planning phase.

The PBS, Torque, Grid Engine, and Xgrid clients submit OGRE jobs to queues on the cluster that it is running. To avoid delaying higher-priority jobs OGRE checks to see if compute resources are idle before submitting a job to the cluster’s queue. The wall-time limit for OGRE jobs is typically set to a small value (currently 15 minutes) to allow a fast OGRE job turnover so that higher-priority, non-OGRE jobs do not sit for long in queues.

4.2 Network Connectivity

OGRE’s network model makes deploying OGRE far easier than with Grid Engine or Condor. OGRE clients initiate all network connections to minimise problems with firewalls (see Figure 5). These clients can work with even very restrictive firewalls. By default a compute client assumes that the MasterNode’s XML-RPC network port is 2004 and the DataServer’s upload and download port is 2002 . Both these protocols use HTTP so a HTTP proxy can be used if port 80 on the firewall is open. If not the proxy can tunnel, using SSH, through port 22 to the MasterNode and DataServer.




Figure 5: OGRE’s SSH-tunneling scheme to pass through restrictive firewalls.

4.3 Job Scheduling

There are two simple schedulers available that clients currently use, a query-based scheduler and a reservation-based scheduler. The reservation-based scheduler tries to reserve a slot in a queue before retrieving the job from the MasterNode. This is because with some cluster configurations it can be difficult to tell whether the queue is full or not. The other scheduler queries any queues to see if there a submitted job will run immediately.

Dynamic linear scaling in genetic algorithms for antenna design

P.J Williams *, T.C.A. Molteno

Abstract: In this paper we present a simple extension to the conventional linear fitness-scaling technique for genetic algorithms. This techniques, *dynamic linear scaling* is implemented, and compared with linear scaling on the automatic design of wire antennas using genetic programming. In these comparisons, dynamic linear scaling gives consistently better antenna designs for the same computational effort.

Keywords: Genetic algorithm, Antenna design, Linear scaling

1 INTRODUCTION

The goal of a genetic algorithm (GA), is to try and find a near optimal result in the most efficient way. A key component of a GA, is the strategies used in the selection process, these have a major effect on the performance of a GA [1, 2].

In antenna designs, the smallest difference in direction or size can result in a major effect on the performance on an antenna design, which when tested can result in a large performance change due to small variation. This can create a situation when a strong performing design, once discovered, can dominate the population of subsequent generations. This reduces the diversity within the population and compromises the performance of the GA.

Various fitness-scaling techniques can be used to prevent an individual having too much influence, the most widely used being linear-scaling (see Section 2). Linear scaling is very sensitive to the choice of scaling parameter C_m , and choosing a good value for this parameter is difficult. In this paper, we introduce a simple modification to linear scaling *dynamic linear scaling* and show that, for the case of GA optimisation of wire antenna, it results in improved performance.

2 LINEAR SCALING

Most GA in antenna design uses a scaling technique called linear-scaling [3, 4]. The scaled fitness, sF , is given by,

$$sF = \alpha \cdot J + \beta, \quad (1)$$

where J is the unscaled fitness,

$$\alpha = \frac{C_m - 1}{\delta} * J_{avg} \quad (2)$$

*Department of Physics. University of Otago. Box 56. Dunedin. New Zealand. pjw@physics.otago.ac.nz

$$\beta = J_{ave} * \frac{J_{max} - sF_{max}}{\delta} \quad (3)$$

$$sF_{max} = C_m * J_{ave} \quad (4)$$

$$\delta = J_{max} - J_{ave} \quad (5)$$

J_{max} is the strongest member in the population, J_{avg} is the population average. The user-defined parameter of the scaling C_m , is the number of expected copies desired of the strongest individuals in the next generation. This needs to be optimized for each individual problem.

C_m is fixed for the whole simulation. It has the flexibility where it can be used with small and large population, with nonlinear fitness landscape. Negative numbers can not be used and must be allocated a positive value before scaling.

2.1 Typical run with optimizing for RHCP antenna

In this example, a GA is run 10 times with a population of 1000 for 500 generations and results are shown in Figure 1. The goal is to optimize for a RHCP antenna [5]. The C_m value used in this example is 1.8. In Figure 2, a typical GA run is broken down into groups of 100¹, and the percentage of the population is picked from each group, at each generation, as shown². During the run, the GA performs well up until about generation 320, where a very strong individual is found. From that point on, only the top 200 individuals are used for crossover to create the next generation. This results in a limited gene pool and means that we are less likely to find the optimal design and that 300 individuals are being calculated each generation, but never having a chance of being selected.

Looking at results of many optimisation runs done using linear scaling, I have found that a optimal C_m value at

¹Where individual 1 is the best performing individual, 500 th individual is the worse performing.

²In Figure 2, the GA is run for 1000 generations to show the effect of a strong individual taking over the gene pool.




Figure 1: Typical results from 10 GA runs using linear scaling.




Figure 2: A single GA run showing the frequency individuals from each group is selected for crossover.

the beginning of a simulation is not the optimal in latter generations. And since no two runs are the same, a fixed C_m can not be optimal.

2.2 Fitness

The fitness, F , used in this paper is maximised for an antenna that has a uniform gain pattern across the hemisphere for a right hand circularly polarized (RHCP) signals. It is defined as

$$F = \sum_{(\theta, \phi)} (G(\theta, \phi) - \overline{G})^2 \quad (6)$$

where $G(\theta, \phi)$ is the gain and \overline{G} is the mean gain. The antenna simulation engine used is Nec2++ [6].

3 DYNAMIC LINEAR SCALING

A modification to current GA procedures, is that after all individuals in a population have been scaled, an extra step has been introduced, where depending on the strength of the top 50 percent of the population, will determine if the scaling factor C_m is altered. By increasing C_m value, a greater influence will be given to the strongest individuals, where a lower C_m value will give more influence to the

poorer performing individuals.

$$\varphi = \frac{\sum_{i=0}^{x/2} P[i]}{x} \quad (7)$$

$$if((\varphi/x) > \Lambda) \quad (8)$$

$$C_m = (C_m - 0.1) \quad (9)$$

$$if((\varphi/x) < \Upsilon) \quad (10)$$

$$C_m = (C_m + 0.1) \quad (11)$$

where P is the number of times the individuals was selected for crossover, φ is the total number of times the top 50 percent of the population is picked, x is the population size, Υ is the “up factor” and Λ is the “down factor”. For each generation, the value of C_m changes depending the previous crossover selections. Before the simulation starts, the user enters in a value for the thresholds Υ and Λ as percentages. For example, if Υ is 80 percent and Λ is 60 percent, then if the top half of the population selected for crossover accounts for more that 80 percent, C_m value is reduced by 0.1, and if the top half of the population selected for crossover accounts for less than 60 percent, C_m is increased by 0.1. C_m is not allowed to fall below 1.0 as this would scale the fitness so that weak individuals had a greater influence than strong individuals, and would lead to poor results.

4 DYNAMIC LINEAR SCALING PARAMETERS OPTIMIZATION

The goal of using dynamic linear scaling is that the value of C_m is not critical, the weight of a optimal design is on getting the “up factor” and the “down factor” correct. To optimal this, a broad search was conducted with a range of parameters. The goal was to find a set of parameters which constantly perform well and would be robust over different populations. The parameters where optimized for a population of 500 over 1000 generations. The optimal parameters were found to be 0.65 for “up factor” and is the 0.75 for the “down factor”. A run shown in Figure 4, shows 10 runs with these parameters. Figure 5 shows a plot of fitness and C_m vs generation for the same parameters.

5 RESULTS

To check the change that the dynamic linear scaling had on the performance of the GA, the GA run shown in Figure 2 was rerun with dynamic linear scaling and shown in Figure 3. By looking at the graphs, it can clearly been seen that more of the population contributes to the next population.

Another set of graphs which shows the difference made by the introduction of the dynamic linear scaling are Figure 6 and Figure 7. In this case where there is a fixed value for C_m (see Figure 6), the individuals selected for




Figure 3: GA run with population selected for crossover shown in groups.




Figure 4: 10 GA runs with dynamic linear scaling.

crossover are balanced to start of with, but over time, the best individual become very stronger compared to the rest of the population, and the linear scaling gives them a very high chance of being selected, so the weaker individuals are not selected latter on in the simulation, which gives only a limited gene pool to find new individuals. In the case where dynamic linear scaling has been use (see Figure 7), a greater number of individuals are selected for crossover, over all generations, due to the ability to reduce and increase the influence of the strong and weak individuals.

Finally, the results shown in Figure 1 vs Figure 4 (Dynamic Linear Scaling), shows that in this set of runs, the best individual had a fitness of 0.3, where with dynamic linear scaling, the same parameters produced a individual with a fitness of over 1.2.

6 CONCLUSION

Using the dynamic linear scaling, our GA has produced better individuals and consistency is also improved, this is probably due to a larger fraction of the population contributing to successive generations. This effectively leads to a greater range of genes exist in the population pool, resulting in better performance overall. An individual with optimised fitness (as define in Section 2.2) produced using a genetic algorithm with dynamic linear-scaling is shown




Figure 5: GA with dynamic linear scaling showing how C_m changes over the generations.




Figure 6: A population of 500, over 1000 generations with a fixed value for C_m (3.2).




Figure 7: A population of 500, over 1000 generations with a moving value for C_m (starting at 3.2).




Figure 8: Best individual after 1000 generations with dynamic linear scaling. The GA has ‘discovered’ a helix-like structure for the reception or RHCP signals.

in Figure 8. This was the best individual after a population of 500 individual antennas was evolved for 1000 generations.

REFERENCES

- [1] D. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, 1989.
- [2] M. Michalewicz, Zbigniew; Schoenauer, “Evolutionary algorithms for constrained parameter optimization problems,” *Evolutionary Computation*, vol. 4, no. 1, 96.
- [3] J. Johnson and Y. Rahmat-Samii, “Genetic algorithms in engineering electromagnetics,” *IEEE Antennas and Propagation Magazine*, vol. 39, pp. 7–21, August 1997.
- [4] Z. Michalewicz, *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, 1992.
- [5] P. Williams and T. Moltano, “A Comparison of Genetic Programming with Genetic Algorithms for Wire Antenna Design,” *International Journal of Antennas and Propagation*, 2008.
- [6] T. Moltano and P. Williams, “NEC2++: High-performance numerical electromagnetics code,” in *Proceedings of the 13th Electronics New Zealand Conference*, pp. 283–287, November 2006.

Design of a FPGA-based multi-channel front-end for synthetic aperture sonar

B. C. Bonnett and M. P. Hayes*

*Acoustics Research Group,
Dept. of Electrical and Computer Engineering,
University of Canterbury,
Private Bag 4800, Christchurch, NZ
Email: michael.hayes@canterbury.ac.nz

Abstract: This paper details the redesign of the transmitter and receiver electronics for the KiwiSAS-IV synthetic aperture sonar operated by the Acoustics Research Group at the University of Canterbury. The new system aims to reduce the size and power consumption of the electronics with the long term aim of mounting the sonar on an autonomous underwater vehicle. Class D amplifiers are used for the transmitted signal, and a highly integrated receiver IC is utilised with the necessary decimation and other echo processing being performed by an FPGA.

Keywords: synthetic aperture sonar, class D, CIC filter

1 INTRODUCTION

Optical methods are often unsuitable for obtaining images of the seafloor due to the high attenuation of light in seawater; sonar mapping of the seafloor is a commonly used replacement. Side-looking sonar, typically preferred over downward-looking sonar due to its increased range, requires the use of low frequencies as seawater attenuates high frequency sound energy. This in turn requires impractically large apertures in order to obtain a reasonable resolution in the resulting image; for example, a 2 cm resolution at a range of 200 m requires an aperture of around 10 m. Synthetic aperture sonar (SAS) gets around this by moving a small aperture receiver with a wide beamwidth along the scene taking overlapping readings; these can then be coherently processed to synthesise a suitably large aperture.

The KiwiSAS-IV synthetic aperture sonar operated by the Acoustics Research Group at the University of Canterbury consists of a 36-element projector array and a 9-element hydrophone array mounted in a neutrally buoyant towfish. The transmitted signal is made up of two superimposed linear chirps, one between 20 kHz and 40 kHz and the other between 90 kHz and 110 kHz; in theory, the low frequency chirp will penetrate mud on the seafloor better while the high frequency chirp will give a higher accuracy [1]. By using pulse compression on the echo signal, a resolution of approximately 2 cm can be achieved by the

KiwiSAS-IV system.

This paper details the redesign of the front-end electronics — that is, the electronics that drive the transducers and receive the echo signal. This forms the first stage of a redesign of the entire system with the intention of making it smaller and more power efficient. The long term goal is to use the sonar as a payload on an autonomous underwater vehicle (AUV); this will allow large and/or remote areas of the seafloor to be mapped without the need for a human presence. It will also enable the imaging of both more dangerous waters where the current towed system cannot be used. Obviously, the redesign will also benefit the towed system in the short term.

2 TRANSDUCERS

The KiwiSAS-IV towfish transmits signals on 36 transducers arranged in a 3×12 projector array and uses a separate 3×3 hydrophone array to receive the echo signals. In order to reduce the size of the system, the new design proposes using the current projector array for both transmitting and receiving signals.

The projector transducers are Tonzpizl oscillators. They consist of a stack of eight acoustically active ceramic discs sandwiched between an aluminium head piece and a mild steel tail piece. The discs are made from lead zirconate titanate (PZT) and are 2 mm thick with an outer diameter of 10 mm. Thin copper electrodes are placed between

each disc and are alternately soldered to two copper bars to form the electrical port; the discs are therefore connected in parallel electrically. A high-tensile steel bolt is placed through the transducer to keep the ceramic discs under compression as they are weak under tension.

The transducers were designed to resonate between 20 kHz and 40 kHz; this was confirmed by impedance and efficiency measurements [2] and finite element method (FEM) modelling [3]. They were subsequently found to have further resonant bands, most notably between 90 kHz and 110 kHz which is one of the frequency bands that KiwiSAS-IV uses. There are additional high-frequency resonances that must not be excited to avoid destroying the transducer.

To determine the location of the transducer resonances, the impedance of the transducers was measured. These measurements were taken by placing a transducer in series with a resistor to create a voltage divider network and applying a sinusoidal signal at the desired frequency. By measuring the magnitude and phase shift of the signal across the resistor, the complex current in the network, and hence the impedance of the transducer, can be determined.

Figure 1a shows the impedance of the transducer between DC and 5 MHz. This reveals a large resonance at approximately 3 MHz, the fundamental resonance of the ceramic discs. The magnitude of this resonance swamps the ‘active’ resonances, i.e., the resonances used in transmitting and receiving the signal. Figure 1b shows the impedance between DC and 200 kHz; this reveals a number of resonances including those used for operation.

3 RECEIVER

The KiwiSAS-IV system uses an INA111 low noise current instrumentation amplifier to boost the echo signal by 40 dB. The signal is then sampled at 312.5 kHz with a 16-bit resolution by an AD7723 delta-sigma (Δ - Σ) ADC. To

reduce the space and power requirements of the receiver electronics, this arrangement was replaced by the Analog Devices AD9271 IC. This chip provides eight channels, each with a low noise amplifier (LNA) and a variable gain amplifier (VGA) offering a total gain that can be controlled between 8 dB and 42 dB. Each channel is then simultaneously sampled by a 12-bit pipelined ADC at a controllable rate between 5 MHz and 50 MHz. The resulting data is given on double data rate (DDR) low-voltage differential signalling (LVDS) buses, one per channel; the clock speed of these buses varies between 30 MHz and 300 MHz depending on the sampling rate.

3.1 Decimation

Since the AD9271 is designed primarily for medical ultrasound, the sampling rate is much higher than necessary to capture the echo signal. Since the highest frequency used in the transmitted signal is 110 kHz, the Nyquist limit is 220 kHz. The sampling rate can be reduced by decimation; this is a two-stage process consisting of a low-pass anti-aliasing filter followed by a downsampler. The reduction of the Nyquist rate by the downsampler causes high-frequency energy to fold back into the passband of the slowed signal, hence the need to filter it prior to the downsampler.

The resolution of the output data can also be increased by decimation. Each decimation by a factor of two adds an extra bit of resolution; in general, for a decimation factor R , the number of extra bits in the output is given by

$$B_G = \lceil \log_2 R \rceil. \quad (1)$$

In theory, the resolution can always be increased by slowing the sampling rate further. However, at some point the extra bits will just contain noise rather than any useful information; this limit is typically imposed by the noise level at the input to the ADC and the linearity of the ADC. A




Figure 1: The admittance of the transducers.




Figure 2: A CIC decimator.

common definition for the effective number of bits in a given system is

$$\text{ENOB} = \frac{\text{SINAD} - 1.76 \text{ dB}}{6.02}, \quad (2)$$

where SINAD is the ratio of the desired signal to the sum of the noise level and the ADC distortion.

3.2 CIC filters

Finite impulse response (FIR) filters are typically preferred as anti-aliasing filters over infinite impulse response (IIR) filters due to their linear phase response [4]. However, directly implementing a FIR filter is computationally expensive, especially at high sampling rates. There are many methods of reducing the workload, including breaking the filtering into multiple stages [5, 6], using the polyphase decomposition to avoid calculating unneeded outputs [7, 8], and using half- or M-band structures [9].

The most computationally efficient filter for decimation is the cascaded integrator-comb (CIC) filter. Introduced by Hogenauer [10], it consists of N integrators followed by a downsampler and then N comb filters with a differential delay M as illustrated in Figure 2. By applying the third Noble identity, its transfer function can be found to be

$$H(z) = \frac{1 - z^{-RM}}{1 - z^{-1}}. \quad (3)$$

Rewriting this as the sum of a geometric series gives

$$H(z) = \left[\sum_{k=0}^{RM-1} z^{-k} \right]^N, \quad (4)$$

from which it can be seen that the CIC filter is equivalent to N cascaded FIR filters, each with RM unitary coefficients.

The CIC is an averaging filter; by considering the worst-case input Hogenauer shows that the register growth of a CIC filter is given by

$$G_{\max} = (RM)^N. \quad (5)$$

This corresponds to

$$B_G = \lceil N \log_2 (RM) \rceil \quad (6)$$

extra bits of resolution in the output data.

Substituting $z = e^{j2\pi f/R}$ into the transfer function gives the response for frequency f , where f is relative to the low sampling rate f_s/R . This shows that the CIC filter has a magnitude response

$$|H(f)| = \left| \frac{\sin \pi M f}{\sin \pi f/R} \right|^N, \quad (7)$$

and a linear phase response

$$\phi(f) = \pi f N \left(\frac{1}{R} - M \right). \quad (8)$$

By inspection, Equation 7 is undefined when $f = iR$ where i is any integer, i.e., it is undefined at DC and multiples of the sampling frequency f_s . To obtain the magnitude response at these points, we must take the limit of Equation 7 as $f \rightarrow iR$ and apply l'Hôpital's rule to give

$$\lim_{f \rightarrow iR} |H(f)| = (RM)^N. \quad (9)$$

Comparing Equation 5 and Equation 9, it can be seen that the gain in magnitude is really just the input signal being expanded to fit the increased bit width of the output data.

4 TRANSMITTER

The signals applied to the KiwiSAS-IV transducers have an amplitude of approximately $200 V_{p-p}$. The desired waveform is digitally stored in memory and output through a DAC. A TDA2030 class AB power amplifier is used to amplify the signal to around $20 V_{p-p}$; the signal is then passed through a 12:120 turn ratio transformer to step the voltage up to its output level.

Class AB amplifiers have a maximum theoretical efficiency of 78.5%; the amplifiers used for KiwiSAS-IV are approximately 50% efficient. Since most of the wasted power is dissipated as heat, the amplifiers require a heatsink to prevent them burning out; along with the various biasing components, this takes up a fair amount of space. The transformers also reduce the efficiency of the overall system and increase the space required.

4.1 Class D amplifiers

The class D amplifier is a switching amplifier, i.e., it uses transistors, typically MOSFETs, to switch the output 'on' and 'off' rather than to directly replicate the signal. Although a MOSFET has an on resistance which leads to some power dissipation, this is much less than the power required to bias the transistors in the class AB amplifier. The other main source of power loss in a switching amplifier is the power required to switch the transistors from off to on and vice-versa. The capacitance of the gate pin in a MOSFET requires a certain amount of charge each time it is switched; hence the faster it is switched, the greater the




Figure 3: The class D waveform for a sinusoidal signal.

overall power dissipation. In theory, the class D amplifier is 100% efficient; in reality, the on resistance and switching losses lower this, although 90% efficiency or above is achievable which is a large improvement over a class AB or similar amplifier.

This switching provides a series of rectangular pulses; the desired signal must be modulated onto these pulses in some way. The average value of a unipolar square wave $s(t)$ over a period T is

$$\bar{s} = DV_S, \quad (10)$$

where D is the duty cycle and V_S the supply voltage. If the switching frequency is greater than the highest signal frequency, the signal can be used to modulate the duty cycle of each period. Passing this waveform through a low-pass filter extracts the desired signal removing the switching frequency and harmonics. An example of a class D waveform for a sinusoidal signal is shown in Figure 3a; in this case the switching frequency is ten times the sinusoid frequency. The spectrum in Figure 3b clearly shows the energy present at both the signal and switching frequencies along with the various harmonics and intermodulation products. Increasing the switching frequency shifts the harmonics higher which allows the low-pass filter to be implemented as a lower-order filter; however it also leads to increased power losses in the MOSFETs. This is the main trade-off for a class D amplifier.

The highest signal frequency that the sonar system transmits is 110 kHz; the switching frequency was chosen as 800 kHz. This ensures that harmonics do not occur in the 3 MHz resonant band shown in Figure 1a while still giving a suitably large transition band for the filter.

4.2 Component selection

Identical N-channel MOSFETs were selected for both the high- and low-side transistors in the amplifier. Although

this requires a floating supply to drive the high-side transistor, N-channel MOSFETs have a lower on resistance than equivalently sized P-channel MOSFETs; additionally, this allows the rise and fall times of the class D output to be matched.

The transmitter circuit is shown in Figure 4. The MOSFETs are International Rectifier IRF7453 devices, designed for high-frequency DC-DC converters and capable of operating at voltages up to 250 V and currents of up to 2.2 A. Their typical on resistance is 200 m Ω and they come in a 6 mm by 5 mm SO-8 surface mount package. To provide the necessary gate drive signals, the International Rectifier IRS20124 chip is used. This is capable of switching the MOSFETs at 1 MHz, and it contains an in-built boost converter to generate the floating supply needed to drive the high-side MOSFET. A DC blocking capacitor is placed at the output of the amplifier to convert the signal from an unipolar 0 to 200 V to a bipolar -100 to 100 V. A two-pole LC low-pass filter with a -3 dB frequency of 159 kHz is used to attenuate the switching frequency and harmonics from the signal before it is output to the transducer. After transmission is complete, both MOSFETs are turned off to present a high impedance load on the transducer and allow the echo signal to be detected.

5 T/R SWITCH

Since the transmitter generates a ± 100 V signal and the inputs to the AD9271 are only rated to ± 2 V, a transmit/receive (T/R) switch is required to protect the AD9271 when the system is transmitting. The circuit shown in Figure 4 [11] provides this protection. The diode bridge clips large incoming signals to within one diode drop of the ± 5 V rails; the back-to-back Schottky diodes provide a second level of protection, further clipping the signal to approximately ± 400 mV. Provided that the diodes used in the bridge are balanced (i.e., they all have the same for-




Figure 4: The front-end for a single channel. The transmitter is shown to the left of the transducer while the T/R switch on the right leads into the inputs of the AD9271.



Figure 5: The PCB designed to evaluate the front-end. The AD9271 is the IC next to the connector.

ward voltage drop), small signals are passed through the switch without alteration. The initial portion of a fast switching signal can get through the switch before the diodes can switch off; however, the low-pass filter on the output of the amplifier prevents the switching from reaching the T/R switch.

6 PCB

A four layer PCB was designed and manufactured to evaluate the performance of this circuit. A single AD9271 chip was used with eight identical transmitter and T/R switches, one for each of the channels of the AD9271. A 100 pin connector was placed on the PCB to allow it to be connected to a Spartan 3E FPGA development board for testing. This connector provides access to a number of the LVDS inputs available on the FPGA which are used to receive the data from the AD9271. The remaining pins are used to control the transmitter circuits. A photo of the PCB




Figure 6: The output of the class D amplifier for a 110 kHz sinusoid with a supply voltage of 30 V.

is shown in Figure 5.

7 RESULTS

In the KiwiSAS-IV system, each channel applies an independent gain and delay to the signal, allowing the transmitted acoustic beam to be steered [12]. In the new design, the duty cycle for each output period is stored in memory, and then output to the channels via a tapped delay line. With the sampling rate of the amplifier being 800 kHz, each tap provides a delay of 1.25 μ s. The hardware multipliers in the FPGA are then used to scale the duty cycle to increase or decrease the amplitude of each channel. The duty cycle is converted to a square wave using a counter and a comparator to decide if the output should be high or low. Higher resolution delays are implemented at this point by delaying the start of the period. The resolution of this delay depends on the frequency the FPGA is running at; with a clock speed of 50 MHz the resolution is 20 ns.

An example waveform taken from one of the transmitter channels is shown in Figure 6. The signal being transmitted is a 110 kHz sinusoid; this is the highest frequency




Figure 7: The output of the T/R switch when the transmitter is outputting a 110 kHz, 30 V_{p-p} sinusoid.

signal that is used in the KiwiSAS system. The distortion present shows that some of the harmonics get through the low-pass filter without significant attenuation. With a simple second-order low-pass filter this was expected, and it is anticipated that the transducer will further attenuate these harmonics since they do not coincide with any resonances.

Figure 7 shows the T/R switch preventing the transmitted signal from reaching the receiver. Before transmission starts, the switch is open and the receiver inputs are sitting at ground. When the transmission begins, the voltage starts rising until it gets to about 300 mV; at this point the Schottky diodes turn on and clip the signal. When the voltage drops after the first half-cycle the switch briefly opens until the voltage reaches about -300 mV at which point it starts to clip again.

8 CONCLUSIONS

This paper has described the design of a multi-channel front-end to drive the KiwiSAS transducers with a 200 V signal and detect echoes. A prototype has been constructed and connected to an FPGA development board for testing. The transmitter has been successfully operated and the T/R switch verified to successfully prevent the high voltages from damaging the receiver IC.

8.1 Future work

Further work is required to program the FPGA to receive the data from the AD9271 IC and verify the operation of the CIC decimation filters. A separate project is currently underway within the department to design a back-end system to take the data received from this front-end and store it for later use. Once the receiver section is working and the back-end design, the circuits will need to be scaled to handle all 72 channels. Due to the size restrictions, this will require multiple PCBs with appropriate communication and synchronisation between them.

REFERENCES

- [1] D. R. Jackson, K. L. Williams, E. I. Thorsos, and S. G. Kargl, "High-frequency subcritical acoustic penetration into a sandy sediment," *IEEE Journal of Oceanic Engineering*, vol. 27, pp. 346–361, Jul 2002.
- [2] J. S. Knight, "Ultrasonic transducer development for a CTFM synthetic aperture sonar," Master's thesis, University of Canterbury, Feb 1987.
- [3] D. W. Hawkins, *Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar*. PhD thesis, University of Canterbury, 1996.
- [4] R. G. Lyons, *Understanding digital signal processing*. Prentice Hall, 2 ed., 2004.
- [5] R. E. Crochiere and L. R. Rabiner, "Optimum FIR digital filter implementations for decimation, interpolation, and narrow-band filtering," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 23, pp. 444–456, Oct 1975.
- [6] M. W. Coffey, "Optimizing multistage decimation and interpolation processing," *IEEE Signal Processing Letters*, vol. 10, pp. 107–110, Apr 2003.
- [7] M. G. Bellanger, G. Bonnerot, and M. Coudreuse, "Digital filtering by polyphase network: Application to sample-rate alteration and filter banks," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, pp. 109–114, Apr 1976.
- [8] R. E. Crochiere and L. R. Rabiner, "Interpolation and decimation of digital signals: a tutorial review," *Proceedings of the IEEE*, vol. 69, pp. 300–331, Mar 1981.
- [9] M. G. Bellanger, J. L. Daguët, and G. P. Lepagnol, "Interpolation, extrapolation, and reduction of computation speed in digital filters," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 22, pp. 231–235, Aug 1974.
- [10] E. B. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, pp. 155–162, Apr 1981.
- [11] Analog Devices, "AD9271 octal LNA/VGA/AAF/ADC and crosspoint switch data sheet (Rev. B)," May 2009.
- [12] M. P. Hayes, "A real-time beam steering system for KiwiSAS-IV," in *Proceedings of Electronics New Zealand Conference, ENZCON2003*, pp. 27–32, Sep 2003.

Digital Receiver for Transient Radio Emission Array Detector Prototype (TREAD-P).

Vyacheslav Kitaev¹, Achala Perera², Gene Souldenkov³

School of Engineering

Auckland University of Technology

Private Bag 92006, Auckland 1142, NEW ZEALAND.

¹Email: slava.kitaev@ieee.org

²Email: achala.perera@aut.ac.nz

³Email: gene.souldenkov@aut.ac.nz

Abstract: The paper provides an insight into development of FPGA-based autonomous digital low frequency receiver for the Transient Radio Emission Array Detector Prototype (TREAD-P). The overall project has been outlined and the details of designed hardware have been provided and discussed.

Keywords: Transient sky, SKA, radio astronomy, array, FPGA, digital receiver

1. INTRODUCTION

Transient radio sky or time-variable radio sources in space have been recognized as one of the key science drivers for the Square Kilometre Array (SKA) what will be reflected in its design and operations (Cordes, 07/2007 (Revised 04/2009)). This science area is marked as “Exploration of the Unknown” what includes the likely discovery of new classes of objects and phenomena (Wilkinson, 2004). Transient Radio Emission (TRE) is characterised by irregular and very short (nanoseconds to seconds) bursts of electromagnetic radiation. Examples of TRE that are known range from 0.4 ns, such as giant pulses from the Crab pulsar, and longer in time scale with apparent brightness temperatures from thermal to 10^{42} K. However, yet we know very little about the overall constituency of the transient radio sky, what even includes the high-energy transient sky.

Currently used methods and techniques for detecting TRE differ depending on the expected nature of the events, their origin, and the types of instruments used for detection. The detection of TRE events, especially the short duration ones (<1ms) in radio spectrum can be severely complicated by the ever growing problem of terrestrial radio interferences (RFI) which may even lead to false detections.

Single beam instruments like e.g. wide-field gamma-ray telescopes can have highly variable event-by-event resolution. Therefore, processing of data, coming from these sources can be done even with a simple threshold test approach: the spectrum is first flattened by integrating it over time and a

threshold limit is applied. Such approach is less effective at radio frequencies of electromagnetic spectrum mostly again due to RFI, especially at the frequencies below 1GHz. Recent advances in data processing algorithms, combined with rapidly increasing computational powers may allow for significant improvements in this area for RFI mitigation.

Another effective way to attack the problem of possible false detections due to RFI is using interferometers (Bannister, 2009). Multi-antenna/station measurements have a possibility for cross-correlating signals which removes local unresolved RFI. Moreover, radio interferometers like LOFAR, MWA and SKA will produce multiple beams and since the observations are simultaneous and synchronized in time, it is possible to compare signals found in different beams to help discriminate TRE signals from RFI.

One of the main tools of RFI mitigation is the time-frequency analysis of received signals with a high temporal (less than 1 microsecond) and frequency (less than 1 kHz) resolution. This approach allows for analysing statistics of the mixture “system noise” + “source noise” + RFI and separate RFI from the presuming Gaussian distribution of “system noise” plus “source noise” combination. In case of terrestrial RFI, which has the highest power threshold due to the proximity to the detectors, it is assumed that the unwanted signal will a) will appear strongly in one beam pointed in the direction of origin; b) will correlate only weakly with the similar noise from another beam. Since multibeam

techniques require phased cross-correlation, the location of the source can also be a factor in identifying the origin of the signal (terrestrial or stellar). For VLBI-like interferometers terrestrial signals are set wide apart and as said can be easily removed.

One of the characteristics of TRE signals, which are short-living and probably aperiodic events, is limited or incomplete information about their statistical properties. Therefore, the detection of such signals fits well into the stochastic analyses domain where the assumption is made by many (Bulsara, 1996; Thompson, 2001; Cordes, 2003) that the noise in noise+signal mix possesses Gaussian qualities and as such can be discriminated against signal. From a purely mathematical perspective the problem of optimal signal detection was solved once it was connected to statistical hypothesis testing (Neyman, 1933).

In case of transient extraterrestrial signals, the objective is to first derive an expression for the likelihood ratio when the signal to be detected can consist of multiple transient signals with unknown arrival times; then compare the performance of a detector based on this likelihood ratio to one with the maximum likelihood ratio detection statistic. Several practically relevant detection problems have been solved by using this methodology, such as maximum likelihood estimators, where Cramer-Rao lower bound is used for thresholding (Leshem, 2000).

Although some statistical methods can greatly increase the overall sensitivity of detection, the greatest disadvantage is that they work well only in RFI “sterile” environment when the interferences are minimized. In an environment contaminated with highly energised RFI these procedures are not robust and statistically unstable in a presence of outliers in the time or frequency domains. However combining these methods with VLBI-like interferometer has the potential to be very effective and be sensitive even below one sigma.

Additional distinct feature of TRE can also be used combined with the statistical methods which is that unlike most of RFI, TRE signal arrives in different time at different frequencies due to the delay of the pulse occurring in the interstellar plasma.

It is important to note that all the above techniques and approaches to TRE detection require significant computational resources as work with large datasets from multiple antennas or sensors and are computationally intensive in general. Defining suitable and effective approaches and algorithms is necessary for the future radio telescopes such as SKA, LO-FAR and MWA. This was the main motivation for setting up an experimental instrument Transient

Radio Emission Array Detector Prototype in New Zealand (TREAD-P).

2. TREAD-P

Figure 1 shows the time-luminosity phase space for radio transients (Cordes, 07/2007 (Revised 04/2009)) as a log-log plot of the product of peak flux S_{pk} in Jy and the square of the distance D in kpc vs. the product of frequency ν in GHz and pulse width W in s. As one can see there are still many areas there a few or none objects have been found yet, including the areas where the large sensitivity is not required, but rather the right combination of field of view, sensitivity, time resolution and ability to differentiate TRE from RFI.

TREAD-P is thought to be positioned in the area indicated by yellow fading triangle on Fig 1. It will consist of a few stations deployed in different locations of New Zealand sufficiently distant from each other to have most of RFI at each station uncorrelated with other stations. Each station will consist of a number of cross-polarised low frequency antennas forming a phased array with effective area approximately 100m² in each polarization (see Figure 2).




Figure 1. Time-luminosity phase space for radio transients (adopted from (Cordes, 07/2007 (Revised 04/2009))). It shows a log-log plot of the product of peak flux S_{pk} in Jy and the square of the distance D in kpc vs. the product of frequency ν in GHz and pulse width W in s. The yellow fading triangle shows the approximate area of detection of unknown objects by TREAD-P.

Each station has a single wide beam providing a field of view 15-35° depending on the frequency. A dual-channel direct sampling wide band digital receiver radiometer DRS at each station is equipped with communication interface for real-time data steaming. TREAD-P will exploit already existing

infrastructure of Kiwi Advanced Research and Education Network¹ and BeSTGRID² to provide real-time 24/7 data streaming and processing. Figure 3 shows planned deployment and operation of TREAD-P.

3. DRS

The DRS is a fully contained system which includes dual-channel synchronized direct sampling receiver and an embedded system with various input/output interfaces and memories. The following is a detailed discussion of the concept and design of DRS.




Figure 2. Computer model of one of the possible solutions for TREAD-P station. The approximate dimensions of the station on the ground are 10.5x10.5 m.




Figure 3. Deployment and operation of TREAD-P in New Zealand.

The DRS has a modular design which generally features the embedded system based around Altera Cyclone II FPGA system, Analogue to Digital Converter (ADC) & Global Positioning System (GPS) & precision clock module, Filter & Pre-amplifier, Amplifier Module, and either So-

lar/Battery Power Supply or main Power Supply (~230V). Figure 4 shows the system modules block-diagram of DRS with mains power supply module.

The main reason for such modularity is to provide an ability to replace either of modules if necessary without a need to redesign the rest of the modules. This is because TREAD-P is a prototype system which will require further considerations and optimizations of e.g. the frequency band, simultaneous sampling band and possibly pre-processing capabilities. Figure 4 shows the functional block-diagram of DRS which is further discussed in details. Figure 5 shows the enclosure of DRS in two EMI shielded boxes. The box on the left contains digital elements of the system with 100MHz fibre-optical data transmission line to reduce harmful interferences to the antenna and analogue electronics.




Figure 4 DRS system module block-diagram. The blocks represent physically separate modules of DRS system.

The first stage amplifier and band pass filtering modules are based around the low noise amplifier ZFL-500LN from Mini-Circuits (G1A and G1B on the diagram) which provides 28dB gain and intended to be located as close to the antenna as possible. The filters BPF 1A, 2A, 1B and 2B are band pass filters to limit the frequencies to the desired. Filters are split as primary filter, just before the amplifier and secondary filter and DC block after the amplifier. Both filters are 3rd order passive L and C filters with sharp cut-off frequency response to minimize aliasing during digitalisation.

The power for the low noise amplifiers G1A and G1B is supplied via the coaxial signal cable from the dual channel amplifier module. The second stage filter units BPF 2A and BPF 2B block the DC voltage going into the amplifier output, but supply the power to a separate connectors of the ZFL-500LN amplifier modules.

The second stage amplifier module G2 is based on AD8330 0-150MHz variable gain amplifier with a gain of 50dB (Figure 6). The amplifier has fully

¹ <http://www.reannz.co.nz/>

² <http://www.bestgrid.org/>




Figure 5. DRS functional block-diagram




Figure 6. Assembled DRS. The box on the left contains digital elements of the system with 100MHz fibre-optical data transmission line to reduce harmful interferences to the antenna and analogue electronics.

differential signal path. The gain of AD8330 is controlled by external voltage source from 0 to 1.5V DC using a stable single transistor variable voltage source. The differential output signal is converted back to single ended by using a high frequency transformer. Also DC voltage is injected to the input line for the module to supply the power to the antenna pre-amplifier G1A and G1B modules.

GPS, ADC and Precision Clock module consists of dual-channel signal conditioning variable gain amplifier AD8330, dual-channel 130MSPS MAX19517 ADC device, high accuracy oven control 10MHz clock oscillator and Trimble Resolution-T GPS module (Figure 7).

The single ended input is converted to differential input in order to feed the single ended signal into differential signal conditioning Pre-ADC amplifier AD8330. The amplifier can be controlled either manually via a hardware switch or from the FPGA using a Digital to Analogue Converter chip (DAC).

The conditioned signal then is fed via balun to match the input impedance to the ADC MAX19515 dual-channel 10-bit asynchronous ADC converter. This ADC chip is capable of baseband and IF sampling up to 400 MHz. Redundancy in the sampling rate is necessary to allow a wider bandwidth at later stages of development.

The regime of ADC output and operation is controlled through SPI interface from the FPGA based embedded system. The ADC clock is provided from the FPGA Phase Lock Loop (PLL). Each channel has 10-bit data and synchronisation clock outputs to the FPGA. The digital data output track's lengths are tuned to optimise the propagation delay of the signal. The digital data then is transferred to the data processing board through high grade ribbon cable.



Figure 6: Two Channel Amplifier G2 module



Figure 7: ADC, GPS and Precision Clock module

The ADC and Clock module contains a high accuracy ovenized controlled 10 MHz oscillator OC14T5A which provides stability $<5 \cdot 10^{-11}$ at 1 second over the wider range of temperatures -25C to 75C.

Another essential module to achieve the highest accuracy and stability clock is Trimble Resolution-T GPS module. This module is connected to GPS antenna and provides Pulse Per Second (PPS) signal with ± 15 ns accuracy. NIMA data stream to FPGA based embedded system is used to obtain the location and time which are then embedded with the output data of DRS.

10 MHz oscillator is fed into a Phase Locked Loop (PLL) block on Cyclone II FPGA. The PPS signal is used to lock the 10 MHz on the PLL and produce high stability and high accuracy clock output which is then fed back to the ADC and used to clock the FPGA design. This method is expected to provide the sufficient accuracy to synchronize ADCs of multiple DRS devices deployed in geographically different locations.

The output of 10-bit ADC data is supplied to the high speed LVDS pins of the FPGA. Output of Trimble Resolution-T GPS module is fed into FPGA based embedded system via a serial UART interface.

Initially the embedded FPGA system was designed using a large Cyclone III device. However later it was found that readily COTS available Altera DE2 development kit based on Cyclone II EP2C35 FPGA is sufficient, rich in I/O interfaces and very inexpensive to academic institutions.

The embedded application written in C is running on NIOS II processor utilizing VGA, Ethernet 10/100, and PS/2 (keyboard) ports as well as the available LCD display and SD-card data storage slot. At this stage the embedded software application is capable of streaming raw data without pre-processing.

4 FUTURE WORK

Based on continuing considerations for the optimal frequency band and the bandwidth the analogue modules of the DRS are likely to be changed. This work will be accompanied by the design of a suitable antenna or array of antennas perhaps similar to those used in MWA (Bowman J. D. et al., 2007).

5 CONCLUSION

We have developed the direct sampling dual channel digital receiver radiometer DRS for TREAD-P. DRS will provide the initial pre-processing of the data acquired from a phased array of antennas and streaming the data into KAREN network for further processing in order to detect TRE signals from several simultaneously and synchronously working TREAD-P stations. Synchronization is provided by a high accuracy GPS and on-board high accuracy oscillator.

6 ACKNOWLEDGMENTS

The work has been partially supported by the grant KAREN CBF.

Authors thank Alan Rogers of Haystack Observatory of the MIT for the consultations about MWA antenna design and allowing using the software generating NEC files for the antenna.

7 REFERENCES

Bannister, K. (2009) "Detecting fast radio transients with next generation interferometers" [High Performance Computing in Observational](#)

- Astronomy: Requirements and Challenges
Volume, DOI:
<http://www.ncra.tifr.res.in/~hpc.astro/>
- Bowman J. D., D. G. B., Frank H. Briggs, Brian E. Corey, Merv J. Lynch, N. D. Ramesh Bhat., S. S. D. Roger J. Cappallo, Brian J. Fanous, David Herne, Jacqueline N. Hewitt., et al. (2007). "Field Deployment of Prototype Antenna Tiles for the Mileura Widefield Array Low Frequency Demonstrator." The Astronomical Journal **133**: 1505Y1518.
- Bulsara, A. R., & Gammaitoni, L. (1996). "Tuning in to Noise." Physics Today(49): 39-45.
- Cassiday, G. L., & et al. (1990). Mapping the u.h.e. Sky in search of point sources. Nuclear Physics B.
- Cordes, J. M. (07/2007 (Revised 04/2009)). "The Square Kilometre Array as a Radio Synoptic Survey Telescope: Widefield Surveys or Transients, Pulsars and ETI." SKA Memo(97).
- Cordes, J. M., & McLaughlin, M.A. (2003). "Searches for Fast Radio Transients." The Astrophysical Journal **596**: 1142-1154.
- Hastings, W. K. (1970). "Monte-Carlo sampling methods using Markov chains and their applications." Biometrika(57): 97-109.
- Hobson, M. P., & McLachlan, C. (2002). "A Bayesian approach to discrete object detection in astronomical datasets." Monthly Notices of the Royal Astronomical Society **338**(3): 765-784.
- Huchra, J. P., & Geller, M. J., & de Lapparent, V., & Corwin, H. G., Jr. (1990). "The CfA Redshift Survey: Data for the NGP +30 Zone." ApJS(72): 433.
- Leshem, A., & van der Veen, A.-J. (2000). "Radio astronomical imaging in the presence of strong radio interference." IEEE Trans. on Information Theory **46**(5): 1730-1747.
- Metropolis, N., & Rosenbluth, A. W., & Rosenbluth, M. N., & Teller, A. H., & Teller, E. (1953). "Equations of state calculations by fast computing machines." Journal of Chemical Physics **21**(6): 1087-1092.
- Morales, M. F., & Williams, D.A., & DeYoung, T. (2004). "An advanced analysis technique for transient searches in wide-field gamma-ray observatories." Astroparticle Physics **20**(4): 485-497.
- Neyman, J., & Pearson, E., (1933). "On the problem of the most efficient tests of statistical hypotheses." Phil. Trans. Roy. Soc. **A231**(9): 492-510.
- Palonpon, A., & Amistoso, J., & Holdsworth, J., & Garcia, W., & Saloma, C. (1998). "Measurement of weak transmittances by stochastic resonance." Optics Letters **23**(18): 1480.
- Thompson, A. R., & Moran, J.M., & Swenson Jr., G.W. (2001). Interferometry and Synthesis in Radio Astronomy. New York, Wiley-Interscience.
- Wilkinson, P. N., Kellermann, K. I., Ekers, R. D., Cordes, J. M., & Lazio, T. J. W. (2004). "Exploration of the Unknown." New Astronomy Review **48**: 1551.

Minimum Hamming Weight Representations for Irregular Symbol Alphabets

William Kamp*, Andrew Bainbridge-Smith*, Michael Hayes*

*Department of Electrical and Computer Engineering, University of Canterbury
Private Bag 4800, Christchurch 8140, NEW ZEALAND.
Email: w.kamp@ieec.org

Abstract:

This paper details an algorithm to recode a number into a redundant number system such that the number of non-zero symbols (digits) is minimised, that is to minimise its Hamming weight. The target number system may have an alphabet with irregularly spaced symbols and any radix. The state machine constructed to perform the conversion is analysed to determine the average minimum Hamming weight of its output representations. The averages obtained for some common number systems are confirmed by other published results.

Keywords:

Redundant number systems, Hamming weight, Shift and Add, Algorithm

1 INTRODUCTION

The demand for performance from digital signal processing systems continues to grow. For example, the SKA (square kilometer array) radio telescope will require an extraordinary amount of signal processing, mostly to be performed digitally and in real time [1]. The entire signal processing system will be re-examined including not only the algorithms used for correlation, but the technology they are built with.

Redundant number systems (RNS) provide opportunities for increasing the performance of digital signal processing hardware. Redundant representations have two fundamental advantages over binary. Firstly, the Hamming weight of a representation can be decreased i.e., a fewer number of nonzero symbols. Employing representations with low Hamming weight can reduce the number of calculations that must be performed in an algorithm. For example, the number of partial products in a multiplication can be reduced [2], meaning fewer additions and hence a shorter delay. Secondly, redundant representations support carry propagation-free addition [3], meaning an addition is performed in constant time, irrespective of the operands' widths [4].

A RNS is a positional number system that has a symbol alphabet, \mathbb{S} , with cardinality greater than the radix, β . Each symbol s in the alphabet has a rational value associated with it, usually an integer. As usual in a positional

system a representation's value, v , is determined as,

$$v = \sum_{i=0}^{N-1} s_{\langle i \rangle} \beta^i, \quad (1)$$

where $s_{\langle i \rangle}$ is the i^{th} symbol (digit) and N is the representation word-length, which may be infinite, although usually restricted in practical implementations.

A simple RNS is binary signed digit (BSD) having the symbol alphabet $\mathbb{S} = \{\bar{1}, 0, 1\}$ with corresponding values $\{-1, 0, 1\}$ and $\beta = 2$ [5]. RNS have multiple representations for most values in their range. For example, the value 15 in BSD can be written as 01111_2 , $1000\bar{1}_2$, $100\bar{1}1_2$, and $1\bar{1}111_2$ among many other representations.

The RNS that we are primarily interested in are those with $\beta = 2$ and alphabets with a cardinality of four or less. A radix of two is chosen since its representations have the highest starting proportion of zero-symbols of any other radix. For example, on average radix-4 representations have a quarter zero symbols, radix-3 – a third, and radix-2 – a half. With a RNS, these proportions can be increased by choosing an appropriate symbol alphabet. One of the smallest radix-2 symbol alphabets that gives a redundant number system is $\{\bar{1}, 0, 1\}$ [6]. With three symbols it requires the combined states of two digital signals to encode the BSD alphabet $\{\bar{1}, 0, 1\}$. This is repeated at every position along the representation. To implement a RNS efficiently in digital logic the number of signals should be minimised, i.e., limited to two per symbol. This

has a flow on effect of reducing the combinational logic required to perform operations like addition. The digital signal pairs can represent four states, hence four symbols. Some work has been done on ternary CMOS circuits, but they are much larger and slower [7].

This paper presents a method to construct a Mealy state machine that serially encodes a non-redundant number representation into a minimum Hamming weight (MHW) representation using a prescribed RNS. The method developed can handle irregularly spaced symbols that other methods do not take into account. The generalised sliding window solution [8] covers a large range of symbol alphabets, however, it limits these to alphabets with consecutive symbols ignoring symbols that are a multiple of the radix. For example, the alphabet $\{\bar{1}, 0, 1, 3\}, \beta = 2$ is valid, where symbol-2 is a multiple of β . Alphabet $\{\bar{1}, 0, 1, 5\}, \beta = 2$ cannot be encoded since the algorithm wrongly assumes there is a 3-symbol. The new algorithm supports MHW encoding these irregular number systems.

The new algorithm consists of three parts:

1. Find all MHW recodings for a relatively small range of values and place in a recoding table.
2. Reduce the recoding table to have the smallest recoding window size to create a mapping table.
3. Build a Mealy state machine that serially implements the recoding table and extends it past the window width.

The linking mechanism between the parts are tables associating a windowed section of the non-redundant input word to its MHW replacement, more specifically the least significant symbol (LSS) of the replacement. Part 1 uses a branch and bound tree algorithm to find the MHW representations and is presented in section 2. Part 2 groups the possible MHW representations, reducing the recoding table to a minimum number of entries and is presented in section 3. Part 3 is presented in section 4, where the mapping table is used to build a MHW recoding Mealy state machine. Finally, the average MHW is calculated from the state machine for use as a metric for the RNS in section 5. Some results are presented in section 6.

2 MHW RECODING TABLE

The recoding table lists all the MHW representations for values in the range $[0, \beta^M)$. These are the values of a windowed section of the input word, $x_{\langle i+M-1 \rangle}, \dots, x_{\langle i \rangle}$, of length M symbols. M is the largest acceptable window width and sets the maximum search space. This needs to be selected such that it is large enough to find a consistent

MHW encoding. The window width used in the final state machine is m and it is found in section 3. Therefore, M needs to satisfy $m \leq M \leq N$.

The recoding table entries consist of the windowed value and a list of MHW representations, which, may have any width. Only the least significant symbol $s_{\langle i \rangle}$ will be used. The rest of the MHW representation will be regenerated by the state machine later. The replacement strings are generated for a redundant number system defined by symbol alphabet \mathbb{S}_r and radix β . To efficiently find all the MHW representations for each $v \in [0, \beta^M)$ a branch and bound technique is employed. A tree such as Figure 3 is built by the algorithm. Each leaf node represents a windowed value. The transitions' outputs along each path back to the root node gives a MHW representation and the inputs give the windowed string. The tree is built breadth-wise, branching to a new level of nodes as the window width is increased. Several attributes are stored by the node data structure of Figure 1.

1. The input string's value along a path of length l from the root node, $\sum_{i=0}^{l-1} x_{\langle i \rangle} \beta^i$. This must be identical for all paths to this node.
2. The propagate value, the difference between the input string and the encoded string, projected to the next tree level, $\left[\sum_{i=0}^{l-1} (x_{\langle i \rangle} - s_{\langle i \rangle}) \beta^i \right] / \beta$. This must be identical for all paths to this node.
3. The Hamming weight of the encoded string along all paths from the root node. This must be identical for all paths and is used as the node merging or removal criteria.
4. A set of the least significant symbols for all paths leading to this node, i.e., the symbol(s) output on each path's transition from the root node.

The tree begins from a root node with the default (zeroed) attributes at level $l = 0$. From each node the input symbol $x_{\langle i+l \rangle}$ can take one of β different values. Each input symbol could result in several non-redundant symbols to be output. For each permutation of input and output symbols a transition to a new node is created if and only if,

$$(x + p) \bmod \beta \equiv s \bmod \beta, x \in [0, \beta), s \in \mathbb{S}_r, \quad (2)$$

where p is the propagate attribute of the node. The transition has an input condition of $x_{\langle i+l \rangle} = x$ and a Mealy output of $s_{\langle i+l \rangle} = s$.

Each transition terminates at a new node in the next level. The node's value and propagate attributes are derived from the transition's input and output symbols as previously described. If two nodes have equivalent value and

```

1: class node()
2:   self.value ← 0   ▷ Value of input symbol string.
3:   self.propagate ← 0 ▷ Difference between value
   and output symbol string value.
4:   self.weight ← 0   ▷ Hamming weight of output
   symbol string(s).
5:   self.lss ← {} ▷ Set of least significant symbols of
   paths leading to this node.
6: end class

```

Figure 1: Definition of a data structure for storing the node attributes for the FINDENCODINGTABLE algorithm.

propagate attributes then the node with the largest Hamming weight attribute is deleted along with its transitions. If the two nodes' Hamming weights are equivalent then the nodes are merged with the union of their least significant symbol sets and incoming transitions. By this mechanism, only the MHW encoding paths are kept in the tree and paths with greater Hamming weight are terminated as soon as a better one is found.

When $l = M$, $P = \lceil \log_{\beta}(\max(\mathbb{S})) \rceil + 1$ more levels are generated with $x = 0$ as the input symbol in (2). This allows the propagate values to converge to zero and as they do so the nodes will merge or eliminate each other until one leaf node for each value remains. The algorithm this far is presented in Figure 2 as pseudo code. An example of a completed search tree for BSD ($\mathbb{S} = \{-1, 0, 1\}$, $\beta = 2$) with $M = 4$ is shown in Figure 3.

Upon completion of the tree there will be β^M leaf nodes, one for each maximum-window value and all with propagate attributes equal to zero. The recoding table summarises the tree, listing the window values against a set of least significant symbols. A set may contain multiple symbols if more than one path back to the root node exists. For example, the BSD number system generates Table 1, where values 3, 11, and 13 have multiple paths and hence two symbols in their least significant symbol sets, $\{\bar{1}, 1\}$.

Notice in Table 1 that for each value $v \bmod \beta = 0$ (the even values when $\beta = 2$) the output symbol is zero. This allows an optimisation in the tree generation where the *root node* does not need to generate a transition for $x_{\langle i \rangle} = 0, s_{\langle i \rangle} = 0$.

3 MINIMUM WINDOW WIDTH

The next stage of the algorithm attempts to find the minimum window width that will give a MHW encoding. This will make the finite state machine smaller and if implemented in hardware, faster and more resource efficient.

```

1: procedure FINDENCODINGTABLE( $\mathbb{S}, \beta$ )
2:    $P \leftarrow \{\text{new node}\}$  ▷ Initialise parent node list
   with default root node
3:   for  $l = 0, l < M, l = l + 1$  do
4:      $P \leftarrow \text{MAKEMINWEIGHTENCODINGLEVEL}(\mathbb{X}, \mathbb{S}, \beta, P, l)$ 
5:   end for
6:   for  $l = M, l < M + P, l = l + 1$  do
7:      $P \leftarrow \text{MAKEMINWEIGHTENCODINGLEVEL}(\{0\}, \mathbb{S}, \beta, P, l)$ 
8:   end for
9:   return  $P$ 
10: end procedure
11:
12: procedure MAKEMINWEIGHTENCODINGLEVEL( $\mathbb{X}, \mathbb{S}, \beta, P, l$ )
13:    $C \leftarrow \{\}$  ▷ Initialise empty child list
14:   for each  $p \in P$  and  $x \in X$  and  $s \in S$  do
15:     if  $s - x - p.\text{propagate} \bmod \beta = 0$  then
16:        $c \leftarrow \text{new node}()$ 
17:       if  $p.\text{lss} = \{\}$  then
18:          $c.\text{lss} \leftarrow \{s\}$ 
19:       else
20:          $c.\text{lss} \leftarrow p.\text{lss}$ 
21:       end if
22:        $c.\text{value} \leftarrow p.\text{value} + x \times \beta^l$ 
23:        $c.\text{propagate} \leftarrow (x + p.\text{propagate} - s) / \beta$ 
24:        $c.\text{weight} \leftarrow p.\text{weight}$ 
25:       if  $s \neq 0$  then
26:          $c.\text{weight} \leftarrow c.\text{weight} + 1$ 
27:       end if
28:        $d \leftarrow \text{node} \in C$  similar to  $c$  else None
29:       ▷ similar means same value & propagate.
30:       if  $d \neq \text{None}$  then
31:         if  $d.\text{weight} > c.\text{weight}$  then
32:            $C \leftarrow C - \{d\}$  ▷ Remove  $d$  from  $C$ 
33:            $C \leftarrow C \cup \{c\}$  ▷ Add  $c$  to  $C$ 
34:         else if  $d.\text{weight} = c.\text{weight}$  then
35:            $c.\text{lss} \leftarrow c.\text{lss} \cup d.\text{lss}$  ▷ Merge  $c$  &  $d$ 
36:            $C \leftarrow C \cup \{c\}$  ▷ Add  $c$  to  $C$ 
37:         end if
38:       else
39:          $C \leftarrow C \cup \{c\}$  ▷ Add  $c$  to  $C$ 
40:       end if
41:     end if
42:   end for
43:   return  $C$ 
44: end procedure

```

Figure 2: Algorithm to find the least significant symbol(s) of the minimum weight representations of the numbers $[0, \beta^M)$ for the number system with alphabet \mathbb{S} , and radix β .

Table 1: Recoding table for binary to MHW BSD. Showing windows of maximum length $M = 4$ and the least significant symbol sets that will give a minimum weight encoding.

v	x_{m-1}, \dots, x_0	\longrightarrow	$s_{(0)}$ set
0	0_2	\longrightarrow	{ 0 }
1	1_2	\longrightarrow	{ 1 }
2	10_2	\longrightarrow	{ 0 }
3	11_2	\longrightarrow	{ $\bar{1}$ 1 }
4	100_2	\longrightarrow	{ 0 }
5	101_2	\longrightarrow	{ 1 }
6	110_2	\longrightarrow	{ 0 }
7	111_2	\longrightarrow	{ $\bar{1}$ }
8	1000_2	\longrightarrow	{ 0 }
9	1001_2	\longrightarrow	{ 1 }
10	1010_2	\longrightarrow	{ 0 }
11	1011_2	\longrightarrow	{ $\bar{1}$ 1 }
12	1100_2	\longrightarrow	{ 0 }
13	1101_2	\longrightarrow	{ $\bar{1}$ 1 }
14	1110_2	\longrightarrow	{ 0 }
15	1111_2	\longrightarrow	{ $\bar{1}$ }

Table 2: MHW mapping table for binary to MHW BSD with a window width of $m = 2$.

G_x	x_{i+1}, x_i	values	\longrightarrow	$s_{(i)}$
G_0	0, 0	{0, 4, 8, 12, ...}	\longrightarrow	0
G_1	0, 1	{1, 5, 9, 13, ...}	\longrightarrow	1
G_2	1, 0	{2, 6, 10, 14, ...}	\longrightarrow	0
G_3	1, 1	{3, 7, 11, 15, ...}	\longrightarrow	$\bar{1}$

The input condition of the Mealy transitions is a m length window of the non-redundant representation, $x_{(i+m-1)} \cdots x_{(i)}$. This consists of a base condition, $x_{(i+m-2)} \cdots x_{(i)}$ and a decision symbol, $x_{(i+m-1)}$. The base condition is common to all transitions in to and out of a state. The base condition is present as the most significant symbols of the input condition into a state $x_{(i+m-1)} \cdots x_{(i+1)}$. It is really a property of the state but is also attributed to the transitions for clarity. The output symbol $s_{(i)}$ for a transition is taken from the MHW mapping table T at the entry value,

$$t = (w_{(i)} + p_k) \bmod \beta^m; \quad (4)$$

$$s_{(i)} \leftarrow T(t). \quad (5)$$

Each state in the finite state machine has a unique pairing of attributes:

1. A value b_k in the range $(0, \beta^{m-1}]$. This is the value of the $m-1$ length base condition. From an incoming transitions condition,

$$b_{k+1} = x_{(i+m-1)} \cdots x_{(i+1)}. \quad (6)$$

2. A propagate value, being the difference of the input and output strings to that point. The next state's propagate, p_{k+1} , is calculated from the current state's propagate, p_k , and the connecting transition's input and output symbols, $x_{(i)}$ and $s_{(i)}$ respectively,

$$p_{k+1} = \lceil p_k / \beta \rceil + \lfloor (x_{(i)} - s_{(i)}) / \beta \rfloor. \quad (7)$$

The propagate value is negative if the output string has borrowed some value from future input windows and positive if the output string's value lags the input and needs to be made up in future transitions.

The attribute pair (b_{k+1}, p_{k+1}) uniquely identifies a state and can be determined from incoming transitions attributes and its outgoing state propagate attribute.

The state machine is built recursively starting from the reset state S_R , where the propagate value $p_{(0)} = 0$. The transitions out of the reset state are special, as they do not have a base condition. The full window of m input symbols is used to determine the next state. Therefore, β^m transitions are created from the reset state, one for each of the mapping table entries. A transition's end state is identified by calculating the propagate value p_{k+1} from (7) and the base condition b_{k+1} from (6). If the state (b_{k+1}, p_{k+1}) does not exist then it is created with β new transitions out to be resolved. The state machine is complete when all transitions' end states are resolved. The MHW encoding Mealy state machine for BSD with a window width of $m = 2$ is shown in Figure 4. A larger encoding state machine is shown in Figure 5 for the number system $\mathbb{S} = \{\bar{1}, 0, 1, 3\}, \beta = 2$.

5 AVERAGE MHW

Representations with MHW are desirable as they maximise the number of zero symbols in a word. A useful metric to compare number system is the expected number of nonzero symbols in a representation of length N that is MHW encoded. That is the MHW averaged over all values. Performing a Markov analysis [9] of the MHW encoding state machine reveals the probabilities of being in each state, hence, the probabilities of the symbols output at each transition from that state.

A Markov chain is a sequence of random variables (states) X_1, X_2, X_3, \dots with the Markov property; the next state depends only on the current state and no previous states [9]. The MHW encoding state machines are designed with this in mind. With the Markov property no history is kept meaning the transitions must have a stationary probability of being selected. In the case of the MHW encoding state machines, the transition probabilities are a constant $1/\beta$, with exception of those from the reset state with probability $1/\beta^m$.

Table 3: The average MHW for some symbol alphabets in radix-2. Maximum window width $M = 7$. Some entries confirmed by the listed references.

S	m	$H_{\min}(V)$	confirmed
$\{0, 1\}$	1	$0.5N + 0.0$	
$\{0, 1, 2\}$	1	$0.5N + 0.0$	
$\{0, 1, 2, 3\}$	2	$0.333333N + 0.111111$	
$\{0, 1, 2, 4\}$	1	$0.5N + 0.0$	
$\{\bar{1}, 0, 1\}$	2	$0.333333N + 0.111111$	[10, 11]
$\{\bar{1}, 0, 1, 3\}$	3	$0.285714N + 0.163265$	[8]
$\{\bar{1}, 0, 1, 5\}$	6+	$0.276458N + 0.236886$	
$\{\bar{1}, 0, 1, 7\}$	7+	$0.281354N + 0.282566$	
$\{\bar{3}, \bar{1}, 0, 1\}$	3	$0.285714N + 0.163265$	[8]
$\{\bar{3}, 0, 1, 3\}$	5	$0.277778N + 0.209877$	
$\{\bar{3}, 0, 1, 5\}$	6+	$0.285887N + 0.351390$	
$\{\bar{3}, 0, 1, 7\}$	6+	$0.273322N + 0.258510$	
$\{0, 1, 3\}$	2	$0.333333N + 0.111111$	
$\{0, 1, 3, 6\}$	2	$0.333333N + 0.111111$	
$\{0, 1, 5\}$	6	$0.324503N + 0.259024$	
$\{0, 1, 3, 5\}$	3	$0.285714N + 0.163265$	[8]
$\{0, 1, 3, 7\}$	5	$0.277778N + 0.195988$	

maximise the number of zeros was found by performing a Markov analysis on its recoding state machine. The results of these trials are presented in Table 3.

The symbol alphabet with the lowest average MHW is $\{\bar{3}, 0, 1, 7\}$ although to achieve this a window width of $m = 6+$ is required to achieve this. The symbol alphabets of $\{\bar{1}, 0, 1, 3\}$, $\{\bar{3}, \bar{1}, 0, 1\}$, and $\{0, 1, 3, 5\}$ have a similar average MHW but only require a window width of 3. Their state machines will be much smaller.

Symbols that are a radix multiple of another symbol do not reduce the Hamming weight. This can be seen by comparing the average MHW of the symbol alphabets $\{0, 1, 2\}$ and $\{0, 1, 2, 4\}$ with binary $\{0, 1\}$ where 2 and 4 are radix-2 multiples of 1. Also with the symbol alphabets $\{0, 1, 3\}$ and $\{0, 1, 3, 6\}$, where 6 is a radix-2 multiple of 3. This is expected as the radix multiples provide the same recoding, just left shifted one position.

7 CONCLUSION

Minimum Hamming weight representations allow a designer to reduce calculations in multiplication like algorithms. To take advantage of this an algorithm is required to generate MHW representations quickly. This paper presented a new method to design a state machine to perform the conversion from non-redundant representation to redundant representation. The new method improves on previously published algorithms by allowing a more diverse symbol alphabet. From the target number system's

radix and symbol alphabet specification a search is performed for all MHW representations of values in the range $[0, \beta^M)$ to generate a table of possible recodings. This table is reduced to a table mapping from a window of the input word to a substitution symbol such that the table is as small as possible and the MHW recodings are maintained. This mapping table is used to build the recoding a state machine that serially transforms the non-redundant to redundant representations. A Markov analysis of the state machine reveals the average Hamming weight of the recoded numbers which match some previously published results.

REFERENCES

- [1] L. D'Addario, A. Gunst, J. Bunton, and A. Boonstra, "SKA signal processing," *SKA Memo 91*, p. 98, Aug. 2007.
- [2] Chip-Hong Chang, Yajuan He, and Jiangmin Gu, "An alternative scheme of redundant binary multiplier," in *Proceedings of the 2004 IEEE Asia-Pacific Conference on Circuits and Systems*, pp. 33–36, 2004.
- [3] Behrooz Parhami, "Carry-free addition of recoded binary signed-digit numbers," *IEEE Transactions on Computers*, vol. 37, pp. 1470–1476, 11 Nov. 1988.
- [4] W. Kamp, A. Bainbridge-Smith, and M. Hayes, "Efficient implementation of fast redundant number adders for long word-lengths in FPGA," *Field Programmable Technology, 2009, International Conference on*, Dec. 2009.
- [5] B. Parhami, "Generalized signed-digit number systems: a unifying framework for redundant number representations," *IEEE Transactions on Computers*, vol. 39, no. 1, pp. 89–98, 1990.
- [6] A. Avizienis, "Signed-digit number representation for fast parallel arithmetic," *IRE Trans. Elect. Comput.*, vol. EC-10, pp. 389–400, Sept. 1961.
- [7] H. Gundersen, R. Jensen, and Y. Berg, "A novel ternary switching element using CMOS recharge semi floating-gate devices," in *Proc. 35th Int. Symp. Multiple-Valued Logic*, pp. 54–58, Citeseer, 2005.
- [8] B. Phillips and N. Burgess, "Minimal weight digit set conversions," *IEEE Transactions on Computers*, vol. 53, pp. 666–677, June 2004.
- [9] U. N. Bhat, *An introduction to queueing theory : modeling and analysis in applications*. Birkhauser, 2008.
- [10] Richard I. Hartley, "Subexpression sharing in filters using canonic signed digit multipliers," *IEEE Transactions on Circuits and Systems—Part II: Analog and Digital Signal Processing*, vol. 43, pp. 677–688, 10 Oct. 1996.
- [11] S. Arno and F.S. Wheeler, "Signed digit representation of minimum Hamming weight," *IEEE Transactions on Computers*, vol. 42, pp. 1007–1010, Aug. 1993.

Advantages of 3D Time-of-Flight Range Imaging Cameras in Machine Vision Applications

Adrian A. Dorrington*, Cameron D.B. Kelly, Shane H. McClure,
Andrew D. Payne, and Michael J. Cree

*Department of Engineering, University of Waikato
Private Bag 3105, Hamilton, New Zealand.
Email: a.dorrington@waikato.ac.nz

Abstract: Machine vision using image processing of traditional intensity images is in wide spread use. In many situations environmental conditions or object colours or shades cannot be controlled, leading to difficulties in correctly processing the images and requiring complicated processing algorithms. Many of these complications can be avoided by using range image data, instead of intensity data. This is because range image data represents the physical properties of object location and shape, practically independently of object colour or shading. The advantages of range image processing are presented, along with three example applications that show how robust machine vision results can be obtained with relatively simple range image processing in real-time applications.

Keywords: Range Imaging, Three-dimensional, 3D, Machine Vision

1 INTRODUCTION

Machine vision has been successfully used in industry in applications such as factory automation. One of the keys to its success is tight control of environmental conditions, particularly lighting control as well as object and background colours. There is also a demand for machine vision in other applications, such as people tracking and security, where it is difficult to control the environment and the objects. These applications, which are abundant in the literature, are less successful because of the difficulty in separating the wanted object from the background. The problem is confounded when the objects of interest have high contrast features, such as a person wearing a light shirt and dark trousers.

The main limitation in the current machine vision object detection problem is essentially the camera technology. Sophisticated algorithms have been developed to detect objects in the 3D world using a 2D camera, thus ignoring the most obvious of useful information, depth. Attempts have been made with reasonable success to use stereopsis (multiple two dimensional cameras with triangulation techniques) to gain the depth information[1]. However, this approaches needs substantial computing power and has limited operating volume (defined by the intersection of the field of views of the cameras).

Time-of-flight (TOF) range imaging is a relatively new imaging technique that can be loosely described as a 3D camera. Such cameras produce a digital photograph or video like output that contains both depth and (monochrome) intensity information simultaneously for every pixel in the image. These cameras have great potential for machine vision applications because they natively capture the depth information that must be inferred with traditional imaging systems. Because the depth information captured with these cameras is (for the most part) independent of intensity and object colour, a machine vision system operating primarily on depth information can separate object from the background with relatively simple algorithms.

We introduce the basic operating principle of range imaging systems, highlight their advantages and disadvantages, and demonstrate some simple applications illustrating how they can be used for machine vision.

2 BACKGROUND

2.1 Range camera operating principle

There are several different techniques that can be used to acquire range images, but by far the most common is the continuous wave amplitude modulated (CWAM)

approach[2, 3, 4, 5]. As this is the only technique used in the currently available commercial products, it is the only one we describe here.

Range imaging is an active image approach. The scene is flood illuminated with an intensity modulated light source. Near infrared LEDs are the illumination source of choice in commercial products with modulation frequencies ranging from 10 MHz to 50 MHz. Laser diodes are used in research devices at modulation frequencies up to 100 MHz, and may make an appearance in future commercial devices because a higher modulation frequency provides better range measurement precision.

The image sensor gain is also modulated at the same frequency as the illumination source. Digital or square wave modulation is often used for both the illumination and sensor gain, so the sensor gain modulation can be thought of a high-speed shuttering during the integration period. This sensor gain modulation makes the pixel sensitive to the phase of the modulation envelope of the illumination scattered from the scene and collected by the camera lens. A change in phase (due to propagation delay) manifests as a change in pixel brightness. The measured phase change is converted to propagation time hence range (with knowledge of the speed of light). Light collected from objects close to the camera generate a bright pixel due to a small propagation delay resulting in a small phase change. More distance objects generate a dull pixel because of a large propagation delay that results in a large phase change.

Of course, pixel brightness is also influenced by object colour and background lighting conditions, so one capture is insufficient to accurately determine distance. Typically four captures are used in which the relative phase difference between the illumination and shuttering is stepped in 90° increments. This allows calculation of the background intensity, b , active intensity, a , and phase shift of the modulation envelope, φ , with

$$a = \frac{\sqrt{(A_0 - A_2)^2 + (A_1 - A_3)^2}}{2} \quad (1)$$

$$b = \frac{A_0 + A_1 + A_2 + A_3}{4} \quad (2)$$

$$\varphi = \tan^{-1} \left(\frac{A_0 - A_2}{A_1 - A_3} \right) \quad (3)$$

where $A_0, A_1, A_2,$ and A_3 are the pixel brightness values from the four successive images. The object range can then be calculated from the phase, modulation frequency, f , and the speed of light, c , with

$$d = \frac{\varphi c}{4\pi f} \quad (4)$$

2.2 Advantages and disadvantages

Using range imaging cameras in machine vision applications has the significant advantage of providing distance

information, thus allowing computer systems to perceive the world in 3D. It also provides the advantage of generating separate background and active intensity images. This may be of advantage in some applications because the active intensity image provides grayscale information that is not influenced by ambient lighting conditions and shadows.

These advantages come with a price. There are currently several significant disadvantages associated with range imaging cameras[6], of which the most most obvious is the need for active illumination. This increases power consumption and physical size, complicates thermal dissipation, but perhaps most importantly, limits the useful operating distance of the cameras. Objects too distant can be poorly illuminated leading to low quality range measurements. This is especially the case in a scene that contains both near and far objects because the illumination levels and camera integration time must be controlled to avoid saturation of the near, bright objects, resulting in very dull distant objects.

Another significant disadvantage is the low spatial resolution of commercially available range imaging cameras. At the time of writing, the highest resolution camera available is the CamCube 2.0 (PMD Technologies GmbH, Siegen, Germany) with 204×204 pixel resolution. However, as with traditional imaging technology, resolution is increasing steadily with each new model offering higher resolution as the technology matures.

Range imaging cameras also suffer from several range measurement artefacts. At the edge of objects light from two or more objects at different ranges may fall on the same pixel causing erroneous measurements. These errors are referred to as mixed pixels or flying pixels. Also, because of the cyclic nature of the phase measurement used to calculate distance, ambiguities can arise if objects occur at distances longer than half the wavelength of the modulation signal. Finally, significant motion can cause corrupt range data because the scene may change during the four acquisitions required to produce one range image. Research is being conducted to solve these shortcomings, and it is likely that future range imaging cameras will include features to minimise or resolve all of these effects.

3 RANGE IMAGE PROCESSING

Many traditional image processing techniques can be applied to range images[7, 8]. Indeed, operating on the range data not only has the potential to simplify operations such as object detection, but also opens new opportunities for more advanced and qualitative measurements. With range data it is possible to determine physical object parameters such as size and surface curvature. Consequently, range data image processing is likely to be much more robust and reliable than traditional image processing.

To illustrate the simplicity of range data image processing, we present a simple example. Figure 1 shows images acquired with the SwissRangerTMSR4000 camera (MESA Imaging, Zuerich, Switzerland). The active intensity image, figure 1a, is similar to a traditional greyscale image but using the built in light source. Note that the object in the foreground has colourful high contrast surfaces and is in front of a visually complex background. The range image, figure 1b, instantly provides additional information about the scene. One of the first noticeable aspects to this image is that the object and the background are of low contrast compared to their appearance in the intensity image, because there is little difference in distance from the camera across their surface.

Segmenting the object in this scene based on the intensity information presents a challenge for traditional image processing, but is a relatively simple task when processing the range data. For example, consider an edge detection on the intensity image, as shown in figure 1c. Because of the complicated nature of the scene, this detection provides little useful information. However, the same processing on the range data, figure 1d, clearly finds the edges of the gnome in the foreground and the bookshelf in the background.

The other aspect that is immediately apparent from the range image is that the gnome is closer to the camera than the bookshelf as darker pixel shades represent range values closer to the camera. This information in particular cannot be determined in the intensity image. In fact, it is possible to determine the shape, size, and relative location of objects in the scene because the SR4000 also provides a 3D output consisting of x , y , and z values calibrated in metres for each pixel.

Segmentation can be performed very simply in one of two ways. Firstly, if this image is the only data available, simple thresholding can be applied to the range image to segment the objects based on their distance to the camera, as shown in figure 1e. Alternatively, simple background subtraction can be used if a range image of the the scene before the objects were added is available.

4 EXAMPLE APPLICATIONS

To illustrate the potential uses of range image processing, we present three examples. In all cases, the image processing code was written in MATLABTM, and is sufficiently simple to operate in real-time on a 2 GHz computer under Microsoft Windows XPTM.

4.1 Height estimation

The first example is an algorithm that identifies objects and estimates their height off the ground. This is primarily aimed at measuring a person's height. Although there are




Figure 1: Intensity image (a), range image (b), edge detection on intensity image (c), edge detection on range image (d), and threshold segmented object mask (e) of a gnome in front of a bookcase.



Figure 2: Screen shot of the height detection software showing a scene containing two people, each with their height and average distance from the camera displayed.

some constraints on object size and dimensions, the software does not discriminate between people and other similar sized objects. The height of a person (or other object) can be determined as long as the top of the object remains in frame, regardless of their distance from the camera. An example output of the final application is shown in figure 2.

The image processing algorithm starts by capturing the background scene without any people. This forms the background image for all subsequent captures and is valid as long as the scene geometry remains unchanged. For each capture a foreground mask was calculated by thresholding the difference between the current capture and the background image, that is a background subtraction algorithm. This marks the people in the scene. Any edges of the scene are then detected by thresholding the gradient of the depth image. These edge pixels were deleted from the foreground mask to remove any mixed pixels and to help distinguish between overlapping groups. In order to measure more than one person at a time, the foreground mask was separated into groups using the watershed transform. The pixels in each group were then analysed to determine the maximum height and average distance from the camera for each person.

4.2 Gesture control user interface

Another example that highlights the advantages of range imaging is a gesture control application. In this application, the user can activate and control a menu system, press virtual buttons, and move virtual slider controls simply by moving an empty hand in free space. A screen shot from the application is shown in figure 3. The action of pressing a button, for example, is very difficult to detect with traditional cameras because movement towards the camera produces very little change in the image. However,



Figure 3: A user interacting with a push button menu by moving their hand in free space.

with depth data it is very easy to detect if an object moves towards or away from the camera.

The range image processing algorithm first uses background subtraction to isolate a potential user. Then, range thresholding produces a mask indicating all pixels that are within 0.2 m from the closest pixel to the camera, thereby segmenting the closest object to the camera. Morphological image processing techniques are used to remove any outlying pixels in the mask. The centre of the mask is found, and the 3D location of the centre is tracked. The location and motion of the centre is used to select and control objects in the user interface.

A push button and a slider bar are two examples of controls available. The push button works by tracking the center of the mask, a button push is recorded when the center is over a button graphic and then moves forward by a configurable distance over at least 5 frames. The slider bar works by tracking the position of the centre relative to the position of the slider. When the centre is close, to the slider's position tracks with the centre. The user in figure 3 is selecting a push button in an interactive menu.

4.3 Robotic arm location and control

The final example demonstrates range imaging being used to identify and calibrate a low cost robotic arm. Range imaging is used because this type of robotic arm is inherently inaccurate in its positioning and has no feedback to provide control compensation. Two ping-pong balls were added to the arm to provide 3D reference markers, and were detected using range image processing techniques. Intensity and range images of the robotic arm are shown in figures 4 a and b. First background subtraction is used to segment the robotic arm. The locations of the ping-pong balls were then identified using a watershed segmentation technique (applied to the range data), as seen in figure 4c, and their 3D centres calculated from the 3D data, as shown in figure 4d. A model of the arm was then aligned with the centres of the ping-pong balls to relate the arms actual position and orientation to the programmed movement.

5 CONCLUSIONS

We have shown how applying image processing techniques to range data, rather than intensity data, can have significant advantages. This is because range data represents the location and shape of objects in the scene, practically independently of their colour or shade. Comparisons of simple edge detection on both intensity and range data highlight marked differences in the outputs.

Three examples of applications using range image processing are also presented. These examples highlight how range image processing can be easily used in real-time applications such as user interfaces, object size detection possibly including security applications, and control system feedback.

REFERENCES

- [1] T. Acharya and A. K. Ray, *Image Processing: Principles and Applications*. John Wiley and Sons, 2005.
- [2] D. Carnegie, M. Cree, and A. Dorrington, "A high-resolution full-field imaging system," *Review of Scientific Instruments*, vol. 76, p. 083702, 2005.
- [3] A. Dorrington, A. Payne, R. Conroy, M. Cree, and D. Carnegie, "Achieving sub-millimetre precision with a solid-state full-field heterodyning range imaging camera," *Measurement Science and Technology*, vol. 18, pp. 2809–2816, 2007.
- [4] J. P. Anthes, P. Garcia, J. T. Pierce, and P. V. Dressendorfer, "Nonscanned lidar imaging and applications," vol. 1936, pp. 11–22, SPIE, 1993.
- [5] R. Lange and P. Seitz, "Solid-state time-of-flight range camera," *Quantum Electronics, IEEE Journal of*, vol. 37, pp. 390–397, Mar 2001.
- [6] M. Hebert and E. Krotkov, "3d measurements from imaging laser radars: How good are they?," *IVC*, vol. 10, pp. 170–178, 1992.
- [7] A. A. Dorrington, M. J. Cree, D. A. Carnegie, A. D. Payne, R. M. Conroy, J. P. Godbaz, and A. P. P. Jongenelen, "Video-rate or high-precision: a flexible range imaging camera," vol. 6813, p. 681307, SPIE, 2008.
- [8] R. Tanner, M. Studer, A. Zanoli, and A. Hartmann, "People detection and tracking with tof sensor," *Advanced Video and Signal Based Surveillance, IEEE Conference on*, vol. 0, pp. 356–361, 2008.




Figure 4: Images of the robotic arm showing the intensity image (a), range image (b), watershed with identified ping-pong balls (c) and 3D data wireframe with fitted shaded spheres (d).

Circular Polarised Patch Antenna for Medical Microwave Holography

Oleksandr Sergieiev¹
& Vyacheslav Kitaev²

Engineering Research Institute
Auckland University of Technology
Private Bag 92006, Auckland 1142, NEW ZEALAND.

¹Email: oleksandr.sergieiev@ieee.org

²Email: slava.kitaev@ieee.org

Abstract: The paper provides an insight into development of the antenna element for the microwave holography for early breast cancer diagnostics. A circular polarised patch antenna has been modelled with in software. Several examples of the antenna have been manufactured and tested. It has been demonstrated that coupling of such antenna elements is within the requirements for the planned array of 12-16 antennas within 300 mm distance from each other.

Keywords: Microwave imaging, patch antenna, antenna array

1 INTRODUCTION

Active microwave imaging exploits non-ionizing electromagnetic waves with the radiated power level which is much lower than that from a conventional mobile phone. Although due to much longer wave length direct microwave imaging do not provide the same spatial resolution as X-rays, it has the potential to offer improved detectability of cancer on early stages due to the high dielectric contrast between normal and cancerous tissue (C. Gabriel, Gabriel, & Corthout, 1996; S. Gabriel, Lau, & Gabriel, 1996a, 1996b). The method is based on irradiating a monochrome wave, which is then reflected from a tumour and scattered back. The scattered radiation is received by an array of antennas then digitized at each antennae element of the array. The image is the reconstructed using the phase information of backscattered signal.

In this work we have analysed the requirements to an antennae element needed for such an array to receive backscattered signal preserving the phase information.

2 ANTENNA ELEMENT

Linearly polarized antenna can accept only the component of electric field with the same orientation as the antenna. Therefore due to unavoidable rotation of polarization of the reflected wave, using of the linearly polarized antennas will significantly decrease amount of energy collected on the receiving antennas. While it is possible to increase power of transmitted signal to compensate losses due to this effect, this will lead to additional energy dissipation into the human tissue what is undesirable. To avoid this, a circularly polarized (CP) antenna type was selected for the antenna element as it will collect more electromagnetic energy crossed the antenna aperture.

CP can be obtained if two orthogonal modes are excited with a 90° time-phase difference between them (Balanis, 1997). Some types of antenna are inherently CP e.g. Spiral (Stutzman & Thiele, 1997) or Conical Spiral antennas

(Hertel & Smith, 2002). Other types like Cross-Bowtie (Yun, Fear, & Johnston, 2005), Cross-Vivaldi (Zhang, Fear, & Johnston, 2009), Square Patch (Balanis, 1997) can be excited with a 90° time-phase difference between two antenna elements to produce CP by means of special feeding circuitry. All this antenna types have their own strengths and weaknesses, the detailed discussion of which is outside of the scope of this paper due to the limited volume.

Further we will describe the details of the design and optimisation of CP patch antenna which was chosen for the array in our application.

2.1 THEORETICAL BACKGROUND

CP patch antenna can be accomplished by adjusting the physical dimensions of a patch and using either single or multiple feeds (Balanis, 1997). The most direct way is to use two separate orthogonal feeds to excite a square-patch antenna. In this work a single feed to generate two orthogonal modes on the nearly square patch has been used. To meet the necessary conditions for circular polarization, the dimensions of the patch and position of the feeding point must ensure that the magnitudes of the two modes are equal. The initial patch is represented on Figure 1.

To select initial dimensions for our antenna model we have to consider bandwidth, efficiency and quality factor. There is always a trade-off between them in deriving an optimum antenna performance. The total quality factor is influenced by radiation, conduction, dielectric and surface wave losses (Balanis, 1997).

Quite narrow bandwidth 3.3% was selected for current antenna to achieve better total quality Q_t in that case for VSWR 2:1 in accordance to (Balanis, 1997)

$$\frac{\Delta f}{f_0} = \frac{1}{Q_t \sqrt{VSWR}} \quad (1)$$

where f_0 is middle frequency

and Δf is difference between frequencies f_L and f_W

related to effective dimensions of our patch L_e and W_e correspondently




Figure 1. Near square patch antenna sketch and variable dimensions for selected basic antenna design

And side ratio between patch sides

$$L = W \left(1 + \frac{1}{Q_t} \right) \quad (2)$$

where L is length

and W is width of the patch

Thus total quality will be

$$Q_t = \frac{1}{0.033\sqrt{2}} = 21.43 \quad (3)$$

And required ratio between patch sides

$$\frac{L}{W} = 1 + \frac{1}{Q_t} = 1.047 \quad (4)$$

Due to fringing effects electrically the patch looks greater than its physical dimensions Figure 2.




Figure 2. Effective patch dimensions are greater than its physical dimensions

Figure 2 shows that effective dimensions of the patch are

$$L_{eff} = L + 2\Delta L \quad (5)$$

$$\text{and } W_{eff} = W + 2\Delta W \quad (6)$$

Thus physical dimensions of the patch could be derived from required frequencies dielectric and magnetic properties of materials.

$$L = \frac{1}{1f_L \sqrt{\epsilon_{reff}} \sqrt{\mu_0 \epsilon_0}} - 2\Delta L \quad (7)$$

$$W = \frac{1}{1f_W \sqrt{\epsilon_{reff}} \sqrt{\mu_0 \epsilon_0}} - 2\Delta W \quad (8)$$

Required ΔL and ΔW could be derived from empirical equations (Kirschning, Jansen, & Koster, 1981)

$$\frac{\Delta L}{h} = \frac{\xi_1 \xi_3 \xi_5}{\xi_4} \quad (9)$$

With

$$\xi_1 = 0.434907 \frac{\epsilon_{eff}^{0.81} + 0.26 \left(\frac{W}{h} \right)^{0.8544} + 0.236}{\epsilon_{eff}^{0.81} + 0.189 \left(\frac{W}{h} \right)^{0.8544} + 0.87} \quad (10)$$

$$\xi_2 = 1 + \frac{\left(\frac{W}{h} \right)^{0.371}}{2.358\epsilon_r + 1} \quad (11)$$

$$\xi_3 = 1 + \frac{0.5274 \arctan(0.084 \left(\frac{W}{h} \right)^{0.371})}{\epsilon_{ref}^{0.9236}} \quad (12)$$

$$\xi_4 = 1 + 0.0377 \arctan(0.067 \left(\frac{W}{h} \right)^{1.456}) (6 - 5 \exp(0.036(1 - \epsilon_r))) \quad (13)$$

$$\xi_5 = 1 - 0.218 \exp\left(-\frac{7.5W}{h}\right) \quad (14)$$

To determine actual dimensions of the antenna we have to know property of the substrate as well as properties of the media. The Rogers TMM10 electrodeposited copper foil $17\mu\text{m}^1$ was selected for the model because it has $\epsilon_r = 9.2$ what is very close to supposed propagation media, fatty tissue (Duck, 1990; C. Gabriel et al., 1996; S. Gabriel et al., 1996a, 1996b). Substrate with the thickness of 0.508 mm has been selected as trade-off between bandwidth, and efficiency.

Using media and substrate with the same properties gives a possibility to decrease amount of variables in the model by approximately 4:1 and calculation time by 16:1. Moreover number of variables is restricted in the modelling tool we used and without above mentioned condition modelling is provided on the cutting edge of the modelling tool ability. To be picking ahead there are other reasons to use media and substrate with the same properties however they are beyond of the current paper scope.

Now we can evaluate the wavelength for the selected frequency $f = 4.95\text{GHz}$, taking in account relative permittivity of the media $\epsilon_r = 9.2$ where our signal is travelling.

$$\lambda = \frac{c}{f \sqrt{\epsilon_r \mu_r}} = 19.96\text{mm} \quad (15)$$

Equations (11), (12), (13), (14) have been developed by (Kirschning et al., 1981) for the patch antenna working in the air or vacuum and have to take in account relative

¹ TMM® Thermoset Microwave Laminates (2008, 30.09.2009). Retrieved 30.03.2009, 2009, from <http://www.rogerscorp.com/documents/728/acm/TMM-Thermoset-laminate-data-sheet-TMM3-TMM4-TMM6-TMM10-TMM10i.aspx>

permittivity of the substrate. In our case it is possible to simplify these equations due to implied equity of the relative permittivity of the substrate and surrounding media. In the simplified form

$$\xi_1 = 0.434907 \frac{1 + 0.26 \left(\frac{W}{h}\right)^{0.8544} + 0.236}{1 + 0.189 \left(\frac{W}{h}\right)^{0.8544} + 0.87} \quad (16)$$

$$\xi_2 = 1 + \frac{\left(\frac{W}{h}\right)^{0.371}}{3.358} \quad (17)$$

$$\xi_3 = 1 + 0.5274 \arctan\left(0.084 \left(\frac{W}{h}\right)^{0.371}\right) \quad (18)$$

$$\xi_4 = 1 + 0.0377 \arctan\left(0.067 \left(\frac{W}{h}\right)^{1.456}\right) \quad (19)$$

Initial dimensions of the patch L and W represented in the Table 1 were obtained by iterative decision of the equations (10), (8) and (9) for target middle frequency 4.95GHz and propagation media with the $\epsilon_r = 9.2$. Coordinates of the feeding point p_L and p_W as point on the diagonal of the patch on the middle between centre of the patch and corner. Size of the feeding point d has been selected in accordance to diameter of the central cord of the feeding cable RG405². External dimensions of the patch A are taken as reasonable in relation to the patch size.

Table 1. Initial dimensions for patch antenna

L	9.26mm
W	8.844mm
p_L	2.565mm
p_W	2.461mm
d	0.5mm
A	16mm

2.2 ESTABLISHING THE MODEL

Modelling has been done using MATLAB Antenna Toolbox (MAT) which is based on the Method of Moments (Makarov & Kempel, 2005). This tool is suitable for the modelling basic metal-dielectric antennas and resonators. Method of Moments relies on Rao-Wilton-Glisson edge elements (Rao, Wilton, & Glisson, 1982). Initially the surface of a metal antenna is divided into separate mesh of triangles (see Figure 3) to establish matrix of the vector basis functions associated with every interior edge (i.e., no boundary edge).




Figure 3. 3-D model of patch antenna. All dimensions are shown in m.

The detailed explanation for the Method of Moments and its implementation is given in (Makarov, 2002).

The multipoint MoM solution is performed on every point of the frequency. The result for complex input impedance derived for the first model with dimensions from the Table 1 is represented on the Figure 4.

2.3 Optimization

There are two picks apparently visible on the Figure 4 pointed by vertical arrows $fr(L)$ and $fr(W)$ related to effective dimensions of the patch L_e and W_e correspondently. Middle frequency is lower than desirable. One can see difference between impedances for these two frequencies what lead to narrowing of the circularly polarisation zone due to broken condition of the orthogonal components equity.




Figure 4. Modelled parameters for initial dimensions given in the Table 1. Two resonance picks pointed by arrows $fr(L)$ and $fr(W)$ related to effective dimensions of our patch L_e and W_e correspondently.

Multiple iterations with dimensions of the patch and feeding point placement have been done to acquire desirable middle frequency 4.95GHz equity for the active part of impedance $R_{in}=50\Omega$ for both resonance frequencies as it is shown on the Figure 5.

² Cable,microwave,copper,1m,RG405. (2009). Retrieved 30.01.2009, 2009, from <http://newzealand.rs-online.com/web/search/searchBrowseAction.html?method=searchProducts&searchTerm=388704>




Figure 5. Final results of the mathematical modelling. Active impedance $R_{in}=50\Omega$ equal for both resonance frequencies

The final result of mathematical modelling for the reflection coefficient provided by Matlab is presented on the Figure 6. Dimensions derived during optimization process presented in the Table 2.




Figure 6. Final results of the mathematical modelling. Reflection coefficient is below -20db

Table 2. Final dimensions for the patch antenna model.

L	9mm
W	8.6mm
p_L	3.55mm
p_W	1.53mm
d	0.5mm
A	16mm

It is possible to improve mathematical model particularly to place negative pick of the reflection coefficient at 4.95GHz strictly between resonance frequencies by playing with thickness and other dimensions, but it would be reasonable on the later stage of the project. Model parameters shown in the Table 2 is sufficient to provide experiments on the real antennas.

3 Results

Antenna elements (Figure 7) were produced from the TMM10 substrate with thickness 0.508mm in accordance to the dimensions given in the Table 2. The outer sheath of the cable RG405 is soldered to the ground plane of the patch antenna and central wire is soldered to the patch. SMA connectors³ were soldered to the opposite end of the cable.




Figure 7. Antennas used in measurements. Dimension of the grid cell is 5x5 mm.

As the antenna element was designed for the measurement needed to be done in the media with close dielectric properties. Two component silicone rubber RHODORSIL® RTV-585⁴ mixed with Barium Titanate powder⁵ was used as the measurement media mixed in proportion 27% of the Barium Titanate and 73% of the RHODORSIL® RTV-585 according to (Cherney, 2005).

Vector network analyser Agilent N5230A was used to measure impedances, reflection coefficient and phase variations. Measurements have been done for single and a couple of antenna elements submerged in the media. The results acquired for impedances of the single antenna elements are shown on Figure 8. The results acquired for the reflection coefficient with the same antenna element are shown on Figure 9. The results for the phase variations are shown on Figure 10. Special interest represented by phase shift around resonance frequency. Multiple measurements of the impedances for the same antenna element have been done while second antenna element was placed at different distances from the measured antenna element with both antennas being excited. Results acquired for several distances between antenna elements from 16mm to 46mm with step 1mm are shown on the Figure 11.

³ Connector,coaxial,RF,SMA,solder,jack,cable,50 Ohm,RG405/U. Retrieved 30.06.2009, 2009, from <http://newzealand.rs-online.com/web/search/searchBrowseAction.html?method=getProduct&R=4683205>

⁴ RHODORSIL® RTV-585. (1998). Retrieved 30.09.2009, 2009, from http://www.precisionconverting.com/downloads/rhodia_techdata/pi585.pdf

⁵ Barium Titanium Oxide, BaTiO₃. (2009). Retrieved 27.07.2009, from <http://www.advancedmaterials.us/5622-ON2.htm>

4 Discussion




Figure 8. Result acquired with network analyser. The Impedance Resistance and Reactance components. Middle frequency shown by the vertical line is 4.84GHz




Figure 9. Reflection coefficient. Middle frequency is shown by vertical line is 4.84GHz




Figure 10. Measured phase variations. Resonance frequency 4.84GHz is shown by vertical line.




Figure 11. Coupling. Dependence of the Real and Imaginary parts of the mutual impedance from the distance between two antenna elements. Horizontal dashed line represents value for the Imaginary part for the single antenna and solid line correspondent value for the Real part.

Although as it was expected the manufactured antennas were somewhat different from the model, it sufficiently close to the given requirements for the microwave holography system. As one can see the middle frequency shown by vertical line on Figure 10 is 4.84GHz what is only 2% difference from the modelled. It is reasonable to suggest that that these deviations are due to the limited accuracy of the used manufacturing method. The same reason could lead to relatively less than expected reflection coefficient (see Figure 9) in comparison with the computed model (see Figure 6).

Coupling between the antennas will determine a constrain for the minimal distance between antenna elements in the array which defines the lower angular resolution α as:

$$\alpha = \frac{\lambda}{D} \quad (20)$$

where λ is the wavelength

and D is separation between the elements in array.




Figure 12. Dependencies of linear resolution for two frequencies 5 and 10 GHz from the distance between antenna elements

The same is plotted on Figure 12 for two selected frequencies 5 and 10 GHz.

Coupling is a result of undesirable energy transferring between neighbour antennas. Therefore coupling leads to measurable changes in the antenna parameters. Measurements of the impedances for the same antenna element on the resonant frequency 4.84GHz have been done while second antenna element was placed on the different distance from the measured antenna element and both elements exited. The results for Real and Imaginary parts of the impedance are shown on Figure 11. The horizontal dashed line represents a value for the Imaginary part for the single antenna and solid line represents the correspondent value for Real part. By the recalculating of the phase deviation from the collected data we can get dependency for possible error value as it is shown on the Figure 13.




Figure 13. Dependency for possible error value in the phase measurements due to mutual coupling from the distance between two antenna elements.

5 FURTHER WORK

As it was mentioned above there are two possible ways to improve linear resolution in this holography method. One is to reduce distance from the antenna array to the object under investigation. However there are limitations related to big angular deflections. Second is to increase the frequency as one can see on Figure 12. Unfortunately 10 GHz perhaps a limit due to very high absorption of microwaves above that what leads to very little penetration. In the future work 10 GHz can be attempted.

Comparing mathematical model with measurement results it is clear that real antenna elements especially feeding point placement required additional optimization. It can be achieved by using more sophisticated modelling software and by producing more samples for measurements. Either way several iterations of modelling-measurement should be done to produce reliable antenna setup.

4 CONCLUSIONS

Although the parameters of developed antenna are slightly different from one optimized in the computational model, the acceptable parameters of CP patch antenna have been achieved at 5 GHz. Thus the used methodology can be trailed at other higher frequencies observing the accuracy of manufacturing.

Antenna element requirement analysis and design discussed in this work required additional development. However it could be perceived as starting point for further development of the microwave holography technique for early breast cancer detection.

5 REFERENCES

Balanis, C. A. (1997). *Antenna Theory Analysis and Design* (2nd ed.). New York: WILEY-INTERSCIENCE.

Cherney, E. A. (2005). Silicone rubber dielectrics modified by inorganic fillers for outdoor high voltage insulation applications. *IEEE Transactions on Dielectrics and Electrical Insulation*, 12(6), 1108-1115.

De Carvalho, S. A., & De Souza Mendes, L. (1999). Scattering of EM waves by inhomogeneous dielectrics with the use of the method of moments and 3-D solenoidal basis functions. *Microwave and Optical Technology Letters*, 23(1), 42-46.

Duck, F. A. (1990). *Physical properties of tissue*. London, UK: Academic Press Inc

Gabriel, C., Gabriel, S., & Corthout, E. (1996). The dielectric properties of biological tissues: I. Literature survey. *Physics in Medicine and Biology*, 41(11), 2231-2249.

Gabriel, S., Lau, R. W., & Gabriel, C. (1996a). The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz. *Physics in Medicine and Biology*, 41(11), 2251-2269.

Gabriel, S., Lau, R. W., & Gabriel, C. (1996b). The dielectric properties of biological tissues: III. Parametric models for the dielectric spectrum of tissues. *Physics in Medicine and Biology*, 41(11), 2271-2293.

Hertel, T. W., & Smith, G. S. (2002). Analysis and design of two-arm conical spiral antennas. *Electromagnetic Compatibility, IEEE Transactions on*, 44(1), 25-37.

Kirschning, M., Jansen, R. H., & Koster, N. H. L. (1981). AC-CURATE MODEL FOR OPEN END EFFECT OF MICROSTRIP LINES. *Electronics Letters*, 17(3), 123-125.

Kulkarni, S., Lemdiasov, R., Ludwig, R., & Makarov, S. (2004). Comparison of two sets of low-order basis functions for tetrahedral VIE modeling. *IEEE Transactions on Antennas and Propagation*, 52(10), 2789-2794.

Makarov, S. N. (2002). *Antenna and EM Modeling with MATLAB: WILEY-INTERSCIENCE*.

Makarov, S. N., & Kempel, L. C. (2005, 2005). *MATLAB Antenna Toolbox. A draft*. Retrieved 10.10.2008, 2008, from <http://ece.wpi.edu/mom/>

Partial Differential Equation Toolbox (2009). Retrieved 1.10.2009, 2009

Rao, S. M., Wilton, D. R., & Glisson, A. W. (1982). ELECTROMAGNETIC SCATTERING BY SURFACES OF ARBITRARY SHAPE. *IEEE Transactions on Antennas and Propagation*, AP-30(3), 409-418.

Stutzman, W. L., & Thiele, G. A. (1997). *Antenna Theory and Design* (2-nd ed.): John Wiley and Sons.

Yun, X., Fear, E. C., & Johnston, R. H. (2005). Compact Antenna for Radar-Based Breast Cancer Detection. *Antennas and Propagation, IEEE Transactions on*, 53(8), 2374-2380.

Zhang, J., Fear, E. C., & Johnston, R. H. (2009). Cross-Vivaldi antenna for breast tumor detection. *Microwave and Optical Technology Letters*, 51(2), 275-280.

Current New Zealand Activities in Radio Astronomy: Building Capacity in Engineering & Science for the Square Kilometre Array

M. Johnston-Hollitt^{*}, V. Kitaev[†], C.P. Hollitt[‡], N. Jones[§], T.C. Molteno[¶]

^{*}School of Chemical & Physical Sciences, Victoria University of Wellington
Box 600, Wellington, NEW ZEALAND.
Email: Melanie.Johnston-Hollitt@vuw.ac.nz

[†]School of Engineering, Auckland University of Technology
Private Bag 92006, Auckland, NEW ZEALAND.

[‡]School of Engineering & Computer Science, Victoria University of Wellington
Box 600, Wellington, NEW ZEALAND.

[§]Department of Physics, University of Otago
Box 56, Dunedin, NEW ZEALAND.

[¶]Department of Computer Science, University of Auckland
Private Bag 92019, Auckland, NEW ZEALAND.

Abstract:

We present an update on the NZ-wide advances in the field of Radio Astronomy & Radio Engineering with a particular focus on contributions, not thus reported elsewhere, which hope to either directly or indirectly contribute to New Zealand's engagement with the international Square Kilometre Array (SKA) project. We discuss the status of the SKA project in New Zealand with particular reference to activities of the New Zealand Square Kilometre Array Research and Development Consortium.

Keywords:

Radio Astronomy, Square Kilometre Array in New Zealand, Radio Engineering.

1 INTRODUCTION

The Square Kilometre Array¹ (SKA) is a billion dollar, multinational project which aims to construct the world's largest radio telescope [1, 2]. A consortium of 17 countries has pledged involvement in this mega-science project which aims which will collect more information on the radio sky in the first seven hours of operation than has been obtained in the first 70 years of radio astronomy. The SKA will be comprised of thousands of individual radio telescopes, with a total collecting area of one square kilometre. These telescopes will be arranged in stations of 20 - 45 antennas. The stations will stretch over thousands of kilometres in order to provide resolution of the order of milliarcseconds which will be comparable to the best optical and infrared instruments. The array will be sited in either Australasia or southern Africa and bids to host the telescope are being coordinated by South Africa and Australia with both countries committed to building 'pathfinder' telescopes on the sites of the proposed SKA core in each country.

¹www.skatelescope.org

Realization of the truly paradigm-shifting nature of the SKA will require significant challenges in radio engineering, data transport & storage, signal processing and power generation to be overcome. Much progress has been made on many of these fronts but there is much still to do and even seemingly small communities can play a vital role in the project. In this paper we describe the current status of the project in New Zealand with particular emphasis on the research engagement from Aotearoa.

2 SKA IN AUSTRALIA

Australian researchers have been involved in the SKA project since its inception over a decade ago. Radio astronomy is the largest single discipline in Astrophysics in Australia accounting for over 30 per cent of the total astronomical research output in Australia [3]. Radio astronomy research groups currently exist in seven Australian universities in addition to the government funded CSIRO division of the Australia Telescope National Facility and radio astronomy accounted for nearly 40 per cent of all permanent jobs in Australian Astronomy in 2005 [3]. With the

advances toward the SKA, this fraction is likely to expand significantly.

To date the Australian Federal Government has committed 125 million AUD for construction of their pathfinder telescope, the Australian Square Kilometre Array Pathfinder (ASKAP) and 80 million AUD for a High Performance Computing (HPC) facility in Perth². In addition, the Western Australian Government has committed a further 20 million AUD to establish the International Centre for Radio Astronomy Research (ICRAR)³ as a joint venture between the state and the University of Western Australia and Curtin University of Technology.

3 SKA IN NEW ZEALAND

In August 2009 New Zealand and Australia signed an arrangement to collaborate on the bid to host the SKA⁴. This marked the first formal step of New Zealand's engagement with the project and will see a joint trans-Tasman effort to work together to co-host the telescope.

3.1 NZ SKA Project Office

The Ministry of Economic Development (MED) is the lead government agency for New Zealand's involvement in the SKA. A New Zealand SKA project office has been established as New Zealand's point of contact for SKA matters and is coordinating New Zealand's Government level programme for engagement in the SKA project. As part of the coordination effort the project office is working closely with the Australia-New Zealand SKA Coordinating Committee (ANZSCC) on which the NZ Government has a representative.

3.2 SKA Research & Development Consortium

In June 2009, just prior to the announcement of the formal arrangement with Australia, the New Zealand Square Kilometre Array Research & Development Consortium (SKARD) was formed⁵. The role of SKARD is to bring together professional researchers involved in SKA related research in NZ, to foster collaboration both within NZ and internationally and to liaise with industry, government and other groups to advance New Zealand's contribution to the SKA. Members are drawn from all of New Zealand's major research universities and have interests in antenna design, signal processing, imaging and inference, high performance computing and radio astronomy.

²www.ska.gov.au/news/Pages/MinisterCarrnamesPawseyHPCCentre.aspx

³www.icrar.org

⁴Arrangement between the Australian Government and the New Zealand Government on a Joint Bid to Secure the Siting of the Square Kilometre Array in Australia and New Zealand'

⁵www.ska.ac.nz

3.3 NZ SKA Industry Consortium

In addition to SKARD, the New Zealand SKA Industry Consortium (NZ SKAIC) has formed to achieve positive economic outcomes for New Zealand from involvement in the SKA project. The group consists of industry partners representing IT software, hardware, networks and services, the Ministry of Economic Development (MED) and New Zealand Trade and Enterprise (NZTE).

4 RESEARCH ENGAGEMENT IN NEW ZEALAND

Although SKARD acts as a coordination point for researchers engaged in SKA related research, research engagement at the individual level has been quite active over 2009 even before the formation of the consortium. Several groups from across the country have come together to form collaborative projects around radio astronomy and related engineering. Highlights of this collaborative work include efforts in radio astronomy, radio engineering and signal processing.

4.1 Radio Astronomy

Currently there are only two universities in which radio astronomy research is conducted in New Zealand. Consequently, direct involvement in pure radio astronomy research is limited⁶, however where such engagement has occurred it has been well placed. In particular, NZ has good representation on forth coming major science and associated technical research associated with the Australian SKA Pathfinder, pan-NZ capability building in radio astronomy, signal processing and engineering and an emerging interest in associated High Performance Computing (HPC).

4.1.1 ASKAP Surveys

Australia has committed to build a 1 per cent demonstrator for the SKA called the Australian Square Kilometre Array Pathfinder (ASKAP). ASKAP, which is in construction now and due for completion in 2012, will be a powerful new telescope in its own right. Located at the core of the Australian & New Zealand site for the SKA, the Murchison Radio Observatory in Western Australia, ASKAP will comprise 36 15 meter diameter radio dishes spread over a 6 km diameter area. The configuration is designed to optimize sensitivity on angular scales of 30 arcseconds, whilst still providing good low surface brightness sensitivity at lower resolutions [4]. ASKAP will operate over 700 - 1800 MHz with a one degree field of view and had

⁶Refereed research publications in radio astronomy for NZ based researchers in the last decade are confined to only two authors: Johnston-Hollitt & Budding.

been specifically designed as a survey instrument for red-shifted observations of neutral hydrogen (HI) [5]. ASKAP will act as a survey instrument for approximately 75 per cent of the first five years of operation (2013 - 2018) and an open call for survey science proposals in late 2008 resulted in 38 expressions of interest from world-wide teams of researchers. An extensive, multi-stage process of international peer review to prioritize the science surveys has just concluded with eight science surveys and two strategic programs selected[6].

Several NZ-based researchers are involved in the successful surveys which were selected for ASKAP. Staff from the Victoria University of Wellington (VUW) are involved in three ASKAP proposals, the A ranked Evolutionary Map of the Universe (EMU)[7] and the A-ranked surveys "POSSUM" and "VAST". Staff from the Auckland University of Technology are part of the strategic project "The high resolution component of ASKAP: meeting the long baseline specifications for the SKA".

Within the commensally observed EMU and POSSUM projects, researchers at VUW will contribute to science related to clusters of galaxies. In particular, understanding the role of environment on the production of radio emission in clusters will be investigated. This will include: the detection and characterization of low surface brightness diffuse synchrotron emission associated with dynamical encounters [8], the use of tailed radio galaxies as barometers of "cluster weather" [9] and as probes of high density regions [10, 11], and an investigation into the properties of AGN and starforming galaxies in different cluster environments [12, 13, 14]. The latter two projects are relatively straightforward in terms of finding and identifying sources. Detection of diffuse, low surface brightness sources however will be a challenge and is likely to require new signal processing techniques. Consequently researchers from VUW are undertaking a broader program of research in signal processing for astronomical research to address the challenge (see [15]).

The particular interest of NZ members of the VAST team is investigating coherent emission from active stars and binary systems. Such observations will probe fundamental parameters of stellar magnetospheres such as magnetic field intensity and topology and electron energy distribution and provide vital information on the relation of this emission to the physical characteristics of the host stars.

4.1.2 Radio Telescopes in New Zealand

There have been a range of substantial low-frequency arrays and small higher frequency ($\nu \leq 3$ GHz) dishes in NZ over the last 50 years. In particular, there has been a long history of polytechnics associated with small dishes including a 10m dish at Hamilton built in 1985 by the Waikato Institute of Technology in association with




Figure 1: Interference thresholds as a function of frequency for continuum observations with several different configurations of radio telescopes. The line for total power describes the spectral power flux density ($\text{dB}(\text{W m}^{-2} \text{Hz}^{-1})$) for a single dish radio telescope. Subsequent lines give the threshold for harmful interference for different configurations of interferometers going from more compact (VLA(D)) to least compact (Merlin) configurations. The least stringent requirements are the present recommendations for Very Long Baseline Interferometry (VLBI). Taken from the International Telecommunication Union Handbook on Radio Astronomy [19]. The outer stations of the SKA, such as those potentially in New Zealand, will be required to conform with VLBI levels [20, 21].

Waikato University[16] and a 5m dish constructed in 1998 by the Central Institute of Technology in Upper Hutt [17]. (See [16] and references therein for a full discussion of the history of small radio dishes in NZ.) The most recent in this long line of small dishes is the AUT 12m dish at Warkwork, located just north of Auckland [18]. Although this dish, which has been constructed but is currently still undergoing testing, will primarily undertake geodetic observations [18], it is hoped it will contribute to science observations of bright radio sources through the Australian Long Baseline Array (LBA).

Additionally, there are several low-frequency arrays operating in NZ for the purposes of undertaking upper atmospheric or meteor detection research, which are not connected to radio astronomy research.

4.1.3 Low Frequency Transient Network Development

Funding has recently been obtained to build a low frequency receiver network deployed across NZ to investigate properties of the transient radio sky. The project Transient Radio Emission Array Detector Prototype (TREAD-P) [22] will employ FPGA devices connected across the Kiwi Advanced Research and Education Network (KAREN) to

test collection and storage of very high data rates and novel techniques for data processing using Bayesian inference with the hope of characterizing the low frequency radio sky. The core of the detector network will be the Digital Receiver Sensor (DRS) which is based around Field Programmable Gate Array (FPGA) technology, which is becoming increasingly important to radio astronomy. Depending on the configuration, each DRS will produce up to 10MB of data per second and the network of 10 devices will produce more than 8.6 Terabytes daily when running at maximum capacity.

Major international radio telescopes like the Low Frequency Array (LOFAR) [23] and ASKAP have identified characterization of the transient radio sky as a high science priority and proposed instruments like the Long Wavelength Array plan to conduct science specifically on the transient radio sky (20 - 80 MHz) in the next decade. Instruments like the DRS will provide vital information on the frequency of both artificial and naturally occurring radio signals which will be vital to the survey design of these large low-frequency telescopes.

In order to achieve this goal researchers from the Auckland University of Technology, University of Otago, Victoria University of Wellington, University of Auckland and the University of Canterbury and industry partners have had agreed to contribute resources. It is expected this will provide opportunities for at least three research students at partner institutions across New Zealand in 2010.

4.2 HPC Facilities

New Zealand has capability in High Performance Computing, with dedicated facilities at the University of Canterbury (BlueFern) and computation clusters in several universities and government research labs. Many of these facilities are coordinated via its national grid initiative BeST-GRID led by the University of Auckland and seven other member institutions, which operates over a 10 Gb/s national research network the Kiwi Advanced Research and Education Network (KAREN), providing essential services for collaboration, computation, and data management.

NZ-based researchers are currently utilizing these tools to develop capability in relation to the SKA. As an example, the low frequency transient project (Section 4.1.3) involving researchers from five institutions has been recently funded to exploit this infrastructure to transport, store and process data at high rates from a network of GPS-synchronized low frequency antennas across NZ for radio astronomy research.

Additionally, researchers at three of NZs universities (Otago, Auckland and AUT) have established a collaboration developing new high performance computing techniques for processing radio astronomy data with a specific focus on reconfigurable computing. AUT, Otago and

VUW maintain HPC clusters and in addition AUT and Otago have reconfigurable FPGA based systems (based in physics and engineering respectively) while BeSTGrid (Auckland) provides an ideal middleware platform to develop techniques for distributed computations which will be required for SKA research.

While this is a foundation, it should be noted however that host countries for the SKA will require significant upgrades to e-Infrastructure including a distributed high-speed fibre optic network between the stations and the science processing facility, mega-watt power for SKA stations and specialised hardware for the central digital signal processing facility. Depending on how the SKA is configured many of these resources may be located in NZ as part of the international project.

5 DISCUSSION: IMMEDIATE GOALS

5.1 Station Locations

In partnering with Australia to host the SKA, New Zealand is hoping to contribute the outer stations of the array. At this stage this is likely to mean two stations of antennas in NZ, though arguments for more stations could be made in the future. In order to achieve the science goals of the SKA even the outer stations of the array will require low levels of radio frequency interference (RFI)[21]. Figure 1 shows the currently required levels specified for outer stations of the SKA in the low to mid frequency bands[20]. While there have been a couple of suggested locations put forward for SKA stations none of these meet the required specifications.

It is probably worth noting here that the vast majority of RFI is man-made and easiest way to achieve the radio quiet levels required for the SKA is to place the antennas in regions of low population density. Figure 2 shows the current population density of NZ as determined from the 2006 national census⁷ in addition to Telecom's current mobile phone coverage⁸. Regions of low population density and no mobile coverage exist in much of the South Island and there are some pockets in the North Island with similar characteristics.

At present, a comprehensive and independent site selection and survey process in close collaboration with Australia is currently being negotiated through the NZ SKA Project Office. As with all telescope site selection processes this will examine a range of metrics before finalising potential sites.

5.2 Research Engagement

With the signing of the arrangement with Australia and formation of SKARD and NZSKAIC, New Zealand is now

⁷<http://en.wikipedia.org/wiki/File:NewZealandPopulationDensity.png>
⁸www.telecom.co.nz



Figure 2: LHS: Population density of NZ from the 2006 Census⁷; RHS current mobile phone coverage for NZ⁸. To reduce the levels of RFI, potential SKA stations should ideally be located in regions of low population density with little mobile coverage.

on a firm footing for researchers from the private and public sector to engage in this project. As the final host nations for the array will not be decided until the early part of 2012 it is important for NZ to maximise benefit in the project via early engagement while simultaneously mitigating the risk of that the project will go to Southern Africa. This means that it is important to identify niche opportunities which are independent of the ultimate location of the array. This can be achieved via expanding existing links between NZ and Australia in radio astronomy related areas via:

- i) continued participation of NZ scientists in the science and design of the SKA,
- ii) building expertise in NZ through the training of graduate students in radio astronomical research - a natural complement to NZ's existing strength in optical astronomy, and
- iii) undertaking significant and ground breaking research which will be used for further enhancement of science with next generation radio telescopes.

The latter speaks directly to the identification and exploitation of niche research opportunities. Such niches have clearly emerged in NZ around science-driven post-processing via novel signal processing [15]. Work of this nature combines NZ's strengths in radio astronomy, signal processing, HPC & middleware. Additionally, capacity building exercises and pan-NZ collaborative work [22] which has commenced will be important to realize opportunities that will arise with hosting the telescope.

On the purely astronomical research front, in order to maximize the science outcomes of both ASKAP and eventually the SKA it is crucial that knowledge from existing

instruments is examined and used to develop expertise vital for survey design on next generation instruments and the research teams thus far involved are well placed to make significant contributions.

ACKNOWLEDGEMENTS

MJ-H acknowledges continuing support towards radio astronomy research by the Victoria University of Wellington.

REFERENCES

- [1] R. T. Schilizzi, P. E. F. Dewdney, and T. J. W. Lazio, "The Square Kilometre Array," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 7012 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Aug. 2008.
- [2] J. Lazio, "The Square Kilometre Array," in *Panoramic Radio Astronomy: Wide-field 1-2 GHz research on galaxy evolution*, Proceedings of Science, p. (PRA2009)058, Sept. 2009.
- [3] B. Gibson, S. Maddison, R. Webster, M. Wals, T. Tzioumis, H. Sim, S. Ricketts, and B. Zealey, "A Demographic study of Australian Astronomy," in *New Horizons: A Decadal Plan for Australian Astronomy 2006-2015*, vol. 2, Nov. 2006.
- [4] N. Gupta, S. Johnston, and I. Feain, "ASKAP Array Configurations: Technical Studies," *ATNF SKA Memo Series*, vol. 016, 2008.
- [5] L. Staveley-Smith, "Resolution and confusion in ASKAP HI Surveys," *ATNF SKA Memo Series*, vol. 020, 2008.
- [6] L. Ball, "ASKAP Survey Science Projects," *ATNF Memo*, 2009.
- [7] R. P. Norris and the EMU team, "ASKAP-EMU: Overcoming the challenges of wide deep continuum surveys," in *Panoramic Radio Astronomy: Wide-field 1-2 GHz research on galaxy evolution*, Proceedings of Science, p. (PRA2009)033, Sept. 2009.
- [8] M. Johnston-Hollitt, *Magnetic fields & diffuse radio emission in A3667 and other Southern rich galaxy clusters*. PhD thesis, University of Adelaide (Australia), July 2003.
- [9] M. Johnston-Hollitt, M. Fleenor, J. Rose, W. Christiansen, and R. W. Hunstead, "Radio imaging of head-tailed galaxies from an in-falling filament of the Horologium-Reticulum supercluster," in *IAU Colloq. 195: Outskirts of Galaxy Clusters: Intense Life in the Suburbs* (A. Diaferio, ed.), pp. 423-425, July 2004.

- [10] M. Y. Mao, M. Johnston-Hollitt, J. B. Stevens, and S. J. Wotherspoon, "Head-tail Galaxies: beacons of high-density regions in clusters," *Monthly Notices of the Royal Astronomical Society*, vol. 392, pp. 1070–1079, Jan. 2009.
- [11] M. Y. Mao, R. Sharp, D. J. Saikia, R. P. Norris, M. Johnston-Hollitt, E. Middleberg, and J. E. J. Lovell, "Wide-angle tail galaxies in ATLAS," *Monthly Notices of the Royal Astronomical Society*, p. (submitted), 2009.
- [12] T. Venturi, S. Bardelli, G. Zambelli, R. Morganti, and R. W. Hunstead, "Radio properties of the Shapley Concentration - IV. The A3528 cluster complex," *Monthly Notices of the Royal Astronomical Society*, vol. 324, pp. 1131–1146, July 2001.
- [13] T. Venturi, S. Bardelli, M. Zagaria, I. Prandoni, and R. Morganti, "Radio emission from the A3571 cluster complex: The final stage of a cluster merger?," *Astronomy & Astrophysics*, vol. 385, pp. 39–54, Apr. 2002.
- [14] M. Johnston-Hollitt, M. Sato, J. A. Gill, M. C. Fleenor, and A. Brick, "Radio observations of the Horologium-Reticulum supercluster - I. A3158: excess star-forming galaxies in a merging cluster?," *Monthly Notices of the Royal Astronomical Society*, vol. 390, pp. 289–303, Oct. 2008.
- [15] C. P. Hollitt and M. Johnston-Hollitt, "Automatic Detection of Supernova Remnants using the Circle Hough Transform," in *Proceedings of the 16th Electronics New Zealand Conference* (T. Molteno, ed.), Nov. 2009.
- [16] M. Head, "A Brief History of Radio Astronomy in Aotearoa New Zealand," www.astronomy2009.org.nz/radio, accessed October 2009, 2006.
- [17] A. W. Dodson and E. Budding, "The Development of the Radioastronomy Project at the CIT," *Southern Stars*, vol. 39, pp. 21–28, Sept. 2000.
- [18] S. Gulyaev and T. Natusch, "New Zealand 12-m VLBI Station for Geodesy and Astronomy," in *NASA Technical Publication -2009-214183: International VLBI Service for Geodesy and Astrometry 2008 Annual Report* (D. Behrend and K. Baver, ed.), pp. 68–73, July 2009.
- [19] International Telecommunication Union Recommendations, "Handbook on Radio Astronomy," *International Telecommunication Union, Geneva (Section 2.3)*, 1995.
- [20] International Telecommunication Union Recommendation, "Protection criteria used for radio astronomical measurements," *ITU-R RA769.2*, 2003.
- [21] SKA Task Force on Regulatory Issues, "Spectrum Protection Criteria for the Square Kilometre Array," *SKA Memo Series*, vol. 073, 2005.
- [22] V. Kitaev, A. Perera, and G. Souldenkov, "Digital Receiver for Transient Radio Emission Array Detector Prototype (TREAD-P)," in *Proceedings of the 16th Electronics New Zealand Conference* (T. Molteno, ed.), Nov. 2009.
- [23] H. Röttgering, A. G. de Bruyn, R. P. Fender, J. Kuijpers, M. P. van Haarlem, M. Johnston-Hollitt, and G. K. Miley, "LOFAR: a new radio telescope for low frequency radio observations: science and project status," in *Texas in Tuscany. XXI Symposium on Relativistic Astrophysics* (R. Bandiera, R. Maiolino, & F. Mannucci, ed.), pp. 69–76, 2003.

Automatic Detection of Supernova Remnants using the Circle Hough Transform

Christopher Hollitt^{*}, Melanie Johnston-Hollitt[†]

^{*}School of Engineering and Computer Science
Victoria University of Wellington
PO Box 600, Wellington 6140
Email: chollitt@ieee.org

[†]School of Chemical and Physical Sciences
Victoria University of Wellington
PO Box 600, Wellington 6140
Email: Melanie.Johnston-Hollitt@vuw.ac.nz

Abstract: The next generation of radio telescopes currently being developed will generate unprecedented volumes of data. A range of automatic source detection algorithms will be required to extract interesting objects from the datasets, as the current method of manual searching will be infeasible. In this paper we discuss a framework for the automatic detection of supernova remnants using Hough transforms. We demonstrate the success of the technique in detecting a representative sample of supernovæ morphologies and verify its superiority over conventional detection methods, potentially allowing for detection of heretofore unknown supernova remnants.

Keywords: SKA, Hough Transform, Supernova Remnants

1 INTRODUCTION

When a massive star at the end of its life becomes a supernova, an energetic shock wave is released from its surface. As the shock travels outward it accelerates electrons in the interstellar medium to produce emission at radio frequencies. In the simplest case the resulting emission would be spherically symmetric as the spherical shock passes through a homogeneous interstellar medium. However, where the shock wave is asymmetric, or there is existing structure in the gas cloud surrounding the supernova, then more complicated supernova remnants are produced. The symmetry, expansion rate and shape of a supernova remnant thus reveal important information about the supernova that precipitated it and about the interstellar medium into which it is expanding including valuable details of the density and ambient magnetic fields[1, 2].

Our galaxy contains a large number of known supernova remnants (SNRs); the result of supernovæ occurring in roughly the last million years. The Molonglo Observatory Synthesis Telescope Supernova Remnant Catalogue (MSC) [3] contains a list of 75 SNRs from the 843 MHz survey of the southern Galactic plane within the area

$245^\circ \leq l \leq 355^\circ$ in Galactic longitude and $\pm 1.5^\circ$ of the Galactic plane. Statistical studies of the SNR population allows much more information to be inferred than would be possible by examination of one SNR alone. In addition to the valuable information gleaned on the average nature of the ISM surrounding supernovæ, statistical analysis shows that there is presently a significant discrepancy between the total number of known SNRs and that expected from stellar lifecycles [4]. Comparison with the expected rate of supernovæ production from OB star counts, pulsar creation rates, iron abundance and observation of supernova in other galaxies suggests that there should be over 1000 SNRs detectable in the Galactic plane [5, 6]. However at present only of order 250 have been found despite and increasing number of surveys of the region[7, 4]. The detection of the population of SNRs is therefore an important part of radio surveys of the galactic disk.

Previous surveys of the radio sky have relied on manual inspection to find SNRs [3, 8, 4]. This is error prone and often results in objects being missed in the initial data release [3, 7]. While this approach has been adequate (although not ideal) in the past, future generations of radio telescopes will generate such a vast stream of data that

manual inspection will not be practical. The increased resolution and field of view that will become routine when using instruments like the SKA [9, 10], forces a paradigm shift that will require a much greater reliance on automated data extraction techniques.

We propose a two stage method for the automatic detection of SNRs. Candidate objects will be first be selected via a fast conventional blob detection algorithm such as Duchamp¹ or VSAD[11], before being processed with a circle Hough transform to identify any circular features in the data. In addition to SNRs, a variety of astronomical objects can produce arc-like features, examples include the so-called double radio relics [12, 13] and a sub-class of Active Galactic Nuclei (AGN) which have their radio jets bent back due to interactions with the medium, these are known as Head-tailed galaxies[14, 15]. However, as these additional arc-like radio sources are all extragalactic in origin, an intimal separation based on proximity to the Galactic plane should be sufficient as 95 per cent of known SNRs are confined to within $\pm 5^\circ$ of the plane [7]. Although we note, a manual classification stage might still be required to resolve any ambiguous objects at the higher Galactic latitudes ($\leq 20^\circ$).

An overview of this method will be presented in more detail in section 2. Preliminary results from the use of the Hough transform to characterise supernova remnants will then be discussed in section 3. Section 4 will present the discussion and conclusions.

2 METHOD

While a stereotypical SNR contains a complete circular shell, real SNRs often display partial shells or include significant distortions of their spherical symmetry. However, almost all SNRs include features that are circular or arc-like. The use of the circular Hough transform (CHT) [16, 17, 18] therefore provides one possible path for their automatic detection, though there exist a set of pathological SNRs that are likely not to be suitable for detection by this approach. Use of the CHT should isolate the location and scale of all circular features in a set of radio data to provide a catalogue of possible SNRs that can be assessed by an astronomer.

The family of algorithms known collectively as Hough transforms are capable of finding and characterising geometrical objects in image or image like-data [16]. While Hough transforms have seen some use in the astronomical community [19, 20], we are not aware of any previous application of the transform to the problem of source detection.

The Circle Hough Transform is a particular variant of the Hough transform that is tailored to the detection of circular objects. It can find partially occluded circular ob-

jects and is known to be robust to the presence of noise [21]. The CHT takes an image as its input and produces a three dimensional array at its output, where two of the array's dimensions represent the set of possible locations for the centres of circles and the third dimension spans the set of possible circle radii. Peaks in the Hough transform therefore correspond to circles in the input image. The action of the transform is perhaps best envisaged as finding the correlation of the input image with all possible circles that could be present in the image. A peak in the correlation then corresponds to a circle that is the best match to the input data. For our application, if we calculate the CHT of a radio image and then find the peak in the Hough transform we will recover the radius and center of any supernova remnant in the input radio map.

Sadly the CHT is a computationally intensive algorithm, particularly when used to characterise circles with unconstrained radii [22]. The CHT is therefore likely to prove too slow to directly process the vast amounts of raw data that will be produced by the next generation of radio telescopes such as the SKA.

We propose a two tier approach to the detection of supernova remnants. The first stage will require the detection of regions of extended radio emission (that is emission from any objects other than point sources). This initial processing must be fast, but doesn't need to be very accurate. Some variety of blob-detection algorithms would be suitable for this purpose. As a fair amount of work has been carried out by the astronomical community on this problem and standard algorithms such as VSAD exist[11], we will not further discuss the details of this method, but will instead assume the existence of an algorithm that can provide a list of possible sites for a SNR detection. Ideally the blob detection algorithm would additionally provide an estimate of the scale of a possible SNR. However, existing astronomical application of blob detection are poor at detecting broken large scale structures, so are likely to provide a poor estimate of both the size and centre of most SNRs.

3 RESULTS

3.1 Known Supernova Remnants

The MSC survey using the Molonglo Observatory Synthesis Telescope (MOST) provides a representative set of supernova remnants against which the new algorithm can be tested. While the MCS survey was conducted at considerably lower frequencies than will be used for SKA observations, its near complete spatial frequency coverage provides a good SKA analogue for our purposes. Existing telescopes currently operating in the mid to low SKA frequency band ($\nu \leq 3$ GHz) necessarily use a relatively sparse sampling of spatial frequencies, which can result in circular calibration artifacts in their output data. Such

¹www.atnf.csiro.au/people/Matthew.Whiting/Duchamp




Figure 1: Four supernova remnants taken from the MCS catalogue [3]. Subfigure a. is G337.3+1.0, b. is G302.3+0.7, c. is G315.4-2.3 and d. is G317.3-0.2

artifacts would confound our SNR detection algorithm.

We have selected a subset of four representative supernova remnants from the MCS catalogue, as shown in Fig 1. This subset was selected to range from a near complete ring (1a), through two partial rings (1b and 1c) to a distorted example (1d). These four SNRs were intended to provide an increasing challenge for the algorithm.

A conventional two dimensional CHT algorithm was performed on each of the four supernova remnants. The four input data sets were each padded to 243 pixels square. Normally a Hough transform is preceded by an edge detection step to highlight the boundaries of objects. However, supernova remnants are inherently shell-like, so we found the edge-detection step to be unnecessary. Comparison of the transformation of the raw and edge-detected data resulted in differences in only a few pixels in the estimated SNR position and radius.

Fig 2 shows a representative Hough transform of a SNR. For clarity we have presented data from the almost complete ring of G317.3-0.2 in this case. While the Hough transform results in a full three-dimensional array this is difficult to present graphically. The figure therefore shows the magnitude of the Hough transform obtained for a small subset of the possible SNR radii. As can be seen for small radii the intensity pattern is well distributed, indicating that there is no strong peak for these radii and hence no small radius circles in the image. However, in the $r = 64$ plane we can see a strong peak toward the centre of the field. This indicates that the input image contains a cir-




Figure 2: Slices through the Hough transform of G317.3-0.2. Slices are taken through several planes corresponding to a set of different possible SNR radii, $r = (16, 32, 48, 64, 80)$ pixels.

cle with radius approximately 64, which is located toward the centre of the image. For larger radii we again see the intensity pattern is spread, so there are no larger circles present.

The Hough transforms of the four SNRs were searched to find their maxima. These peak in the transforms correspond to the circle that is most consistent with the input data in each case. Extraction of the coordinates of the peaks therefore results in an estimate of the location and radius of the supernova remnants.

The effectiveness of the CHT in identifying the circular structures in the four supernova remnants can be seen in Fig 3. In all four cases the best fitting circles returned by the algorithm are in good agreement with the radio data.

3.2 Robustness to Noise

To explore the robustness of the CHT to noise when targeting supernova remnants we contaminated the supernova remnants in Fig 1 with increasing amounts of Gaussian noise and processed the resulting images until the SNRs could no longer reliably be detected. As an example of this procedure Fig 4 shows the SNR in 1c contaminated with additional noise until the Hough transform detector was 50% reliable. For this object the magnitude of the noise has been increased by a factor of 10 from that in the original image. As shown in the figure, the estimated location of the SNR in the contaminated image is slightly different to that in the uncontaminated image. However,




Figure 3: Results of running the four supernova remnants in Fig 1 through the Hough transform code. The overlaid circles are those that best correlate with each supernova remnant.

as we are more interested in detecting SNRs rather than characterising them, this is not perceived as a significant problem.

Experimentation with different sources showed that the Hough transform technique was able to identify a SNR in the presence of noise roughly twice that of a trained human.

4 CONCLUSIONS AND FURTHER WORK

One of the current problems in understanding the so-called ‘Galactic Ecology’ is the discrepancy between the expected and detected rate of supernovæ. At present it is believed we are missing close to 75 per cent of SNRs and in particular those that fall into two important classes; those which are young and distant and those which are old and nearby [7, 4]. In the former case the issue is simply a lack of resolution in current surveys of the Galactic plane leading to such small angular sized, bright sources to be misclassified as point sources. This issue will be readily addressed and resolved by the next generation of radio telescopes such as the SKA. In the latter case however, the SNRs will be both faint and large and if the present ‘eyeballing’ techniques continue to be used they are likely to remain elusive.

We have presented a framework for the automatic detection of supernova remnants and other arc-like structures in radio astronomical data. The framework uses a two




Figure 4: The input and results when the noise level in the supernova remnant from Fig 1c is increased until the Hough transform method is 50% reliable. The solid line indicates the location of the SNR as estimated by the Hough transform and the dashed circle indicates the location of the SNR as predicted in the absence of additional noise.

stage process, where a fast source location algorithm is first used to locate sources of extended radio emission. A Circle Hough Transform technique is then used to analyse the candidate objects to characterise the locations and parameters of any arc or circle-like features in the radio emission.

We have showed that the Hough transform is effective at extracting objects from an existing survey of the galactic plane. As this survey shares many features with the expected telescopes of the future we have confidence that the method should prove a viable option for future development.

The robustness of the method can be improved by the development of a more effective method for searching the Hough transform for peaks. Simply finding the maximum value as we have done here is unlikely to be the optimal technique. In future the method will therefore be augmented with a more sophisticated algorithm for this part of the technique.

The method described in this paper will then be used for a blind analysis of the entire dataset generated in the Molonglo Galactic Plane Survey of which there are two parts being the first and second epoch surveys, MGPS & MGPS-2 [23, 24]. A catalogue of known supernova remnants has been manually compiled from that survey, allowing ready comparison with the results of the automatic detection. This will importantly allow assessment of the technique’s false positive and false-negative error rates. This comparison also admits the possibility of the detection of as-yet undiscovered supernova remnants and in particular those in the class of being old and faint which have been predicted via several independent indicators but remain, as yet, undetected.

ACKNOWLEDGEMENT

MJ-H acknowledges support for radio astronomy research by the Victoria University of Wellington. The MOST is operated by The University of Sydney with support from the Australian Research Council and the Science Foundation for Physics within The University of Sydney.

REFERENCES

- [1] B. M. Gaensler and P. O. Slane, "The Evolution and Structure of Pulsar Wind Nebulae," *Annual Review of Astronomy & Astrophysics*, vol. 44, pp. 17–47, Sept. 2006.
- [2] S. Orlando, F. Bocchino, F. Reale, G. Peres, and O. Petruk, "On the origin of asymmetries in bilateral supernova remnants," *Astronomy & Astrophysics*, vol. 470, pp. 927–939, Aug. 2007.
- [3] J. B. Z. Whiteoak and A. J. Green, "The MOST supernova remnant catalogue (MSC).," *Astronomy & Astrophysics Supplement*, vol. 118, pp. 329–380, Aug. 1996.
- [4] C. L. Brogan, J. D. Gelfand, B. M. Gaensler, N. E. Kassim, and T. J. W. Lazio, "Discovery of 35 New Supernova Remnants in the Inner Galaxy," *Astrophysical Journal Letters*, vol. 639, pp. L25–L29, Mar. 2006.
- [5] Z. Li, J. C. Wheeler, F. N. Bash, and W. H. Jefferys, "A statistical study of the correlation of Galactic supernova remnants and spiral arms," *Astrophysical Journal*, vol. 378, pp. 93–105, Sept. 1991.
- [6] G. A. Tammann, W. Loeffler, and A. Schroeder, "The Galactic supernova rate," *Astrophysical Journal Supplement Series*, vol. 92, pp. 487–493, June 1994.
- [7] D. A. Green, "Galactic supernova remnants: an updated catalogue and some statistics.," *Bulletin of the Astronomical Society of India*, vol. 32, pp. 335–370, Dec. 2004.
- [8] M. D. Filipović, M. Stupar, P. A. Jones, and G. L. White, "Galactic SNRs in the PMN Survey," in *Neutron Stars in Supernova Remnants* (P. O. Slane & B. M. Gaensler, ed.), vol. 271 of *Astronomical Society of the Pacific Conference Series*, pp. 387–+, 2002.
- [9] R. T. Schilizzi, P. E. F. Dewdney, and T. J. W. Lazio, "The Square Kilometre Array," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 7012 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Aug. 2008.
- [10] J. Lazio, "The Square Kilometre Array," *ArXiv e-prints*, Oct. 2009.
- [11] J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick, "The NRAO VLA Sky Survey," *Astronomical Journal*, vol. 115, pp. 1693–1716, May 1998.
- [12] M. Johnston-Hollitt, *Magnetic fields & diffuse radio emission in A3667 and other Southern rich galaxy clusters*. PhD thesis, University of Adelaide (Australia), July 2003.
- [13] A. Bonafede, G. Giovannini, L. Feretti, F. Govoni, and M. Murgia, "Double relics in Abell 2345 and Abell 1240. Spectral index and polarization analysis," *Astronomy & Astrophysics*, vol. 494, pp. 429–442, Feb. 2009.
- [14] M. Johnston-Hollitt, M. Fleenor, J. Rose, W. Christiansen, and R. W. Hunstead, "Radio imaging of head-tailed galaxies from an in-falling filament of the Horologium-Reticulum supercluster," in *IAU Colloq. 195: Outskirts of Galaxy Clusters: Intense Life in the Suburbs* (A. Diaferio, ed.), pp. 423–425, July 2004.
- [15] M. Y. Mao, M. Johnston-Hollitt, J. B. Stevens, and S. J. Wotherspoon, "Head-tail Galaxies: beacons of high-density regions in clusters," *Monthly Notices of the Royal Astronomical Society*, vol. 392, pp. 1070–1079, Jan. 2009.
- [16] R. O. Duda and P. E. Hart, "Use of the hough transformation to detect lines and curves in pictures," *Commun. ACM*, vol. 15, no. 1, pp. 11–15, 1972.
- [17] D. Kerbyson and T. Atherton, "Circle detection using hough transform filters," in *Image Processing and its Applications, 1995., Fifth International Conference on*, pp. 370–374, 1995.
- [18] H. Birk, L. Grevenkop-Castenskiold, J. Lind, T. Moeslund, K. Miller, and K. Pedersen, "Circle detection by image analysis using the Dynamic Generalized Hough Transform," January 1995.
- [19] P. Ballester, "Hough transform and astronomical data analysis," *Vistas in Astronomy*, vol. 40, no. 4, pp. 479–485, 1996. Vision modeling and information coding.
- [20] A. J. Storkey, N. C. Hambly, C. K. I. Williams, and R. G. Mann, "Cleaning sky survey data bases using hough transform and renewal string approaches," *Monthly Notices of the Royal Astronomical Society*, vol. 347, pp. 36–51, Jan. 2004.
- [21] E. R. Davies, "A modified hough scheme for general circle location," *Pattern Recogn.*, vol. 7, no. 1, pp. 37–43, 1988.

- [22] H. Yuen, J. Princen, J. Illingworth, and J. Kittler, "Comparative study of Hough transform methods for circle finding.," *Image and Vision Computing*, vol. 8, no. 1, pp. 71–77, 1990.
- [23] A. J. Green, L. E. Cram, M. I. Large, and T. Ye, "The Molonglo Galactic Plane Survey. I. Overview and Images," *Astrophysical Journal Supplement Series*, vol. 122, pp. 207–219, May 1999.
- [24] T. Murphy, T. Mauch, A. Green, R. W. Hunstead, B. Piestrzynska, A. P. Kels, and P. Sztajer, "The second epoch Molonglo Galactic Plane Survey: compact source catalogue," *Monthly Notices of the Royal Astronomical Society*, vol. 382, pp. 382–392, Nov. 2007.

A comparative study of ‘inversion’, ‘optimization’, and ‘inference’ as frameworks for imaging

Colin Fox¹ and Nicholas Dudley Ward²

¹ Electronics Group, Department of Physics, University of Otago, PO Box 56, Dunedin, New Zealand

² Groundwater Resources Research, 54 Dudley Street, Christchurch, New Zealand

Email: fox@physics.otago.ac.nz & nick@groundwater.org.nz

Abstract: We give four examples that examine the features, and foibles, of performing imaging by *direct inversion*, *optimizing an objective function*, and *model-based inference*.

Keywords: Imaging, regularization, Bayesian inference, heat transfer, hydrology

1. INTRODUCTION

Millman and Grabel noted in their 1987 textbook [1], that the diverse activities encompassed by the term ‘electronics’ share the common property of processing information. In fact, a definition of modern electronics could be ‘the processing of information represented as electrical signals’.

A central part of information processing is the analysis of measurements. We often interpret measurements in terms of physical models as a means of probing the world around us. The question we address in this paper is: what framework should guide us in doing that analysis? We present four examples that display the features, and foibles, of three common frameworks. We use the term ‘imaging’ to refer to data analysis because the primary unknown in our work is usually a spatially-varying quantity that is usefully displayed as an image.

The three frameworks we compare in this paper are *direct inversion*, *optimizing an objective function*, and *model-based inference*. The tools used in ‘signal processing’ largely come from the first two of these, and have become essential knowledge in electrical and electronic engineering. The third framework is now seen as the ‘gold standard’ in imaging, and we are starting to see the first measurement equipment built around inferential methods.

It might come as a surprise to practitioners in electronics that statisticians have played a key role in developing these frameworks, and are largely responsible for the framework of signal processing as practised in electrical engineering. For example, it was the statistician John Tukey who brought us

the Cooley-Tukey FFT algorithm, and also the word ‘bit’.

The connection between imaging, and statistics arises because every measurement is corrupted to some degree by noise. Hence any measurement process is naturally described using probability, and the task of estimating the unknown true image is naturally one of statistical inference. In fact the discipline of statistics has its origins in such problems, with Laplace’s work in astronomy, and the later work by Jeffreys in geophysics [2].

2. INVERSION

2.1 Example 1: Image deblurring

Figure 1 shows a test image, with 200×200 8-bit gray scale pixels, and the result af-



Figure 1: A gray scale test image (left) and a blurred version (right).

ter blurring using the Matlab command `blurred = conv2(test,ones(5,5)/25,'same');`

Most students of electronics will have taken some courses in mathematics, and so have been exposed

to the idea of inversion of a function. This is by far the most often suggested framework for imaging by those who have little or no experience in practical data analysis. In this section we look at using direct inversion to perform deblurring, and see why it will (always) fail for practical problems.

In this case the blurring, or forward map, is given by a convolution, so the image f and data d are related by

$$d = f * h$$

when h is the ‘point spread function’. A standard result is then that the Fourier transforms F , D , and H of f , d , and h , respectively, are related simply by

$$D = F \times H, \text{ or } F = D/H.$$

In Matlab this deblurring by Fourier division looks like `f = real(ifft2(D./H))`; and the result is shown in the left-hand panel of Figure 2. As you can see, the result is not good.




Figure 2: Left: image evaluated by direct inversion. Right: image evaluated by Tikhonov regularized least-squares inversion, using $\lambda = 0.14$.

2.2 What went wrong with the inverse?

In this case the forward map is invertible, so the route of direct inversion should work. The magnitude of the components of H (which are also the singular values of the forward map) are plotted in Figure 3, and though some singular values are small, none are actually zero or even near machine precision. We can’t blame the failure of direct inversion on noise in the data since we added no noise (apart from round-off error) to the blurred image, so everything should work, right? Well, no. The problem is that the forward map is not *exactly* a convolution, because we trimmed the data to have the same size as the image. More than 5 pixels from the edge of the picture the forward map is exactly a convolution, but it is not quite for a band near the edges. That small *model error* is enough to make direct inversion near useless, with the effect that mis-modelling at the edges propagates catastrophically through the image.




Figure 3: The singular values of the forward map, sorted in descending order.

When the data does contain noise, direct inversion will also fail because the small singular values ensure that the reconstructed image is dominated by noise.

3. OPTIMIZATION OF AN OBJECTIVE FUNCTION

A common route around the problems of direct inversion is to apply a regular approximation to the inverse of the forward map A . The usual way to do that is to minimize an objective function such as

$$Q(f) = \|Af - d\|^2 + \lambda^2 \|f\|^2$$

that balances the data misfit $\|Af - d\|^2$ against a measure of image quality – in this case the Tikhonov regularizing functional $\|f\|^2$. The balance is controlled by the regularizing parameter λ . When $\lambda = 0$ we get usual least-squares fitting to data, which reduces to direct inversion when A is invertible. More generally the minimum is the solution of the *generalized deconvolution equation*

$$(A^T A + \lambda I) f = A^T d.$$

In signal processing this linear equation is known as the Wiener filter.

The usefulness of this approach can be seen in the right panel of Figure 2 that shows the deblurred image evaluated with $\lambda = 0.14$, using the Matlab `F = D.*conj(H)./(abs(H).^2+lambda.^2);`
`f = real(ifft2(F));`

3.1 Stochastic bias in regularized inverses

For non-linear problems, least-squares estimates and regularized estimates suffer from *stochastic bias*, i.e. a systematic offset in estimated values resulting from errors on measurements. Stochastic bias may be displayed in a simple form, as follows: If x is a (scalar) random variable from any

distribution with mean $E(x) = \mu$ and variance σ^2 , then $E(x^2) = \mu^2 + \sigma^2$. That is, in the presence of noise on variable x , the mean of the square equals the square of the mean *plus the variance of the noise*. The bigger the noise, the bigger the offset. Stochastic bias is a very real effect that, for example, is one of the mechanisms that allows profit to be made in volatile markets, whether increasing or decreasing [3]. The following example shows that stochastic bias occurs in (regularized) least-squares estimation.

3.2 Example 2: Estimating heat capacity from impulsive heating

Consider the problem of determining the heat capacity c in a 3-dimensional homogeneous medium of large extent for which the thermal conductivity k is known. We consider the stylized problem where the medium is subject to unit impulsive heating at $r = 0$ and time $t = 0$, and the temperature at the point of heating is subsequently measured. For this problem the forward map is available analytically, with the noise-free temperature at time t

$$T(t) = \frac{c^{1/2}}{(4\pi kt)^{3/2}}. \quad (1)$$

If measurement d_i is made at time $t = it_s$, $i = 1, 2, \dots, K$, and is subject to additive noise with zero mean and variance σ_n^2 , then the least-squares estimate of c is

$$\hat{c} = \arg \min_c \sum_{i=1}^K \left(d_i - b \frac{\sqrt{c}}{i^{3/2}} \right)^2$$

where $b = (4\pi kt_s)^{-3/2}$. The normal equations are solvable in this case, and give

$$\hat{c} = \left[\frac{\sum_{i=1}^K d_i / i^{3/2}}{b \sum_{i=1}^K 1 / i^3} \right]^2 = w^2.$$

The term in the square brackets, denoted w , is the sum of many random variables and hence w is a random variable with mean \sqrt{c} and variance $\sigma_n^2 / \left(b^2 \sum_{i=1}^K 1 / i^3 \right) \rightarrow \sigma_n^2 / (b^2 \zeta(3))$ as $K \rightarrow \infty$. Here $\zeta(3) \approx 1.2021$ is the Riemann zeta function evaluated at 3. Using the result quoted above, for a large (infinite) number of measurements, the least-squares estimate of c has expected value $\hat{c} = c + \sigma_n^2 / (1.2021 \times b^2)$. That is, in the presence of measurement error the least squares estimate is systematically biased, for all K . Note that reducing the measurement error gives less bias. However, increasing the number of measurements makes the estimate worse (not better as is often claimed) by increasing bias.

In this case the bias may be easily removed (by subtracting $\sigma_n^2 / (1.2021 \times b^2)$) since we were able to calculate the bias analytically. The result is a better estimator, having the same variance but lower bias, and is explicitly not the least-squares estimator. Unfortunately the obvious conclusion, that best-fit to data is not the same as best-fit to parameters, is not commonly observed in the imaging literature. In most imaging problems we are not able to determine the bias analytically, making the least squares estimate both biased and difficult to fix. The application of regularization actually compounds this problem. For example, for this stylized example the Tikhonov-regularised estimate may be calculated analytically, and has bias that is dependent on the *unknown* value of c , leaving an implicit problem to remove bias.

It is instructive to note that the quantity w is an unbiased estimator for \sqrt{c} , since the data is a linear function of \sqrt{c} . Hence the least squares estimate of \sqrt{c} makes a good estimate, but its square is not a good estimate of c . This apparently paradoxical behaviour is an example of how the algebra of (random) variables with uncertainty is quite different to the algebra of deterministic variables (see e.g. [4]). For this reason it is necessary to track the *distribution* of possible values that a variable can take, not just the single ‘best’ estimate. Maintaining and summarizing distributions over variables is the basis of Bayesian inference.

Anticipating section 4., we briefly describe how a Bayesian approach could solve this example. The likelihood function combines the forward map in equation 1 and the distribution over measurement noise. As is typical in inverse problems, the range of the forward map is a small fraction of data space, and so the noise statistics may be determined from the measurements. In this case, at large times when $T(t) \approx 0$ the data solely consists of measurements of the noise. Hence the noise distribution can be determined. The form of the likelihood function depends on the noise statistics, and in this way the Bayesian approach can deal with *any* noise distribution. For this simple single parameter estimation problem, an ‘objective Bayesian’ analysis is feasible, by choosing the Jeffreys type [2] prior distribution that is invariant to choice of units for heat capacity, giving $\pi_p(c) \propto c^{-1/2}$. This is an ‘improper’ prior distribution and would require modification if only a few measurements were available, though then the data so poorly constrains c that no approach would give good results. When the data is adequate, the posterior mean gives a suitable point estimate for c . In the ideal case where the noise is independent identically distributed (iid) zero-mean Gaussian, and measurements are very accurate, there is little to choose between the least squares estimate and the posterior mean, except in

computational cost. In most other circumstances the posterior mean does a much better job of estimating the unknown true heat capacity.

3.3 Example 3: Uncertainty in estimates with uniform noise: a counter example to estimators

The ability to calculate data-dependent or posterior variance is a distinct advantage of a Bayesian approach to inverse problems. In contrast, methods such as least squares are justified on the basis of the variance of the *estimator*, i.e. the average variance over all possible measurements. Yet, in many practical inverse problems the variance of the estimator has nothing to do with the uncertainty in parameters estimated from the data set at hand. This example demonstrates that issue in a very simple setting, where the forward map is the identity function with uniform measurement errors.

Consider the simple case where a scalar quantity μ is measured directly, subject to uniform noise, with mean zero and width 2. Then the i^{th} measurement is

$$d_i = \mu + n_i$$

where each $n_i \sim U(-1, 1)$, i.e. is uniformly distributed over the interval $[-1, 1]$. Since $\mu - 1 \leq d_i \leq \mu + 1$ for all i , it follows that $\max\{d_i\} - 1 \leq \mu \leq \min\{d_i\} + 1$. In fact these bounds are exactly what the data tell us about μ . We note that the likelihood function precisely expresses these bounds, and so they are automatically included in a Bayesian analysis.

The least squares estimate of μ from K measurements is easily seen to be

$$\hat{\mu}_{\text{ls}} = \frac{1}{K} \sum_{i=1}^K d_i.$$

It is instructive to note that this estimate can lie *outside* the interval $[\max\{d_i\} - 1, \min\{d_i\} + 1]$, in which case it is not even consistent with the measured data. For $K = 10$, this happens a little more than 30% of the time, so in 1 out every 3 experiments the least squares estimate is not even a possible value. More troublesome, in practice, is that the error often quoted for the least-squares estimate has little to do with the actual uncertainty in the value of μ as determined by the data. From the considerations above, we see that $\mu \in \frac{1}{2}(\max\{d_i\} + \min\{d_i\}) \pm \frac{1}{2}(2 + \min\{d_i\} - \max\{d_i\})$ so the uncertainty is (certainly) $1 + \frac{1}{2}(\min\{d_i\} - \max\{d_i\})$. The mean-square-error for the least squares estimator of $1/\sqrt{3K}$ is often quoted as the error in the least-squares estimate. Note that it is independent of the data. Three simulations for the

case $\mu = 0$ and $K = 2$, returned the values $(d_1, d_2) = (-0.7477, 0.6688)$, $(-0.6112, -0.6136)$, and $(0.6278, -0.0376)$ giving estimates with posterior error ± 0.2918 , ± 0.9988 , and ± 0.6673 . This is sometimes larger, and sometimes smaller than the least-squares error of ± 0.4082 .

Posterior error estimates are particularly informative when recovering spatially-varying parameters such as the conductivity of an inhomogeneous material. It is clear on physical grounds that the spatial dependence of uncertainty in a reconstruction must be data dependent. For example, when a region of low conductivity is surrounded by a region of high conductivity, the outer region shields the inner region from external sources. Hence the conductivity of the inner region cannot be accurately determined from measurements based on external sources. However, if the whole medium has similar conductivity, then energy can flow through all regions, resulting in more accurate estimation. Since the data reflects the distribution of conductivity, the spatially-varying error must be dependent on the data. As mentioned above, the mean square error of the least-squares estimate, or of any other fixed estimator, gives no clue to this effect.

4. BAYESIAN INFERENCE

4.1 Example 4: Aquifer Parameters from Pump Test Data

In this example, we consider the analysis of ‘pump test’ data used to determine the groundwater availability at a site, and to predict the effects of drawing water from a bore at that site. A pump test consists of pumping groundwater from a borehole and monitoring the groundwater level, or drawdown, at a second borehole at some distance. Estimation of aquifer parameters from pump test data is traditionally based on a graphical approach: pump test data is plotted and parameters from a hydraulic model are varied to achieve a ‘best’ fit to the data. With the advent of the personal computer this process is commonly automated using a least-squares fit.

In practice pump test data is corrupted by noise which means that even under ideal conditions two pump tests from the same site will result in the different drawdown traces, and hence different least square parameter estimates. In this example we concentrate on measurement error, in contrast to the model error that dominated in the Example 1. We assume then that we observe noise perturbed measurements d of the ‘true’ drawdown s in a monitoring bore:

$$d = s + \text{noise}. \quad (2)$$




Figure 5: T-walk sampling for the Theis model, after burn-in. Left: output traces of transmissivity T (top) and storage S (bottom). Right: posterior histograms for transmissivity T (top) and storage S (bottom).




Figure 4: Synthetic pump test data for the Theis model

For simplicity we take the noise to be iid Gaussian distributed with mean 0 and standard deviation σ .

We simulated data using the classical Theis hydraulic model [5], parametrized by transmissivity and storage parameters. Let $s = s(r, t)$ be drawdown at a distance r from the pumping well at time t . Then s solves the Theis equation

$$S \frac{\partial s}{\partial t} = T \nabla^2 s + Q \delta(x) \delta(y),$$

where ∇^2 is the Laplacian, S is storage, T is transmissivity, Q is a constant pumping rate and $\delta(\cdot)$

is the Dirac δ -function. The Theis model has the analytic solution

$$s(r, t)|(T, S) = -\frac{Q}{4\pi T} \text{Ei}\left(-\frac{Sr^2}{4Tt}\right) \quad (3)$$

where $\text{Ei}(\cdot)$ is the exponential integral. The notation $s(r, t)|(T, S)$ is read ‘ s of r and t given T and S ’ and indicates that the true drawdown is parametrized by transmissivity and storage parameters.

Equation 3 defines the *forward map* from the parameters (T, S) to noise-free measurements of drawdown. We assume that measurements are made at times t_1, t_2, \dots, t_N that we denote d_1, d_2, \dots, d_N . Under the assumption of Gaussian noise, above, the probability density over measurements has the form

$$\pi(\{d_i\} | T, S) = \frac{1}{(2\pi\sigma^2)^N} \exp\left\{-\frac{\sum_{i=1}^N (d_i - s(r, t_i)|(T, S))^2}{2\sigma^2}\right\} \quad (4)$$

For a numerical example we suppose that we measure noise perturbed measurements of drawdown in an observation well 50 m from a well pumped over three days with a pump rate of $Q = 1,000 \text{ m}^3/\text{day}$. The true transmissivity is $T = 1,000 \text{ m}^2/\text{day}$ and storage is $S = 1 \times 10^{-4}$. Figure 4 shows one set of measurements simulated




Figure 6: Scatter plot showing the joint distribution of parameters

by evaluating the forward map and then adding zero mean Gaussian noise with a standard deviation of 0.005 m.

The posterior distribution over unknown parameters is easily explored using Markov chain Monte Carlo (MCMC) algorithms [6]. Even as recently as five years ago, implementing an MCMC for any problem, including this simple two parameter case, would require learning how to write an MCMC algorithm, and then *tuning* the proposal densities to give reasonable performance. Fortunately now there are a few black-box sampling algorithms available where the user simply provides the (log) target density and the black-box does the rest. Here we use the t-walk [7], which is not the fastest algorithm in terms of computation, but is easily the fastest in terms of time between forming the posterior distribution and having results that solve the problem. Figure 5 shows output traces from the t-walk sampling, and posterior histograms for T and S .

More importantly than finding ‘best’ estimates for the unknown parameters, the sampler output can be used to determine joint distribution over parameters that is implied by the data and model. Figure 6 shows a scatter plot of the joint distribution over S and T , showing that the data asserts a negative correlation.

The purpose of performing pump tests is, ultimately, to predict the effect on the water table of drawing water from the borehole. Since the data and model define a (posterior) distribution over parameters, accordingly the data implies a distribution over possible drawdowns. Figure 7 shows the drawdown prediction in the monitoring well after 50 days of pumping. The variability of prediction could be quantified in terms of a mean and standard deviation. This will give the actual




Figure 7: Histogram over drawdown predictions in the monitoring well after 50 days of pumping

degree of certainty in drawdown. In contrast, the drawdown predicted by the least-squares estimate, as noted earlier, will exhibit a stochastic bias with no indication of accuracy.

5. CONCLUSION

The frameworks compared for imaging of *direct inversion*, *optimizing an objective function*, and *model-based inference* each have advantages and drawbacks. At present the gold standard are the Bayesian inferential methods. These methods are now straightforward to apply, but remain costly in terms of computer time.

6. REFERENCES

- [1] Millman J and Grabel A 1987 *Microelectronics*, McGraw-Hill
- [2] Jeffreys H 1931 *Scientific Inference*, Cambridge University Press
- [3] Sinclair E 2008 *Volatility Trading*, Wiley
- [4] Springer M D 1979 *The Algebra Of Random Variables*, John Wiley & Sons
- [5] Theis C V 1941 The effect of a well on the flow of a nearby stream, *American Geophysical Union Transactions* **22** (3) 734-738
- [6] Liu J S 2001 *Monte Carlo Strategies in Scientific Computing* New York: Springer
- [7] Christen J A and Fox C 2009 A General Purpose Sampling Algorithm for Continuous Distributions (the t-walk), submitted to *Bayesian Analysis*

WaterLogged, an assistive device for Waka Ama

V.A. Fletcher and D.A. Carnegie

School of Engineering and Computer Science
Victoria University of Wellington
P.O. Box 600
Wellington, New Zealand
vincent.fletcher@vuw.ac.nz

Abstract: Team water sports such as Rowing, Kayaking, Canoeing and Waka Ama (Polynesian outrigger canoe) require crew to paddle, or row, in complete synchronicity with one another. In long distance training and racing, fatigue may cause crew members to paddle asynchronously which limits boat speed dramatically.

This report details the design, construction and testing of a handheld waterproof data logging device for use onboard Waka Ama. Designed to provide important training/racing information, the device logs and displays data from two sensors, a GPS and a triaxial capacitive accelerometer. The data of interest include time, date, ground speed, bearing, velocity, roll, pitch and cadence (strokes per minute). Data are logged as comma separated values in a Windows compatible format and may be overlaid onto Google Earth without any additional processing steps. The device also features an LCD Screen, USB programmability, USB charging, and a menu system for calibrating the device and saving settings for future use.

1. Introduction

Waka are Polynesian watercraft, usually canoes which are human-powered (through paddles). Waka Ama is a variant of the Waka in which an outrigger is attached to increase the stability of the (typically) unstable main hull, transforming the boat into an outrigger canoe. Traditionally, Maori Waka Ama were made from wood, today however, they are generally manufactured of composites such as fibreglass, or in some racing boats, carbon fibre [1].

Waka Ama come in various sizes and can be paddled by just a single person, or up to 12. Larger crews do occur for festival events where the large ornately carved, traditional wooden waka are employed, however the purpose of this project is to investigate the racing variants of these craft. In the sprint race events, variants of the wakas include one person, two person and six person craft which can compete in 500 m, 1000 m or 1500 m events. In 1000 m and 1500 m events the boats must complete a 180 degree turn every 250 m. Endurance races include the same number of crew members but distances typically range from 20 km to 30 km. Typical speeds of a 6 person crew will range from 12 km/h for a novice team to around 14.5 km/h for an experienced crew. To reduce fatigue, paddlers will typically change sides every 12 to 20 strokes. This can be a dangerous time for waka, since if a paddler fails to change when his team-mates do, tipping can occur. Unlike rowing there is no coxswain onboard Waka Ama so if paddlers tire and begin to paddle out of time with one another, it is up to the paddlers themselves to correct their mistakes. Asynchronous paddling

seriously limits speed. Indeed this ability to paddle in perfect time is more important than having large/powerful paddlers on the craft. Paddlers can sometimes feel the thrust occurring out of sync, but lack knowledge of who is out of time, and by what margin. Having a real-time diagnostic on-board the craft would be a valuable training aid.

2. Project objective

The goal of this project is to design a device capable of logging performance data for the Waka Ama. Such data would be useful for both displaying in real-time and in post-training analysis. These data include: water speed (not yet implemented), ground speed, position, bearing, tilt from side to side (x axis), cadence (strokes per minute) and paddler synchronisation information. Accuracy of the paddler synchronisation is critical, the sensor used to measure this must be able to detect asynchronous paddling reliably to inform paddlers. Accuracy of ground speed, position and bearing will be limited to the accuracy of the GPS module used in the project (which is limited to +/- 10 m in all locations due to ionospheric effects, and errors in the receiver's almanac). This error is not critical and can be tolerated in the project. The future implementation of a water speed sensor would provide valuable additional training data.

The datalogger developed in this project needs to have an operational time suitable for use in endurance training. The battery life of the device must be able to reliably accommodate operation for four or more hours. Batteries degrade over time

also so the system should facilitate easy battery replacement and this is taken into account when choosing and positioning the battery unit.

As this device is targeted toward water sport, it must be waterproof, and ideally, be able to float. This is a very important aspect of the project as it will govern the size of the device as it must fit into existing, available waterproof cases. Manufacturing a specifically designed case for this project would be restrictively expensive.

Finally, the device must have the ability to be expanded with further features at a later date. Expansions may include additional sensors or software updates. As sensors such as accelerometers and GPS modules become cheaper and more accurate it may be advantageous to swap them and so the board designed must enable a future developer access to these components. The ability to update the onboard software means the device must be easily reprogrammable.

3. Existing Products

3.1 SpeedCoach

The SpeedCoach is a widely used, portable, waterproof device which uses an external impeller to detect water speed and cadence. This device is targeted toward rowers and kayakers, however there is no reason why it would not work for Waka Ama. SpeedCoach includes basic logging of cadence data but does not include a GPS. It is not possible for this device to inform the user if paddlers (or rowers) are in or out of time. SpeedCoach retails for around \$280 NZD.

3.2 The iPhone

The iPhone 3G contains both a GPS module and accelerometers which through a large number of the 75,000+ apps now available may be logged to the iPhone. The iPhone GPS also includes another improvement, the ability to increase its accuracy based on Cell Phone station triangulation, and triangulation of WiFi hotspots. There have been a large number of data logging applications written for the iPhone, however one in particular is applicable to this application, Speedcoach Mobile. Speedcoach Mobile is based largely upon the operation of the SpeedCoach device (above), however it uses the iPhone's accelerometers to detect Strokes/Minute and also logs GPS data which can be overlaid onto Google Maps. This application would be the closest rival to the device being developed, however the extremely high price point would make it unattractive to many consumers. This software is not capable of telling the user if paddlers/rowers are in, or out of time. It also does not include the ability to measure water

speed, so all speed and distance measurements are taken from the GPS which provides ground speed. Also, being on the iPhone platform, it is not easy to include new, or to change existing sensors. Any software developed for the iPhone would be largely confined to the hardware already embedded in the device. Lastly, battery life estimates provided by the Speedcoach mobile developers suggest that iPhone battery life will only support around 3 to 4 hours of continuous operation – only marginally within the required endurance specifications.

SpeedCoach Mobile software sells for \$84.99 NZD. The software requires an iPhone 3G which is \$1179.00 NZD, and a waterproof, floating iPhone Case which can range from \$50-\$80 NZD.

4. The Device Produced

The system developed includes the following key components:

4.1 Microcontroller (MCU)

The AT89C51AC3 ('AC3) contains nearly everything required for this project. It provides adequate Flash Memory for programming (64 Kbytes) and has an onboard 10-bit ADC which enables theoretical resolution of 3.2 mV. It is able to be powered by a supply of between 3 V and 5 V which suits the logic levels of the other devices, and is readily provided by Lithium Ion (LiPoly) batteries. It also contains all of the necessary interfaces in order to be able to communicate with the other devices in the circuit. The MCU's only shortcoming in relation to this project is its single UART (devices such as the AT89C51RE2 contain two) which means that the μ Drive, GPS and RS232-Programming will need to be multiplexed to ensure cross-talk does not occur.

4.2 Accelerometer

There are four main points to consider when selecting an accelerometer device [2]:

1. Analogue vs Digital: Analogue accelerometers output a continuous voltage that is proportional to acceleration. Digital accelerometers (usually) use pulse width modulation (PWM) for their output where the duty cycle reflects the acceleration detected. As the 'AC3 microcontroller contains analogue inputs, an analogue interface is preferred as it is simpler to implement in code.
2. Number of Axes: This project requires two, one facing in the direction of motion (to measure the applied force) the other facing across the surface of the water (to detect the tilt). It is unlikely that vertical motion information will be useful.
3. Maximum Swing: This determines the range of the accelerometer (for example: +/- 3 g). It is generally recommended that for measuring tilts

under Earth's Gravity, anything above +/-1.1 g is more than enough. However, it is not the maximum acceleration that is of interest, but small changes that occur when acceleration is applied.

4. Sensitivity: Sensitivity is of significant importance as the device will need to detect tiny differences in acceleration, initially estimated to be around 0.01 g.

For the reasons outlined above, the ADXL330 (+/-3g) Capacitive Acceleration Sensor (CAS) was purchased. This device provides analogue 3 axis measurements, consumes 200 μ A at 2.0 V (important due to battery life) and provides onboard temperature compensation. Sensitivity is relative to the supply voltage, for instance at $V_{DD} = 3.3$ V, output sensitivity is typically 330 mV/g. As the theoretical resolution of the 'AC3 ADC is around 3.2 mV, the theoretical resolution of the accelerometer is ~ 0.01 g [3]. Only two of the three axes are required for this project, however the third, y, is setup to measure the tilt along the length of the vessel.

4.3 GPS Module

GPS modules handle their own satellite connections and can be setup to output the data required at an appropriate interval [4]. A common protocol for GPS ASCII sentences is that of *National Marine Electronics Association* 0183 (NMEA) [5]. Various NMEA sentences, each of which contain different information, can be output from most GPS modules to the host device, usually via some form of serial communication [6].

For this project, the GPS Module chosen needs to have the following four properties. Firstly, interfacing needs to be done serially via either UART or SPI, as these are the methods of serial data communication available on the selected MCU. Secondly, access to the NMEA Recommended Minimum Sentence (RMC), is required as this provides all of the GPS data required for the project (RMC provides the following data: UTC Time, Latitude, Longitude, Ground Speed, Bearing and Date). Thirdly, the ability to query the GPS for data rather than at a set interval. This would help considerably when additional features are added to the project which makes reacquisition considerably slower. Finally, the device needs to be powered by 3-5 V. The EM-408 GPS Module was chosen for its relatively low price and because it met all of the above requirements.

4.4 MMC interface

To interface with a memory card and program data into FAT16 formatting, a large library of functions would be required which would take a large amount of programming space. Also, the device

would need to handle this processing itself which may take some time and limit the ability to take 10 Hz accelerometer readings. Instead the μ Drive-uSD-G1 was selected to handle the intensive FAT16 formatting required for a Windows readable memory card. The μ Drive is an embedded DOS micro-drive module which features a UART interface for communication. It contains everything necessary to begin reading and writing data to a micro-SD card in FAT16 formatting.

4.5 LCD Screen

Ideally, a high resolution colour touch screen would have been implemented into the system, however time constraints dictated that something much simpler be used. A simple, backlit, 20 \times 4 character HD44780 compliant LCD screen was incorporated. The screen chosen is the EA-DIP204-4, chosen for its availability, applicability and ability to be viewed in direct sunlight.

4.6 User Interface

The attitude behind this device is that very little human input is required for operation. Therefore, the user interface developed is simple, but enables the ability for expansion in the future. Two push-buttons were implemented, one which is used as a 'select' button, and the other as a 'down' button. This enables a user to navigate through, and change settings within a simple menu system.

4.7 Power Supply, Battery and Charging

The operational voltage of all components except the LCD controller is 3.3 V, the controller requires 5 V. Due to the ready availability of 3.7 V LiPoly batteries and charging IC's (MAX1555), a 3.7 V 1100 mAh battery was used for this project. This size was chosen as power consumption estimates (based on datasheets) found the circuit would draw just under 100 mA when fully operational. This battery, in theory should provide 11 hours of power, easily adequate for common 3-4 hour training sessions.

4.8 Programmability:

To permit improvements and future upgrades on this device the ability to program the microcontroller had to be included. This is easily accomplished via a USB interface and so a FT232RL chip was chosen as the UART \rightarrow USB Interface. This device converts the USB differential measurements to the standard Universal Asynchronous Receiver Transmitter on the 'AC3 microcontroller.




Figure 1- Battery charging (3.7V) circuit via USB (5V) and MAX1555, and regulated supply voltages.

4.9 Waterproof Casing

The Otterbox 1601 Case with Magnet Base provided internal dimensions capable of fitting the components required for the device and the board was then produced to match the interior specifications of this case. The magnet base provides a mounting mechanism to any large metallic surface, this could easily be mounted in a Waka Ama canoe.

5. Hardware Layout and Schematics

Due to space limitations, only the schematics for the power supply (Figure 1) could be included. The operational diagram (Figure 2) provides an overview of the device connections.

5.1 Charging Circuit

This section covers the battery, the charging circuit and the regulators that provide the required 3.3 V

and 5 V voltage levels. The schematic for this circuit is provided in Figure 1.

Firstly, the 5 V line available via USB is used as the input voltage to the circuit. This has been implemented because it is assumed that users will want to reprogram the device, and also charge the batteries at the same time. The USB line, therefore provides both charging and programming accessibility via D-/D+ (connections from D-/D+ to FT232RL are not shown in this diagram to avoid cluttering).

The 5 V USB is then run into the MAX1555 chip which controls charging of the battery. The MAX1555 IC can accept input voltages up to 7 V, it then drops these down to 3.7 V which it uses to charge the battery (BAT in Figure 1). The MAX1555 also displays the charging status via the CHG/USB LED. If the light is on, the battery is charging. If it is off, either the battery is charged, or there is no battery connected. The MAX1555 also lowers charge currents as the battery becomes full to prevent overcharging.




Figure 2 - Operational diagram of the WaterLogged device

The power switch (PWR_SW in Figure 1) provides the battery power to two separate regulators, a 3.3 V regulator and a 5 V DC-DC Step Up converter. The 3.3 V regulator chosen is the TPS73633 low-dropout regulator. This component was used because it can accept input voltages as low as 3.4 V (LiPoly output voltages can drop as they discharge). The 5 V DC-DC converter chosen is built on a *breakout* board. It uses the NCP1400 DC-DC converter along with several other components. The breakout board was used as it had been purchased for an earlier projects and it was more economical to continue to use it. This DC-DC converter accepts an input of between 1 and 4 volts and outputs a regulated 5 V line capable of driving up to 100 mA (five times that needed by the LCD screen driver)

5.2 Main Operational Circuit

As expected, the MCU interfaces with all of the connected devices. Because the MCU is communicating with two UART devices the two devices are time multiplexed. Port 1 on the 'AC3 provides the ADC ports and is therefore connected directly to the 3 axes output of the accelerometer. The LCD Screen requires a total of 11 ports, (8 data and 3 control), 3.3 V and 5 V as well as an external potentiometer to adjust the screen contrast. A standard N channel MOSFET has also been included so that the micro can turn on/off the backlight to lower power consumption and increase battery life.

5.3 Completed Device

The initial design of the device was laid out on several breadboards to ensure that various aspects would be compatible, and to write and debug software. The device was then ported to a purpose built PCB board designed to meet the size restrictions of the waterproof casing (Figure 3).

6. Software

A large portion of this project was to develop software which works reliably with the MCU and external components.

In order to achieve the process outlined above, software was written to handle every aspect of the device's operation. First, all initialisations are performed, these are required to setup each of the components correctly. This also includes reading a settings file to reload the settings from the previous use. The system then goes into a loop in which it:

1. Queries the GPS for a data sentence
2. Splits up the sentence into segments and creates pointers to each of the 9 pieces of data



Figure 3 - The completed device inside the water proof case (Otterbox)

3. Applies all corrections to date and time (NZ is UTC+12).
4. Samples all three accelerometer axes at 10 Hz for the duration corresponding to the stroke rate of the boat
5. Detects Strokes; the device detects strokes by checking if 5 of the last 6 z-axis samples have been above a certain value, and uses a timer to calculate stroke rate.



Figure 4 - Logging file uploaded directly from microSD to GPS Visualizer

6. Records all data to microSD card in comma separated value (csv) format which has been




Figure 5 - 8 real-world kayak strokes with overlaid stroke detection

formatted correctly to meet the requirements of *GPS Visualizer* [7] which allows all logged data to be overlaid on Google Earth/Google Maps, Figure 4.

7. Results

Tests of both stationary accuracy, and changes in accuracy with time demonstrate the GPS is accurate to within 11 ± 2.4 metres and has shown to maintain this accuracy over a period of 6 hours. Battery life has been found to last 5 hours and 47 ± 28 minutes. From a cold start the device is able to *lock* onto satellites and begin recording data in 50 ± 12 seconds. The devices processing speed has also been gauged and has been found to take just 0.29 seconds to acquire data from GPS, apply formatting corrections, sample the ADCs and log data to the memory card theoretically providing detection of stroke rates up to 150 strokes per minute.

Cadence (strokes per minute) detection has been implemented into the device and through real-world testing has been demonstrated to have a detection accuracy of $\sim 70\%$ over a test period of 10 minutes (an example of 8 recorded strokes is shown in Figure 5). Although this is not ideal, the device possess the ability to improve through the use of the embedded auto-calibration utility. Due to time constraints real-world testing has only been performed once. During testing the device was exposed to salt water as the case was open to adjust settings. This produced an intermittent fault in the device which prohibited extensive testing. However, in analysing the results recorded, it was possible to apply the auto calibration utility to the data which provided slightly better stroke detection, 75% correctly identified.

8. Conclusion

The completed device's battery life exceeds that of the longest usual training sessions by 90 minutes.

The battery is easily accessible and replaceable if it begins to degrade in the future, and may be recharged by mini USB cable, removing the need for an application specific charger.

WaterLogged contains a number of features surplus to project requirements and objectives. These have been included to improve the device functionality and usability. The ability to save/load device settings from the memory card has been implemented so users are not confined to default, programmed values when they start the device. Data formatting has been provided which enables logged data to be uploaded directly to *GPSVisualizer.com*, without any pre-processing. This has been implemented so logged positional information can be overlaid onto Google Maps giving the device the potential to become a *trip* logger for a wide variety of sports.

9. References

- [1] Waka Ama (Nga Kaihoe O Aotearoa INC), last accessed September 26th 2009
<http://www.wakaama.co.nz/>
- [2] Westerer P, K. R. *Physical activity assessment with accelerometers*, 1999, Department of Human Biology, Maastricht University
- [3] Thomas Borkin, *Microcontroller Assistance for Waka Ama*, 2007, Victoria University of Wellington, New Zealand
- [4] Payne. A. *Design and Construction of a Pair of Cooperating Autonomous Mobile Robots*, 2004, Waikato University, New Zealand
- [5] National Marine Electronics Association 0183 (NMEA), last accessed September 10th 2009
<http://www.nmea.org/>
- [6] Introduction to the Global Positioning System for GIS and TRAVERSE, last accessed October 8th 2009,
<http://www.cmtinc.com/gpsbook/index.htm>
- [7] GPS Visualizer: Do-It-Yourself Mapping, last accessed September 29th 2009
<http://www.gpsvisualizer.com>

Comparison of Optimized Low-Power LNA topologies for 866 MHz UHF RFID

¹Jie Li and ²S. M. Rezaul Hasan

Centre for Research in Analogue and VLSI Microsystems Design (CRAVE)
School of Engineering and Advanced Technology (SEAT)
Massey University, Albany Highway, Gate 4, Building 106, Albany, Auckland

Email: ¹J.Li.2@massey.ac.nz ²hasanmic@massey.ac.nz

Abstract: The design topologies of optimized CMOS low noise pre-amplifier (LNA) for UHF RFID reader at 866 MHz is investigated. Simulations and analysis based on inductively degenerated common-source (CS) single-end telescopic LNA, folded-cascode LNA and differential LNA are performed in detail. The significant feature of the LNA design is the simultaneous impedance and minimum F_{min} noise matching at a low UHF radio frequency, achieved at a very low power drain of 0.95 mW from a 0.7 V supply voltage for single-ended, 1.33 mW from a 0.5 V supply voltage for folded-cascode and 2 mW from a 0.7 V supply voltage for differential topology. Design using sub-1V supply voltage is quite challenging owing to inductor size and bias drain related noise factor degradation. The LNA was simulated in IBM 130nm 6-metal 1 poly RFMOS process technology to deliver a power gain (S_{21}) of over 16 dB, a reverse isolation of -40dB and an input power reflection of -16dB for all design topologies at 866 MHz. It had a minimum pass-band NF of around 1.2 dB (for all topologies) and a 3rd order input referred intercept point (IIP3) of -11.5 dBm for single-ended telescopic topology and -9.5 dBm for differential topology.

Keywords: CMOS LNA, RFID, impedance matching, low power consumption, noise matching, narrow-band.

1 INTRODUCTION

RFID (radio-frequency identification) is one of fastest growing wireless communication technology nowadays for commercial products tracking. A RFID system usually consists of a transponder tag and a reader as shown in Fig. 1. The low noise amplifier (LNA) is the first block in the front-end of the RFID reader that is tuned at a certain trans-receiver frequency. A longer distance can be achieved under UHF band at low power dissipation. However, the LNA need to be designed optimally to minimize the noise for following stages and avoid the distortion of the source signal (requires good linearity).

Considerable research on CMOS LNA design in sub-micron technologies has been approached in recent years: from topology investigation[1, 2] and guideline[3] to various novel ideas on the design improvement of multiple band[4], low-noise figure[5-7], high power gain[8], low power consumption[7, 9-11], and high linearity. The LNA designs at 900 MHz also have been reported by authors in [2, 4, 5, 7, 12]. A lower frequency standard operated at 866 MHz for UHF RFID is also implemented in Europe, Africa and New Zealand. Narrow-band CMOS LNA design at such sub-GHz frequencies has not been widely reported so far. Low power dissipation is a

significant design criterion for RFID applications which will be synthesized by finding the trade-off between gain, low noise figure, input and output impedance matching and high linearity.

In this paper we discuss the complete design and comparison results for three integrated 130 nm CMOS 866 MHz LNAs. One is a single-ended common-source telescopic cascode LNA using an enhanced PCSNIM technique from [2], a folded-cascode LNA operated at 0.5V supply voltage and a differential LNA.




Fig.1 The Simulation Environment of Return Link for RFID system (Receiver, Wireless channel, and Tag)

2 OPTIMIZED LOW-POWER UHF RFID LNA TOPOLOGIES

2.1 Single-Ended telescopic cascode LNA

A single-ended 866 MHz LNA is shown in Fig. 2. This amplifier has the commonly used cascoded architecture form, which provides a good isolation between the input and output stages. The input inductive degeneration topology formed by M_1 and L_s is employed here to get better noise performance for the narrow band applications. L_g is an external inductor for input reactance matching with L_s and C_e . The output impedance matching can be obtained by tuning the L_d , C_d , and R_d at 866 MHz with an angular frequency bandwidth of $(R_d * C_d)^{-1}$. Typically, both the input and the output impedances are required to match to R_s (50 Ω). M_3 in conjunction with R_{ref} controls the DC bias current through telescopic cascode while a large R_b ensures that noise and any ac signal pick-up at the gate of the current mirror is blocked and decoupled. C_{dc} is the DC blocking capacitance at both terminals.




Fig.2 The Completed telescopic cascode LNA topology with DC bias voltage

Power constrained simultaneous noise and impedance matching technique (PCSNIM) is used as the basic design methodology. A smaller device size of M_1 is chosen for minimizing power dissipation. The extra capacitor

C_e is in parallel with C_{gs1} of M_1 in order to achieve minimum noise figure F_{min} with power constrain and higher ω_T [2]. Optimal impedance Z_{opt} for noise match can be derived to achieve theoretical F_{min} $\left(\cong 1 + \frac{2}{\sqrt{5}} \frac{\omega C_{gs1}}{g_{m1}} \sqrt{\delta \gamma (1 - |c|^2)} \right)$ [2]. With regard to the input impedance matching, we can easily derive:

$$Z_{in} \cong R_t + R_g + s(L_s + L_g) + \frac{1}{sC_t} + \frac{g_{m1}}{C_t} \frac{L_s g_{m2}}{(g_{ds1} + g_{m2})}$$

$$\cong \frac{g_{m1}}{C_t} \frac{L_s g_{m2}}{(g_{ds1} + g_{m2})} \quad (\text{at resonance})$$

where R_t is the series resistance of the gate inductor, and R_g is the gate resistance of the transistor M_1 . R_t can be neglected for high Q inductors and R_g can be reduced to insignificant levels by inter-digitating the device [1]. C_t here equals to $C_{gs1} + C_e$ and the effect of finite device conductance $g_{ds} = (1/r_0)$ is included at deeply scaled channel length.

At the central frequency of 866 MHz, the imaginary term of Z_{in} will be zero to match Z_s , which gives

$$s(L_s + L_g) + \frac{1}{sC_t} = 0$$

Optimum values can be calculated from above equations to satisfy PCSNIM technology. The noise figure NF is defined as

$$F = F_{min} + \left[\frac{\gamma}{\alpha g_{m1} R_s} \right] \left[1 - \frac{Q_{opt}}{Q_s} \right]^2$$

The typical value of α is about 0.85, the optimum source quality factor (conductance) Q_{opt} and the actual source quality factor Q_s are defined as

$$Q_{opt} = \alpha \sqrt{\frac{\delta}{5\gamma} (1 - |c|^2)}$$

and

$$Q_s = \frac{1}{\omega C_t R_s}$$

The expression for the width of the optimum device under power-constrained minimum noise figure is given by

$$W_{opt} = \frac{3}{2} \frac{1}{\omega L C_{ox} R_s Q_{sp}} \approx \frac{1}{3\omega L C_{ox} R_s}$$

The calculated width might give an extraordinary large value for short-channel CMOS technology (130nm) under low UHF frequency range. Hence, smaller value is adopted in order to satisfy the power-constrained technology practically. Furthermore, the impact of this input matching technique on linearity can also deteriorate under short-channel CMOS. A larger MOSEFT or a larger capacitor C_e can enhance the linearity by reducing L_g and L_s at the same time, but the NF may deteriorate.

2.2 Folded-Cascode LNA

The significant criteria for RFID application is not just low power consumption but low supply voltage as well. A folded-cascode topology can be adopted under ultra low supply voltage condition.

The basic circuit schematic of folded-cascode LNA consists of one NMOS transistor and one PMOS transistor as shown in Fig. 3. The voltage supply can be significantly reduced compared to the above telescopic cascode circuit. Similar to telescopic cascode topology, M_1 is the input transistor to provide the transconductance (g_m) of the LNA and M_2 is used as the current buffer to minimize the Miller effect and improve the response at high frequencies. However, one inherent feature of the folded cascode is that it allows simple gain control. This can be used to improve the overall linearity of the receiver. This is achieved by varying the gate voltage of M_2 to affect the impedance looking from the source (Z_p), hence adjusting the overall gain of the LNA. The overall voltage gain (A_{total}) of the tuneable LNA can be represented to be

$$A_{total} = A_v G_{tune} \cong g_{m1} R_o G_{tune}$$

where A_v is the fixed voltage gain of a conventional cascode LNA, R_o is the impedance of the output RLC resonant tank, and G_{tune} is the gain tuning factor that represents the portion of the AC signal current generated by input transistor M_1 , which flows into the source of M_2 . It is given by,

$$G_{tune} = \frac{Z_{blk}}{Z_{blk} + Z_p} = \frac{Z_{blk}(1 + g_{m2}r_{o2})}{Z_{blk}(1 + g_{m2}r_{o2}) + r_{o2}}$$

where $Z_{blk} = j\omega L_{blk}$ is given a large value to represent “blocking impedance” and Z_p is

$$Z_p = \frac{1}{g_{m2}} \parallel r_{o2} \cong \frac{1}{g_{m2}}$$

where g_{m2} and r_{o2} are the transconductance and the output resistance of M_2 respectively. It is evident that the A_{total} can be adjusted by varying g_{m2} , which is controlled by V_{bias} . Another advantage is that this gain control technology does not affect input matching.




Fig.3 The folded-cascode LNA topology

2.3 Differential LNA

The single-ended and folded cascode architectures have the shortcoming of being sensitive to parasitic ground inductance. Differential architectures, on the other hand, will be somewhat immune to the common-mode interference from substrate or supply perturbations[1]. Therefore the differential LNA is also presented here in Fig. 4. The degenerating inductors (L_s) are connected together at the “virtual ground”. At this point, a current source is usually placed to provide a current that is twice the current flowing down from each side of the LNA section. Any parasitic reactance in series with the bias current source is largely irrelevant. Two “ideal” balun (balanced to unbalanced) transformers are used here to supply a transformation between differential and single-end signals for both input and output. Two AC sources can also be applied with opposite polarity here to achieve the same effect. Again, they will both be matched to 50 Ω at 866 MHz as in the singled-ended LNA design.

Another advantage of applying differential topology is its ability to reject common-mode disturbance. In mixed-

signal applications, both the supply and substrate voltages may be noisy. Hence, this attribute becomes particularly important.

However, for equal noise figure, the power consumption of this amplifier is twice that of its single-ended counterpart. Offsetting this disadvantage is the improved linearity that is attained by dividing the input voltage between two devices.




Fig. 4 The differential LNA topology


3 SIMULATION RESULTS

The results for the proposed 866 MHz UHF RFID LNA simulated using Cadence Spectre, Virtuoso and Assura are now discussed. The results of the single-ended telescopic LNA are shown in Fig. 5. With a supply of 0.7 V, at 866 MHz, the input reflection coefficient S_{11} = -15 dB, forward power gains S_{21} is around 16 dB, and noise figure is 1.2 dB. The power dissipation is only around 0.95 mW. The supply voltage can go down further to reduce power dissipation, but with a smaller gain, on the other and, a higher gain also results in more power consumption. Fig. 5 (a), (b) and (c) also demonstrate the S_{11} , S_{21} and noise figure behaviour under both 0.7 V and 1 V voltage supplies. Due to the PCSNIM technique, noise figure is minimized at resonance frequency at around 1.1 dB. This example proves how to optimize the trade-off between all those specifications in LNA designs. The IIP3 and the 1-dB compression point were determined using a two-tone test at 866 MHz with 1 MHz tone separation. The results indicate an IIP3 of -11.5 dBm and a 1dB compression point of -16.1dBm as shown in fig. 6.


Fig. 7 shows the S_{11} , S_{21} and noise figure simulated results for folded cascode architecture as well as their different behaviours under different PMOS bias voltage from 0V to 0.1V. With this voltage control technique, most s-parameters maintain the similar results compared with single-end cascode topology. However, with 0.5 V supply voltage, the folded cascode LNA consumes 1.33 mW, which is slightly higher than single-ended architecture.

In order to maintain low noise figure as in single-ended telescopic cascode topology, the differential topology is duplicated symmetrically by two single-ended telescopic cascode LNA to give the same noise figure of 1.2 dB as shown in Fig.8 (b). The power consumption is also doubled but we can identify the improvement of 1 dB compression and IIP3 values of -9.1 dBm and -9.4 dBm compared with single-ended case with values of -16.1 dBm and -11.5 dBm respectively.


Completed specifications of all the proposed LNA designs are summarized in Table I. Table I compares the performance of the proposed UHF LNAs with other UHF LNA designs indicating very low NF achieved at sub-mW power dissipation.



(a)



(b)



(c)

Fig.5 Simulated S- parameter results of the single-ended LNA in Fig.2 with two different supply voltages: 0.7V and 1V (a) S_{11} (b) S_{21} and (c) NF




Fig.6 Simulated IIP3 and 1dB compression point results of the singled-end LNA in Fig.2









Fig.7 Simulated S-parameter results of the folded cascode LNA in Fig.3 with different PMOS bias voltage: 0V, 0.005V and 0.1V (a) S_{11} (b) S_{21} and (c) NF



(a)



(b)



(c)

Fig.8 Simulated S-parameter, noise figure, IIP3 and 1dB compression point results of the differential LNA in Fig.4 (a) S_{11} and S_{21} (b) NF and (c) IIP3 and 1dB compression point

4 CONCLUSIONS

The optimized UHF RFID LNA at 866 MHz has been designed and simulated results demonstrated. The design achieves very low NF of 1.2 dB with only power dissipation of 0.95 mW for singled-end, 1.33 mW for folded cascode and 2 mW for differential topologies respectively.

TABLE I

SUMMARY OF THE PERFORMANCE OF PRESENTED UHF RFID LNAs AND COMPARISON WITH PREVIOUS UHF LNA DESIGNS

	This Work			[2]	[4]	[5]	[6]	[7]	[9]	[10]	[12]	[13]
Process Technology (μm)	0.13			0.25	0.18	0.35	0.24	0.35	0.18	0.18	0.35	0.18
Topology	┆	π	\parallel	π	┆	┆	┆	┆	┆	┆	\parallel	\parallel
S11 (dB)	-16	-36	-21	-18	-12	-10	-38	-10	N/A	-29	-11	N/A
S21 (dB)	16.4	16.3	17.8	12	14	17.5	8.8	14	15.85	12.5	13.4	15
S12 (dB)	-39	-41	-42	N/A	N/A	N/A	N/A	-34	-51.3	-60	N/A	N/A
NF (dB)	1.20	1.18	1.19	1.35	2.3	2	1	1.05	0.9	0.7	3.2	2.9
IIP3 (dBm)	-11.5		-9.4	-4	-14	-6	7.1	0	-5.01	-4	10.8	N/A
P1dB (dBm)	-16.1		-9.1	-15	N/A	N/A	N/A	-12	-19.5	-9	1.4	-15
Power Dissipation (mW)	0.95	1.33	2.0	2	7.5	21.6	7.5	9	0.5	3.9	33	4.32
Supply Voltage (V)	0.7	0.5	0.7	1.25	1.8	2.7	2	1.8		1.8	3.0	1.8
Centre Frequency (MHz)	866			900	900	900	800	920	404	915	900	900
N/A: not available ┆: single-ended cascode topology π : folded-cascode topology \parallel : common-gate topology \parallel : differential topology												

5 REFERENCES

[1] T. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge: Cambridge University Press, 2004.

[2] N. Trung-Kien, K. Chung-Hwan, I. Gook-Ju, A. M.-S. Y. Moon-Su Yang, and A. S.-G. L. Sang-Gug Lee, "CMOS low-noise amplifier design optimization techniques," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 52, pp. 1433-1442, 2004.

[3] M. Yehia, A. Nieuwoudt, and R. Tamer, "Automated Design Solutions for Fully Integrated Narrow-Band Low Noise Amplifiers," in *System-on-Chip for Real-Time Applications, The 6th International Workshop on*, 2006, pp. 109-114.

[4] D. Vu Kien, B. Quang Diep, and P. Chul Soon, "A Multi-band 900MHz/1.8GHz/5.2GHz LNA for Reconfigurable Radio," in *Radio Frequency Integrated Circuits (RFIC) Symposium, 2007 IEEE*, 2007, pp. 69-72.

[5] F. Gatta, E. Sacchi, F. Svelto, P. A. V. P. Vilmercati, and R. A. C. R. Castello, "A 2-dB noise figure 900-MHz differential CMOS LNA," *Solid-State Circuits, IEEE Journal of*, vol. 36, pp. 1444-1452, 2001.

[6] G. Jung-Suk, A. Hee-Tae, D. J. Ladwig, Y. Zhiping, T. H. Lee, and R. W. Dutton, "A noise optimization technique for integrated low-noise amplifiers," *Solid-State Circuits, IEEE Journal of*, vol. 37, pp. 994-1002, 2002.

[7] G. Gramegna, A. Magazzo, C. Sclafani, M. A. P. M. Paparo, and P. A. E. P. Erratico, "A 9mW, 900-MHz CMOS LNA with 1.05dB-noise-figure," in *Solid-State Circuits Conference, 2000. ESSCIRC '00. Proceedings of the 26th European*, 2000, pp. 73-76.

[8] S. Asgaran and M. J. Deen, "A novel gain boosting technique for design of low power narrow-band RFCMOS LNAs," in *Circuits and Systems, 2004. NEWCAS 2004. The 2nd Annual IEEE Northeast Workshop on*, 2004, pp. 293-296.

[9] H. S. Savci, Z. Wang, A. Sula, N. S. Dogan, and E. Arvas, "A 1-V UHF low noise amplifier for ultralow-power applications," in *Circuits and Systems, 2006. ISCAS 2006. Proceedings. 2006 IEEE International Symposium on*, 2006, p. 4 pp.

[10] X. Wu, L. Sun, and Z. Wang, "Low-Power 915MHz CMOS LNA Design Optimization Techniques for RFID," in *Microwave and Millimeter Wave Technology, 2007. ICMMT '07. International Conference on*, 2007, pp. 1-4.

[11] M. M. E. Kholy, "Ultra Low Power Narrow Band LNA," in *International Design and Test Workshop, 2007 2nd*, 2007, pp. 38-42.

[12] R. Hyoung-Hwan, P. Kyoung-Tae, O. Ha-Ryong, S. Yeung-Rak, P. Jun-Seok, and K. Min-Soo, "A common gate low noise amplifier with high linearity over UHF RFID bands," in *Electromagnetic Compatibility and 19th International Zurich Symposium on Electromagnetic Compatibility, 2008. APEMC 2008. Asia-Pacific Symposium on*, 2008, pp. 88-91.

[13] K. Ickjin, E. Yunseong, B. Heemun, C. Kyudon, J. Sangyoon, J. Sungjae, L. Donghyun, and L. Heungbae, "A Single-Chip CMOS Transceiver for UHF Mobile RFID Reader," *Solid-State Circuits, IEEE Journal of*, vol. 43, pp. 729-738, 2008.

A Digital CMOS Sequential Circuit Model for Bio-Cellular Adaptive Immune Response Pathway Using Phagolysosomal Digestion: A Digital Phagocytosis Engine

S. M. Rezaul Hasan

Center for Research in Analog and VLSI microsystem dEsign (CRAVE), School of Engineering and Advanced Technology (SEAT), Massey University, Auckland, New Zealand
hasanmic@massey.ac.nz

Abstract: Living systems have to constantly counter micro-organisms which seek parasitic existence by extracting nutrition (amino acids) from the host. Phagocytosis is the ingestion of micro-creatures by certain cells of living systems for counter nutrition (breakdown of the micro-creature into basic components) as part of cellular adaptive immune response. These particular cells are called phagocytes, all of which are different types of white blood cells or their derivatives. Phagocytes are activated by certain components of the micro-creatures which act as an antigen, generating antibody secretion by the phagocyte. This paper develops a digital CMOS circuit model of phagocytosis: the immune response biochemical pathway of a phagocyte. A micro-sequenced model has been developed where the different stages in phagocytosis are modeled as different states clocked by circadian time intervals. The model converts the bio-chemical immune system digestive pathway into a cascade of CMOS multi-step logical transformations from micro-creature ingestion to the secretion of indigestible residuals. This modeling technique leads to the understanding of cellular immune deficiency diseases of living systems in the form of logical (electrical) faults in a circuit.

Index Terms - Systems biology, CMOS circuit, phagocytosis, silicon mimetic.

I. Introduction

Looking at bio-chemical pathways of biological systems from a micro-electronic circuit diagram perspective is an interesting analytical paradigm in developing the convergence of biosciences, biotechnology and electrical engineering. In an effort towards understanding the bio-chemical functions of living systems at molecular interaction level researchers working in the developing field of systems biology [1][2] have endeavored to provide models that integrate the molecular interactions between

proteins[3], enzymes[15][16], DNA[4], RNA[15], ions etc. within individual bio-cells of different kinds and their derivatives.

Immune response is one of the most fundamental bio-chemical activity of living systems. Micro-creatures constantly harboring towards living systems for parasitic existence are subjected to counter nutrition mechanisms of certain type of immune system cells, e.g., white blood cells or their derivatives usually known as phagocytes. This process of immune response is known as phagocytosis. It is a cascade of molecular interactions within the phagocyte which result in the breakdown of the micro-creature into its basic components, parts of which are recycled (digested) for nutrition and the rest eliminated from the system as residual waste. Circuit theoretic techniques of modeling this biochemical immune system pathway within a phagocyte compared to other mathematical modeling techniques [5][6] is the central theme of this paper. In recently published papers [7][8][9] this author described how the states of DNA protein interactions within the bio-cellular operations can be modeled using integrated circuits (*silicon mimetic* model [10][11]). In addition, models for mitochondrial respiration related electron transfer pathway [12], and cell signaling pathway using G-Protein and Phosphorylation cascade [17] has been reported by this author. These models based on integrated circuits are amenable to the large variety of mature simulation tools [13] available in the Very Large Scale Integration (VLSI) Computer aided design (CAD) arena. This type of modeling thus enhances the knowledge required for the realization of the “dream” of designed biochemical pathways and gene circuits [10] with revolutionary outcomes for gene and cell therapy (“*nano-medicine*”). In this work a digital CMOS[14] micro-sequenced model for phagocytosis has been developed where the sequence of biochemical processing stages are clocked by circadian time (hours) [5] intervals.

II. Phagolysosomal Immune Response Bio-chemical Pathway

Infection of living systems by micro-creatures results in various types of white blood cells becoming phagocytes and moving into infected tissue. These phagocytes enlarge and develop into a macrophage engulfing and devouring micro-creatures in a counter nutrition effort. Fig. 1(a) shows the microphotograph of such a macrophage digesting micro-creatures (red bacterium). Fig. 1(b) shows the mechanism of phagocytosis in terms of the bio-chemical pathway in a phagocyte. The bio-chemical pathway of Phagocytosis [15][16] can be divided into several states such as chemotaxis (or phototaxis), adherence, opsonization, ingestion, digestion and residue ejection. Chemotaxis (or, phototaxis) by chemical (or optical) stimuli is the mechanism of approach through “run” and “tumble” of micro-creatures towards favorable host living systems. Phagocytes are attracted by the micro-creatures through their antigenic components, resulting in the release of antibody serum proteins that opsonize the microcreatures. Opsonized microcreatures then easily adhere to the plasma membrane of the phagocyte which is referred to as a state of adherence. After adherence, there is ingestion, whereby, the phagocyte extend its flagella type “arms” called pseudopods which surround and then engulf the microcreature. Once ingested, the surrounding pseudopods fuse and enclose the micro-creature in a fluid sack called “phagosome” or phagocytic vesicle. The membrane of this vesicle has enzymes which pump hydrogen ions (H^+) into the vesicle, thereby reducing the pH to around 4, so that hydrolytic enzymes can be activated for the breakdown of the micro-creature. The next stage of phagocytosis is the digestive pathway which begins with the phagosome “pinching off” from the plasma membrane and entering the cytoplasm. Inside the cytoplasm the phagosomes attach with lysosome sacks containing digestive enzymes and microbicidal substances. On attachment, the membranes of lysosome and phagosome merge forming a larger composite sack called phagolysosome. The degenerative components inside the phagolysosome take only sub-circadian time interval (less than an hour) to completely breakdown (“kill”) micro-creatures inside it. This digestive pathway contains several main and sub-processes. Lysosomic enzyme lysozyme directly attacks the cell-wall of the micro-creature and initiates the hydrolysis of the cell-wall. A host of other lysosomic enzymes are also active at the same time, so that the micro-creatures macro-components are disintegrated, such as its lipids by the enzyme lipase, its proteins by protease, its ribonucleic acids by ribonuclease, and, its deoxyribonucleic acids by deoxyribonuclease. Lysosomes also contain certain enzymes which can initiate a process known as oxidative burst which result in toxic oxygen products

such as super oxide radical (O_2^-), hydrogen per oxide (H_2O_2), hydroxyl radical ($OH\bullet$) and singlet oxygen ($^1O_2^-$). Other lysosomic enzymes combine with these toxic oxygen products to produce secondary microbicides which hydrolyze and breakdown the micro-creatures. For instance, the enzyme myeloperoxidase (MY) converts chlorine ions (Cl^-) and hydrogen per oxide (H_2O_2) into highly toxic hypochlorous acid (HOCl) which contain the well-known anti-microbial agent hypochlorous ions.

III. A Micro-sequenced Model of Phagocytic Biochemical Pathway

A *silicon mimetic* [9] model of the bio-chemical pathway of phagocytosis has been developed and is shown in Fig. 2. The discussion in the previous section indicates the existence of many intermediate steps in the bio-chemical digestive pathway of phagocytosis whose proper co-ordination is an important factor in the functional phagocyte of living systems. Similar to the concept introduced in [17], the intermediate steps in the bio-chemical path of phagocytic digestion of micro-creatures can be modeled as states in Delay flip-flops (Registers) which indicate the states of molecular interactions (binding, hydrolysis and breakdown) that constitute the behavior of the phagocyte. Using this modeling technique malfunction in the immune system (immune deficiency syndrome: *cause of numerous diseases*) can be conveniently modeled as circuit faults, such as, outputs of logic gates or flip-flops (registers) are stuck at logic “0” or at logic “1”, thus unifying the approach to solving faults in electronic or biological circuits (in this case the biological circuit of the immune response system). At any time-instant n , presence of Gram positive (GP) or Gram negative (GN) micro-creatures in close proximity to the phagocyte is stored as $u1(n)$ in the D-FF_1. The presence of chemo-tectile micro-creatures due to the presence of chemical signals is stored as $u2(n)$ in the D-FF_2. On the other hand, the presence of photo-tectile micro-creature due to the presence of optical signals is stored as $u3(n)$ in the D-FF_3. The presence of chemo-tectile or photo-tectile micro-creature opsonized by anti-body serum protein (SP) is stored as $u4(n)$ in the D-FF_4. The presence of ingested micro-creatures engulfed by pseudopods (PP) is stored as $u5(n)$ in the D-FF_5. The state of the ingested micro-creatures within phagosomes with H^+ ions pumped inside the phagocytic vesicle is indicated by $u6(n)$ in the D-FF_6. The “pinching off” of phagosome from the plasma membrane and entering

the cytoplasm is registered as $u7(n)$ in D-FF_7. The attachment and merger of lysosome with the phagosome in the phagolysosome, and the resulting state of the digestive degradation of the microcreatures is represented by,

$$u14(n) = \sum_{m=10}^{13} um(n-1) \quad (1)$$

The hydrolysis of micro-creature inside the phagolysosome due to hypochloric acid (a byproduct of the oxidative burst) is indicated by the state $u8(n)$ stored in the D-FF_8. The breakdown of the micro-creature's cell-wall by the enzyme lysozyme (LYE) is indicated by the state $u9(n)$ in the D-FF_9. The breakdown of lipids inside the degenerated microcreatures by the lysosomic enzyme lipase is indicated by the state $u10(n)$ stored in D-FF_10. Next, the breakdown of the bacterial proteins inside the degenerated microbe by the lysosomic enzyme protease is indicated by the state $u11(n)$ stored in D-FF_11. The breakdown of the bacterial ribonucleic acids inside the degenerated microbe by the lysosomic enzyme ribonuclease (RNE) is indicated by the state $u12(n)$ stored in D-FF_12. And, finally the breakdown of the bacterial DNA by the lysosomic enzyme deoxyribonuclease (DNE) is indicated by the state $u13(n)$ stored in D-FF_13. Based on the micro-step sequenced model of Fig. 2, the following discrete-time state equations can be written for the phagocytic digestive pathway:

$$u1(n) = GP + GN \quad (2)$$

$$u2(n) = CS * u1(n-1) \quad (3)$$

$$u3(n) = OS * u1(n-1) \quad (4)$$

$$u4(n) = [u3(n-1) + u2(n-1)] * SP \quad (5)$$

$$u5(n) = u4(n-1) * PP \quad (6)$$

$$u6(n) = u5(n-1) * H^+ \quad (7)$$

$$u7(n) = u6(n-1) \quad (8)$$

$$u8(n) = [u7(n-1) * \{MY * Cl^- * (LE1 * H_2O)\}] \quad (9)$$

$$u9(n) = [u7(n-1) * LYE] \quad (10)$$

$$u10(n) = [u8(n-1) + u9(n-1)] * LPE \quad (11)$$

$$u11(n) = [u8(n-1) + u9(n-1)] * PTE \quad (12)$$

$$u12(n) = [u8(n-1) + u9(n-1)] * RNE \quad (13)$$

$$u13(n) = [u8(n-1) + u9(n-1)] * DNE \quad (14)$$

$$u14(n) = [u10(n-1) + u11(n-1) + u12(n-1) + u13(n-1)] \quad (15)$$

Modeling immune response deficiency as electrical faults: Using this circuit model immune deficiency

diseases can be modeled as electrical faults. For example, if the serum protein (anti-body) is not produced, the invading micro-creatures cannot be engulfed, and, the state at $u(4)$ will be stuck at "0", resulting in all the down-stream digestive degradation pathway states being stuck at "0", i.e., becomes non-functional. The solutions can then be designed in terms of CMOS circuit functionality and converted into the equivalent biochemical solution ("nano-medicine").

IV. Micro-sequenced Phagocyte Pathway States and Silicon Mimetic Signal Transduction Results

Fig. 3 displays how the immune system response bio-chemical events form a pipeline through the *silicon mimetic* model of phagocytosis in a digital CMOS micro-sequenced logic simulation process. It also shows the sequence of intermediate states (D flip-flop states) corresponding to whether a certain biochemical state is activated. The sequence of molecular interactions are considered to be taking place in the circadian time-scale (hours or few minutes). The D flip-flops are thus driven by a clock with circadian time period. Fig. 3 shows molecular events occurring over 24 circadian time periods during which microcreatures (bacteria or microbes) arrive towards phagocytes until the period $n=8$. This causes the immune response pathway to activate. This results in the presence of either chemo-tectile (D-FF_2) or photo-tectile (D-FF_3) micro-creatures during the circadian interval $\langle 2,9 \rangle$. Also, opsonized microcreatures are present during the interval $\langle 3,10 \rangle$, microcreatures are ingested during the interval $\langle 4,11 \rangle$, digestive vesicles are formed near the phagocyte membrane during the interval $\langle 5,12 \rangle$, vesicles are detached and moves into the cytoplasm during the interval $\langle 6,13 \rangle$, phagolysosome is formed and the microcreature is hydrolyzed (state in D-FF_8) and microcreature cell-wall is hydrolyzed (state in D-FF_9) during the interval $\langle 7,14 \rangle$. The breakdown of lipids (state in D-FF_10), proteins (state in D-FF_11), ribonucleic acid (state in D-FF-12) and deoxyribonucleic acid (state in D-FF_13) take place simultaneously in the interval $\langle 8,15 \rangle$. The undigested micro-creature artifacts (state in D-FF_14) are removed during the interval $\langle 9,16 \rangle$. As the microcreatures are is not present during the interval $\langle 9,12 \rangle$, there is a short pause in the digestive bio-chemical process which is evident from the diagonal zero states in Fig. 3(b). Also, beginning the interval $\langle 13,16 \rangle$ the presence of a short burst of microcreatures results in a short pipeline of digestive pathway activities (reactions and interactions) that is evident through the diagonal array of "1"s as shown in Fig. 3(b).




Fig. 2. A micro-sequenced digital CMOS model of the bio-chemical pathway of phagocytosis using logic gates and D-flip-flops.

A micro-step sequence for phagocytosis

Path-way state	Circadian Time-Instant											
	n=0	n=1	n=2	n=3	n=4	n=5	n=6	n=7	n=8	n=9	n=10	n=11
D-FF_1	RESET	microb	0	0	0							
D-FF_2	RESET	0	1	1	1	1	1	1	1	1	0	0
D-FF_3	RESET	0	1	1	1	1	1	1	1	1	0	0
D-FF_4	RESET	0	0	1	1	1	1	1	1	1	1	0
D-FF_5	RESET	0	0	0	1	1	1	1	1	1	1	1
D-FF_6	RESET	0	0	0	0	1	1	1	1	1	1	1
D-FF_7	RESET	0	0	0	0	0	1	1	1	1	1	1
D-FF_8	RESET	0	0	0	0	0	0	1	1	1	1	1
D-FF_9	RESET	0	0	0	0	0	0	1	1	1	1	1
D-FF_10	RESET	0	0	0	0	0	0	0	1	1	1	1
D-FF_11	RESET	0	0	0	0	0	0	0	1	1	1	1
D-FF_12	RESET	0	0	0	0	0	0	0	1	1	1	1
D-FF_13	RESET	0	0	0	0	0	0	0	1	1	1	1
D-FF_14	RESET	0	0	0	0	0	0	0	0	1	1	1

(a)

A micro-step sequence for phagocytosis (contd.)

Path-way state	Circadian Time-Instant											
	n=12	n=13	n=14	n=15	n=16	n=17	n=18	n=19	n=20	n=21	n=22	n=23
D-FF_1	0	microb	microb	microb	microb	0	0	0	0	0	0	0
D-FF_2	0	0	1	1	1	1	0	0	0	0	0	0
D-FF_3	0	0	1	1	1	1	0	0	0	0	0	0
D-FF_4	0	0	0	1	1	1	1	0	0	0	0	0
D-FF_5	0	0	0	0	1	1	1	1	0	0	0	0
D-FF_6	1	0	0	0	0	1	1	1	1	0	0	0
D-FF_7	1	1	0	0	0	0	1	1	1	1	0	0
D-FF_8	1	1	1	0	0	0	0	1	1	1	1	0
D-FF_9	1	1	1	0	0	0	0	1	1	1	1	0
D-FF_10	1	1	1	1	0	0	0	0	1	1	1	1
D-FF_11	1	1	1	1	0	0	0	0	1	1	1	1
D-FF_12	1	1	1	1	0	0	0	0	1	1	1	1
D-FF_13	1	1	1	1	0	0	0	0	1	1	1	1
D-FF_14	1	1	1	1	1	0	0	0	0	1	1	1

(b)

Fig.3. Sequence of intermediate steps (D flip-flop states) in the phagocytic bio-chemical pathway, (a) for circadian time-instants n=0 to n=11, and, (b) for circadian time-instants n=12 to n=23.

Channel Sounding with Software Defined Radio

Harry W. H. Jones, Pawel A. Dmochowski, Paul D. Teal

*School of Engineering and Computer Science, Victoria University of Wellington
NEW ZEALAND.

Email: {harry.jones, pawel.dmochowski, paul.teal}@ecs.vuw.ac.nz

Abstract:

We present an investigation into the channel sounding of wireless communications systems using a Software Defined Radio (SDR) platform. Using the Universal Software Defined Radio (USRP), in conjunction with the GNU Radio SDR system, we implemented and evaluated channel sounding capabilities using a variety of signal sequences. The final implementation utilized the IEEE 802.11a Orthogonal Frequency Division Multiplexing (OFDM) preamble for the tasks of carrier frequency and timing estimation, as well as channel sounding. We show that the system is able to characterize 1.8 MHz of spectrum reliably, with a resolution of 34 kHz. The bandwidth restrictions caused by the hardware bottlenecks were the main source of limitation.

Keywords:

Channel sounding, Software defined radio, GNU Radio, USRP

1 INTRODUCTION

In all mobile communication systems, it is essential to take into account the rapidly changing wireless environment. An estimate of the channel impulse response (CIR) or transfer function (CTF) can be used to equalise its effects at the receiver and to optimize transmission. The estimation of these characteristics is referred to as *channel sounding*. Providing accurate channel state information (CSI) to the transmitter can further enhance the use of the channel, although some form of prediction may be required to compensate for the latency in providing this information. Multiple-Input Multiple-Output (MIMO) systems particularly rely on reliable CSI, and most MIMO schemes require CSI at the transmitter.

Modern communications systems are required to be highly adaptive, due to the changing conditions of the wireless environment. Parameters within a system that may be adapted include modulation, power allocation and coding technique. Many functionalities, such as modulation, pulse shaping, constellation mapping and coding are typically done in hardware. Such architectures lack the ability to adapt during operation, and also complicate the task of system re-design. Software Defined Radio (SDR) offers a solution to a number of the issues described above by abstracting the tasks of constellation mapping, coding and pulse shaping from a hardware-based issue to a software one. This allows the engineer to change the specifications of a system by simply re-writing code, not by

creating a new piece of hardware.

The aim of the research elucidated in this paper is to utilize these new devices and the powerful GNU Radio framework to build a fully functional Single-Input Single-Output (SISO) channel sounder. The final SISO sounder is able to characterise 1.8 MHz of unique bandwidth with a resolution of 34 kHz reliably, utilizing the IEEE 802.11a Orthogonal Frequency Division Multiplexing (OFDM) preamble.

2 SOFTWARE DEFINED RADIO

A generic SDR requires a Radio Frequency (RF) front end, a sampling stage and a high speed Field-Programmable Gate Array (FPGA) processor for digital signal processing. The basic idea is that the computer does most of the waveform specific manipulations, while the hardware provides sampling and amplification capabilities. Purely software radios are not yet used in real systems, but a progressively higher proportion of transceiver functionality is becoming software based. Software defined radios are often used in testing and prototyping new wireless systems. There are still some drawbacks to these devices however, mainly to do with limited hardware capability. Analogue to Digital and Digital to Analogue (ADC/DAC) devices able to sample bandpass signals are generally not price effective, leading to many SDRs implementing a super-heterodyne down conversion before digitisation. It is also

difficult for these devices to sample low power signals and thus a system normally requires a low-noise amplification stage before digitisation. There is also the bottleneck between the hardware front end and the software platform itself, in this research a USB cable. However, as the technology improves one can see software defined radio becoming the future of communications. This is due to the fact that the powerful software abstraction offered by SDR allows a robust system to be built easily depending on circumstances.

3 IMPLEMENTATION

The Universal Software Radio Peripheral (USRP) [1, 2] is designed to allow users to program a high bandwidth radio using software. Only the essential functionalities, namely sampling conversion and RF amplification, are implemented in hardware. The sampling conversion is done in both an Altera Cyclone FPGA and an AD9862BST ADC/DAC chip. The USRP offers up to 4 RF front ends, which are implemented on *daughterboards*. The front end is based on a superheterodyne structure. The USRP is designed to work with the GNU Radio open-source platform [1].

3.1 Hardware



Figure 1: USRP motherboard with the low frequency receive and transmit daughterboards shown on the left, and the RFX2400 shown on the right.

The USRP motherboard, with two low frequency daughterboards and one high frequency transceiver connected, is shown in Figure 1. The USRP transmits and receives baseband data from the computer via a USB2.0 connection. A maximum rate of 32 MB/s is able to be transmitted across the USB2.0 connection [1]. This limit was not always achieved due to the limited hardware of the host computer.




Figure 2: ADC/DAC and FPGA sections of the USRP, showing the sampling conversion stages and the multiplexer used for routing data. [1]

Figure 2 shows the ADC/DAC and FPGA sections of the USRP. For transmitted data, baseband data from the computer at a certain sampling rate is transferred to the USRP. It is then routed by a demultiplexer to an up-converter and onto a DAC conversion stage sampling at 128 MS/s, before being transmitted via a daughterboard front end. On the receive end, the USRP receives the analogue waveform based at an Intermediate Frequency (IF) from the daughterboard and digitizes this at 64 MS/s. The data is then routed by a multiplexer to a down-conversion stage, in order for it to be transmitted across the USB interface. The down-conversion is done on the FPGA embedded within the USRP, which implements Finite Impulse Response (FIR) Cascaded Integrator-Comb (CIC) and Half-Band (HB) filters in order to change sampling rate and to avoid aliasing. The up-converter consists of a CIC stage within the FPGA, as well as a final HB stage implemented in the DAC chip. The conversion stages are based on the standard I/Q quadrature multiplier.

The CIC and HB filters of the sampling conversion stages attenuate the spectrum of the baseband waveform near the Nyquist limit significantly. Figure 3 shows the magnitude response of the CIC-HBF receive filters, with a decimation rate of 32. Utilizing complex I/Q receiver sampling at 64 MS/s, this leads to a maximum sounding spectrum of 2 MHz. This was kept for the rest of the research, as it avoided data corruption.

The daughterboards act as the analogue radio frequency front end of the USRP. The USRP motherboard offers 4 different slots, two for receive and two for transmit. Each slot has access to two of the AD/DA slots on the AD9862BST. The RFX2400, which has a bandwidth between 2.4–2.5 GHz, was used for the final channel sounder.




Figure 3: Combined receiver CIC-HBF filter response from DC–2MHz, with a decimation rate of 32. [3]

3.2 Software

The system was run using the NetBSD operating system, in conjunction with the GNU Radio platform [1], which contains a substantial library of prewritten C++ signal processing blocks. The blocks, representing sources, sinks and processing blocks have input and output ports connected together via the high level language Python. The GNU Radio platform was utilized in conjunction with MATLAB, which performed all offline signal processing.

4 IMPLEMENTATION

4.1 Synchronisation techniques

The synchronisation of carrier frequency offset and timing are critical to the operation of the channel sounder. Carrier frequency offset (CFO) is caused by the inaccuracy of the receiver and transmitter oscillator clocks, as well as by the Doppler shift [4] caused by relative movement in the system. This causes the received signal to exhibit a low frequency modulation at baseband, even after the full down-conversion process. The basic effect of the CFO is that the demodulated signal is not centered at DC, rather it is centred at δf , which is a function of the Doppler shift and oscillator offset. It is thus essential to take this into account when decoding data. In the project, a Minimum Mean Square Error (MMSE) estimate for the CFO, Moose’s method [5, 6], was used.

Similarly, it is important to know the optimal timing instant to reduce the effect of any intersymbol interference on the received signal. A popular algorithm used in OFDM systems to both detect received packets and to estimate the initial sample number of the packet is the Schmidl-Cox detection algorithm [6, 7].

Moose’s method Moose’s method is based around taking a L length sliding window correlation of the received data, where L is the repetition length of a certain sequence. A minimum of two repeated sequences are required to obtain an estimate of the CFO. In the project, this sequence was the 64-bin OFDM preamble used in the IEEE 802.11a standard [8]. The L length, L lag sliding window correlation product is defined as

$$P_l = r_{l+L}^H \times r_l \quad (1)$$

where the received sequence of length L , L ahead of sample l , is correlated with the delayed sequence at sample l . This measures the phase offset between successive samples and by doing so over the whole sequence, one obtains a rate of change of phase. The MMSE estimate for the carrier frequency offset in Hz is given by [5]

$$\hat{\nu} = \frac{1}{2\pi LT} \arctan \frac{\Re(P_l)}{\Im(P_l)}. \quad (2)$$

An example of Moose’s method, in the 2.4–2.5 GHz range, is shown in Figure 4. The algorithm estimates a carrier frequency of approximately 8.6 kHz, within the ± 5 ppm oscillator frequency uncertainty for 2.45 GHz quoted for the USRP [1].




Figure 4: Moose’s method estimates a CFO of 8.6 kHz as the packet is received at sample number 900.

Schmidl-Cox detection algorithm The Schmidl-Cox detection algorithm measures the energy difference between two successive repeated sequences and the current one. It uses the sliding window of (1) and divides it by the energy of the later sequence, given by

$$R_l = r_{l+L}^H \times r_{l+L}. \quad (3)$$

The Schmidl-Cox detection merit is then given by [7]

$$\rho_{sc} = \frac{|P_l|}{R_l} \quad (4)$$

where figure of merit ρ_{sc} will increase from 0 to 1 as the sliding window correlator encompasses a larger portion of the two repeated L length sequences. The detection merit will reach unity when the correlator encompasses both of the L length sequences, as P_l will be the same energy as R_l . Thus, the value of l that corresponds to the unity level is the initial sample of the packet. With this known, timing synchronisation is achieved. An example of the Schmidl-Cox detection algorithm output is shown in Figure 5. The detection merit rises to unity as a packet is received.




Figure 5: Schmidl-Cox detection algorithm increases from 0 to 1 as the packet is received at sample number 900

4.2 Channel sounding approach

Initial investigations focused on a time-domain cross-correlation technique [9] to find the CIR of a coaxial cable wired channel, utilizing the low frequency daughterboards. The sounding signals used were complex chirps and pseudo-random noise. At an early stage in the research, the response of Figure 3 was not fully understood and thus attenuated received signals. Without knowledge of the internal USRP filter response, the research moved toward a frequency-domain approach to channel sounding. This involved calculating and averaging the CTF, given by $H(\omega) = \frac{Y(\omega)}{X(\omega)}$, in conjunction with IEEE 802.11a OFDM preambles. The effect of the USRP filters was eventually understood and taken into account. In order to reduce the spectral attenuation caused by the USRP filters near the Nyquist bandwidth, the preambles were interpolated in software via zero-padding. Utilizing a high quality coaxial cable and the low frequency daughterboards, a wired equalisation was performed to take into account the action of these filters. The assumption that the filter effect is similar on the low and high frequency boards is a valid one

since they share the same ADC/DAC paths. The response of the the USRP was assumed to have linear phase, due to the FIR CIC/HB filters. The response of the coaxial cable was assumed to be spectrally flat and distortionless. Thus, only the magnitude response was used to equalise the effects of the transmit and receive filters. It was now possible to implement a wireless system, utilizing the 2.4–2.5 GHz daughterboards.

The final channel sounding approach was based on a waveform with three distinct sections, similar to the IEEE 802.11a OFDM standard. Figure 6 shows the received, CFO compensated, waveform.




Figure 6: CFO derotated received waveform starting at sample number 900, showing the three distinct sections of the sounding sequence.

The first section of the sounding waveform consists of 10×64 -bin 802.11a OFDM preambles encoded with BPSK. This was used for the synchronisation algorithms explained above. The MATLAB script applies both algorithms to the received data. It initially finds a detection peak in the Schmidl-Cox algorithm output, and then uses the averaged estimate of the CFO from Moose’s method to correct the residual baseband oscillation. The 64-bin preamble was used because it offered an accurate CFO estimate and avoided the possible aliasing effects of Moose’s method. The second section is used for calculating and averaging $H(\omega)$. It consists of the same 64-bin preambles, but these are interpolated up to a fraction of the Nyquist bandwidth before transmission. The preamble is initially zero padded to a certain percentage of the Nyquist bandwidth in the frequency domain, before being transformed into the time domain via the Inverse Fast Fourier Transform. Only 53 bins of the preamble were utilized, due to the original spectrum consisting of numerous high frequency nulls. Thus, to sound a bandwidth of 70%, the length of one preamble was 75. This waveform was then repeated 10 times, as in the first section. Figure 7 shows the spectrally flat 53-bin preamble, interpolated to 70% of the available bandwidth. The null DC component for

power reduction is clearly visible and is seen in all subsequent results. The first preamble is used as a cyclic prefix, while the next 9 are used to calculate $H(\omega)$. This also inherently takes into account the possible initial phase offset between the transmitter and receiver systems, as this is averaged over numerous runs. The third section consists of 5×53 -bin QPSK encoded OFDM data symbols, used to ascertain whether an accurate estimate of the CTF has been found. As with the second section, these symbols are interpolated up to the same bandwidth as the sounding signal. The first symbol is used as a cyclic prefix, with the next 4 being used to check the CTF. The script applies a zero-forcing equalisation to the received synchronised data, then ascertains whether the correct QPSK data has been received with a threshold detection device.




Figure 7: Spectrum of 53-bin 802.11a OFDM preamble interpolated to 75 bins, 70% of the original bandwidth.

5 RESULTS

The channel sounding system was set up in a small office, with two stationary USRPs separated by 3m. An example of the averaged magnitude CTF at 2.45 GHz is shown in Figure 8. This depicts 70% of the available 2 MHz spectrum being sounded. It includes the channel sounding result with and without the equalisation of the internal USRP filters. As shown, a significant amount of the bandwidth is attenuated by the CIC and HB filters on both transmit and receive ends. The equalised CTF is shown to exhibit flat fading in this particular region of spectrum. The phase response over this same bandwidth is shown in Figure 9. It is clearly linear, leading to the conclusion that there was no significant distortion through the channel.

In order to ascertain the validity of the CTF estimate, the constellation diagram of the received QPSK encoded OFDM was plotted. An example of this is presented in Figure 10, showing the correct received data. Although not significant enough to cause an erroneous decision, the




Figure 8: Magnitude CTF response centered at 2.45 GHz, with 70% of available bandwidth being sounded.




Figure 9: Phase response of CTF centered at 2.45 GHz, with 70% of available bandwidth being sounded.




Figure 10: Constellation diagram of equalised received QPSK encoded OFDM data sequence

data occasionally exhibited a phase rotation of up to $\frac{\pi}{8}$. This was due to the inaccuracy of Moose's method and the phase drift between the CTF averaging section and the received data. The channel sounder was able to consistently sound up to 90% of the Nyquist bandwidth of the system. A fully automated system was programmed to allow the user to define the data encoded in the third section of the packet, while creating the correct waveforms for 5 different bandwidths. The wired response of the USRP filters to these bandwidths was also measured, in order to equalise their response.

6 CONCLUSIONS AND FURTHER WORK

The project aim was to build a fully functional channel sounder utilizing the USRP and the GNU Radio software platform. Initially, complex chirps and pseudo-random noise were used to sound a coaxial cable wired channel. The final wireless implementation used the IEEE 802.11a preamble, along with the high frequency daughterboards.

The final form of the channel sounding waveform consisted of three distinct sections. The first used the 802.11a preamble for a robust synchronisation algorithm, involving the Schmidl-Cox detection merit and Moose's method for carrier frequency offset estimation. The second was for calculating and averaging the CTF, utilizing the same preamble. The third section was an appended QPSK encoded OFDM data waveform, for testing the validity of the CTF. The response of the USRP filters was equalised utilizing a high quality coaxial cable, allowing the actual CTF to be measured. By checking the received QPSK data with a threshold decision device, the validity of the estimated CTF was verified. The hardware setup had two main bottlenecks, namely the USB cable and the PC itself. This meant that a minimum decimation rate of 32 had to be used. Due to the receive end sampling at 64 MS/s, this led to a maximum sounding bandwidth of 2 MHz. Numerous percentiles of that bandwidth were able to be sounded with the system, depending on the relative importance of range versus precision. Using the full bandwidth and the 53-bin IEEE preamble leads to a resolution of 38 kHz. This implementation was not as reliable as the system that sounded 90% of the available Nyquist bandwidth, leading to a sounding range of 1.8 MHz and resolution of 34 kHz. With this system operational, a fully functional SISO channel sounder for use in the 2.4–2.5 GHz range was built using the USRP and the GNU Radio platform.

The following extensions and improvement are envisaged as a continuation of this research:

- Investigate MIMO channel sounding with the use of orthogonal sets of sequences.

- Implementation of a GNU Radio C++ block for rapid detection using the Schmidl-Cox detection algorithm.
- Implement the channel sounder in the FPGA itself, bypassing the USB bottleneck and reducing the effect of the internal USRP filters.
- Upgrade to the USRP2 platform [1], which improves all the ADC and DAC paths to allow far higher bandwidths. This system is based on an Ethernet interface, rather than USB.
- Use the XCVR2450 daughterboards to sound the 5 GHz band.
- Investigate other carrier frequency offset and detection methods that could improve accuracy.

REFERENCES

- [1] <http://gnuradio.org/trac>.
- [2] <http://www.ettus.com/>.
- [3] F. A. Hamza, "The USRP under 1.5X Magnifying Lens!," 2008.
- [4] J. G. Proakis, *Digital Communications*. McGraw-Hill, 2001.
- [5] P. H. Moose, "A Technique for Orthogonal Frequency Division Multiplexing Frequency Offset Correction," *IEEE Transactions of Communications*, Vol. 42, No. 10, 1994.
- [6] A. J. Coulson, "Maximum Likelihood Synchronisation for OFDM using a Pilot Symbol: Algorithms," *IEEE Journal On Selected Areas In Communications*, Vol. 19, No. 12, 2001.
- [7] T. M. Schmidl and D. C. Cox, "Robust Frequency and Timing Synchronization for OFDM," *IEEE Transactions of Communications*, Vol. 45, No. 12, 1997.
- [8] <http://www.ieee802.org/11/>.
- [9] R. J. Pirkel and G. D. Durgin, "Optimal Sliding Correlator Channel Sounder Design," *IEEE Transactions on Wireless Communications Vol.7, No.9*, 2008.

Design of a Low-power Narrow Band amplifier for HF Radio Applications using CMOS FETs

Muhammad Khurram¹ and Rezaul Hasan²

Center for Research in Analog and VLSI Microsystems Design (CRAVE),
School of Engineering and Advanced Technology (SEAT), Massey University,
Auckland, New Zealand

Email: ¹m.khurram@massey.ac.nz; ²hasanmic@massey.ac.nz

Abstract: This paper reports the design of a narrow band amplifier, operating in HF (high frequency) radio spectrum using 250um long-channel discrete CMOS FET (complementary metal oxide semiconductor field effect transistor) devices. A novel g_m -boosted common-gate topology is presented to provide 50Ω input match at the frequency of operation with improved noise figure. In the design of the amplifier, current reused technique is adopted with a novel enhancement to provide g_m -boosting with constrained power consumption. The amplifier provides around 16 dB of gain with excellent noise figure of 1.6-2.4 dB. The complete amplifier constructed using the discrete CMOS FET devices consumes approximately 50mW of power at 9V power supply.

Keywords Discrete CMOS, Common Source, Common Gate, Active Input Matching, low-noise amplifier, g_m -boosted amplifier, noise optimization

1. INTRODUCTION

In a typical single transistor narrow-band tuned amplifier, several solutions based on the common source (CS) and the common gate (CG) configurations are well known. The noise figure (NF) of the CS amplifier is linear with increasing frequency and that of the CG amplifier is almost independent of frequency of operation with higher noise floor as compared to the CS amplifier [1]. Using g_m -boosting techniques, noise contribution of the CG amplifier can be optimized [2]. Due to very high input impedance of the CS amplifier, the CG amplifier is a better alternative at frequencies in HF range. The input impedance of long channel CG topology is given by $1/g_m$ at low frequencies; where g_m is the device transconductance and can be matched to the HF source by carefully selecting the transistor's operating point. Due to better input matching and noise performance, CG configuration is a suitable choice for long channel CMOS amplifier. This paper presents a novel design of a discrete g_m -boosted CG narrow band amplifier to reduce noise and power consumption employing an enhancement of the current reused technique, operating in the HF range.




Fig.1. (a) Conventional CG amplifier, and, (b) g_m -boosted CG amplifier.

2. G_m -BOOSTED CG AMPLIFIER

2.1 Principle of Operation

Fig.1 shows the conventional and the g_m -boosted CG amplifier configurations.

2.2 Proposed Novel Current Reused g_m -Boosted CG amplifier

As in [5], the pMOS CS stage can be used to provide the inverting gain to boost the transconductance of the CG narrow band amplifier. The noise contribution of pMOS CS stage would be negligible as compared to that of the CG stage. DC power consumption could be optimized by sharing the bias current between the CS and the nMOS CG amplifying stages.

In Fig.2, a novel g_m -boosted CG amplifier architecture is shown. Here, the CG nMOS input stage M_1 and CS pMOS inverting gain stage M_2 shares the same bias current supplied by the current source M_3 . The input HF signal is fed to both the transistors through large bypass capacitors C_1 and C_2 that act as short circuit at high frequencies. The inductor L_1 acts as a short circuit at DC and provides enough reactance at high frequencies to decouple the two stages. Inductor L_2 and capacitor C_3 are used as the series resonant LC tank and hence a low impedance path between drain of M_2 and gate of M_1 is established at resonance. C_4 is another bypass capacitor to decouple the DC levels at the drain of M_2 and output.

For the circuit in Fig. 2, (2) can be written as:

$$F = 1 + \frac{\gamma}{\alpha (1 + g_{m_2} R_{Ld_2})} \Big|_{(1 + g_{m_2} R_{Ld_2}) g_{m_1} R_s = 1} \quad (4)$$

Here, R_{Ld_2} is the real part of the impedance at the drain terminal of M_2 .




Fig.4. Amplifier test board with output buffer and SMA input output connectors.

3. EXPERIMENTAL IMPLEMENTATION AND RESULTS

The complete schematic of the proposed circuit is shown in Fig. 3. Its implementation on a test board is shown in Fig.4. Dual N-Channel and Dual P-Channel Matched MOSFET Pair ALD1103 [6] with $V_{THN}=0.7V$ and $V_{THP}=-0.7V$ are used for this purpose. The MOSFETs are fabricated as 250um long channel devices for a supply voltage in the range of 2V to 12V. A 9V power supply was used in this implemented circuit draining 5.5mA reuse bias drain current through the NMOS and the PMOS devices. Inductor L_1 is chosen to be 1mH to provide isolation between M_1 and M_2 . Series LC tank is tuned at 2 MHz HF frequency. 50Ω matched output buffer is also integrated on the test board. To increase the gain of the amplifier, the channel width 'W' of the active devices is increased by implementing them using multiple transistors connected in parallel. The nMOS devices, M_1 and M_3 , are implemented using two n-type MOSFETs connected in parallel. Considering the low mobility characteristics of p-type channel carriers, pMOS device M_2 in the circuit is implemented by connecting four p-type MOSFETs in parallel fashion. Fig.5 shows the time domain wave form trace of tone test at 2 MHz center frequency, taken from the oscilloscope which indicates high linearity affirmed by the spectrum analyzer photo in Fig. 6. The complete frequency domain behavior of the amplifier is shown in the next few figures. The input reflection had a low of -26dB as shown in Fig. 7 which compares favorably to many recent CG narrow-band designs. The amplifier provides a reverse isolation of around -45dB (S_{12} as shown in Fig. 8) and a forward gain of around 14dB (S_{21} as shown in Fig. 9) at the tuned HF of 2MHz. The output reflection had a low of around -24dB as shown in Fig. 10. The amplifier also shows very good noise performance with the NF_{min} of 2.4 dB.

4. CONCLUSION

An HF amplifier using improved g_m -boosted current reused technique is designed and integrated on board using 250um long channel discrete CMOS transistors. The bandwidth of the amplifier is around 500 KHz with center frequency of 2MHz. and it consumes around 5.5mA current at 9 volts power supply with a 2.4dB minimum noise figure.

5. REFERENCES

- [1] Yuna Shim, Chang-Wan Kim, Jeongseon Lee, and Sang-Gug Lee, "Design of full band UWB common-gate amplifier," IEEE Microwave and Wireless Components Letters, vol. 17, no. 10, pp. 721-723, Oct. 2007.
- [2] Maxwell, D., Sungyong Jung, Heechan Doh, Gao, J. and Youngjoong Joo, "A two-stage cascode CMOS amplifier for UWB wireless systems," 48th IEEE Midwest Symposium on Circuits and Systems, Vol. 1, pp. 627-630, Aug. 2005.

- [3] T. H. Lee, The design of CMOS radio-frequency integrated circuits, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [4] Xiaoyong Li, Sudip Shekhar, and David J. Allstot, "Gm-booster common-gate amplifier and differential colpitts VCO/QVCO in 0.18um CMOS," IEEE Journal of Solid State Circuits, Vol. 40, No. 12, December, 2005
- [5] Chung-Yu Wu, Yi-Kai Lo, and Min-Chiao Chen, "A 3.1~10.6 GHz CMOS direct-conversion receiver for UWB applications," IEEE ICECS, 2006.
- [6] ALD1103 Datasheet, http://www.adlinc.com/adl_mosfetarrays.htm, Advanced Linear Devices Inc., Sunnyvale, Ca, USA.




Fig. 5. Oscilloscope trace of 2 MHz tone test at the output of the amplifier.




Fig. 6. Frequency spectrum of the Oscilloscope trace in Fig. 5.




Fig. 7. Input reflection coefficient S_{11} .




Fig. 8. Reverse isolation S_{12} .




Fig. 9. Forward gain S_{21} .




Fig. 10. Output reflection coefficient S_{22} .




Fig. 11. Noise Figure variation with frequency.

Life Sign Detection on a Disposable Robotic Platform as Part of a Three-Tier System for Urban Search and Rescue Operations

B.M.M. Drayton, D.A. Carnegie

School of Engineering and Computer Science
Victoria University of Wellington
P.O. Box 600
Wellington, New Zealand
ben.drayton@gmail.com

Abstract:

The successful inclusion of robots in search and rescue applications has so far been unrealised. Barriers such as high cost and the requirement for an expert human operator have prevented widespread use. A three tiered system under development at Victoria University aims to overcome these barriers. This paper focuses on the development of a proof of concept for the bottom tier of robots. Key issues are: small size to allow it to penetrate rubble; low cost as the robots are designed to be disposable; and the intelligent combination of several life detecting sensors.

Keywords: Human detection robot

1 INTRODUCTION

Urban Search and Rescue (USAR) scenarios are often characterised by victims trapped under the debris of a multi-storey structure. Current search techniques for such victims include the use of Ground Penetrating Radar and void identification and clearing. These techniques are both time and resource intensive. Indicatively, using traditional techniques, it takes on average, ten people, ten hours to locate and rescue one trapped person. There is only a 48 hour window in which rescue is likely to result in a live person being extracted. The introduction of autonomous robotic search teams could greatly reduce the time required to find these trapped victims, with the potential for dramatically reduced casualties [1].

Robotic teams have been implemented in the past [2] [3] however due to high costs and low success rates they have not been widely accepted. The implementation of a multi-tiered hierarchical system can greatly reduce the development and implementation cost. The requirement to have a trained human operator, which has been a large barrier to implementation in the past, can be removed by the introduction of a control tier while the use of a large number of low cost bottom tier robots increases the chances of finding a trapped victim.

Victoria University of Wellington has a three tiered system currently under development to tackle this issue. The highest tier of robot, designated Grandmother, arrives at the boundary of a USAR site and deploys the other robots. In this proof-of-concept, shown in Figure 1, the Grandmother takes the form of a multi-terrain

tracked mechatron capable of carrying a payload in excess of 100 kg. It is equipped with powerful PC computer and possesses the bandwidth to facilitate the coordination of several hundred other robots [4].



Figure 1: Grandmother proof of concept robot

When the Grandmother arrives at a USAR site it deploys 4-6 second tier robots, designated Mother. These mother robots, shown in Figure 2, must be rugged and able to navigate the difficult, rubble strewn terrain typical of a USAR environment. Potentially, because of falls or tipping, these mother robots must also be able to operate inverted. The specifications of this hierarchical USAR application require each mother to contain a payload of 30 third tier Daughter robots which will be distributed over the disaster zone.



Figure 2: Mother proof of concept robot

The prototype versions of the daughter robots are expected to be approximately the size of a cell phone and are designed to penetrate beneath the surface of the rubble. They are to be low power, disposable robots equipped to detect the presence of trapped humans. They must be inexpensive as no attempt will be made to retrieve them. These robots must be capable of coordinating a variety of sensors in order to detect the trapped victim. Specific sensors include: heat, CO₂, motion, and sound. Once a person is identified, this information, along with the robot's approximate position, is relayed back to the Mother robot.

Both the Grandmother [5] and Mother [6] robots have been developed to the proof of concept stage and significant work has been done on networked robot control and coordination [7]. This paper discusses the development of the lowest tier robots, the Daughters, to a proof of concept stage.

2 SENSOR SYSTEMS

As mentioned in the previous section, the proof of concept design for the Daughter robot focuses on the intelligent combination of a variety of different sensor inputs in order to detect the presence of trapped human beings. The robot must be self powered, contain sufficient intelligence to interpret its environmental data, and be able to communicate with a nearby mother robot. To be useful in this application, the daughter must also be inexpensive since they are disposable, and have an effective operating time of several hours.

An Atmel AT89C51AC3 microcontroller is used to provide intelligence to the robot. It provides four PWM outputs for motor control and four ADC inputs to read sensors. It has both an idle and power down mode to decrease power usage and the Victoria University faculty have significant experience with it, easing the development.

A ZTP-135S thermopile was selected to provide non-contact directional heat measurements. The ZTP-135S

was selected as it provides adequate sensitivity while being cheap and easily available. It includes a flat IR filter and a thermistor for temperature compensation. The sensor response after amplification is shown in Figure 3. This shows the sensor response to a reference heat source, a human hand, placed at set distances from the sensor.




Figure 3: Characteristic response of ZTP-135S thermopile

The AT89C51AC3 has a 10 bit ADC with an input range of 0 V to 3 V resulting in 3 mV resolution. From the characterisation data, the heat from a human hand is distinguishable up to 90 cm.

As it is a passive sensor, a large gain non-inverting amplifier stage with controllable gain and offset is used to interface the sensor with the microcontroller. The thermopile suffers from significant sensor drift so potentiometers must be adjusted to calibrate the sensor before use. As the Daughter robots will be disposable, only a single automatic calibration will be required, therefore this is not a huge impediment.

Even rail to rail operational amplifiers have a dead band large enough to degrade the signal from the thermopile. An LMC7660IN charge pump is used to provide a negative voltage rail to the operational amplifiers to counteract this problem. It also makes amplification of the signal from the microphone simpler.

A second order low pass filter is used to remove high frequency noise from the sensor signal, particularly a peak at 6 kHz introduced by the charge pump. Because the large gain stage creates sizable random fluctuations an integrator is implemented to smooth its response.

An MS401 pyrodetector is used as a motion sensor. It combines a PIR sensor, an operational amplifier and a comparator to trigger on changes in the infrared spectrum, such as the body heat in moving people. The internal circuitry of the MS401 is shown in Figure 4.




Figure 4: MS401 motion detector internal circuit

With testing the MS401 it was found to be too non-directional. A high level of directionality is desirable as it allows the Daughter robot to move towards a potential trapped human and use its other sensors for confirmation. A paper sheave was added during testing to increase the directionality of the sensor.

To characterise the motion detector in a way meaningful to USAR environments, three levels of motion are defined. Small movements, defined as movement of the fingers, medium movements, defined as limb movement and large movements which are full body movements, for example jumping, were used for testing. The results of these initial experiments are illustrated in Figure 5.




Figure 5: Characteristic response of MS401 motion detector with (top) and without (bottom) paper sheave

Measurements were taken on a grid using 0.5 m spacing perpendicular to the sensor's axis and 1 m in line with the sensor. Each grid space was assigned a value depending on the smallest class of movements it would reliably respond to. The experiment was repeated with and without the paper sheave. Figure 5 shows both sets of results.

An AOM-4542P-R omni directional PCB mount microphone is included to detect human vocal sounds. A second order low pass filter restricts the bandwidth of the microphone to 3.4 kHz and a DC bias stage is used to prevent the microcontroller's ADC from clipping the microphone signal. The circuit used is shown in Figure 6.




Figure 6: Microphone circuit schematic

The microphone response was characterised using a W4-654S woofer speaker [8]. The frequency response of the speaker is shown in Figure 7 and the microphone circuit response in Figure 8.




Figure 7: Frequency response of W4-654S speaker [8]




Figure 8: Microphone response

The Cooley-Tukey FFT algorithm is used to transform sampled microphone data to the frequency domain. The maximum achievable sampling rate with the AT89C51AC3 is 4 kHz. While larger bandwidth is required to uniquely identify individual speakers this low bandwidth is acceptable for differentiating vocals from background noise. Due to the processing limits of the microcontroller the size of the FFT is limited to 64 samples. Using a 4 kHz sampling rate this gives a 62.5 Hz frequency resolution, which is acceptable for the peak detection algorithm used. Peak detection is implemented by comparing each bin to the average magnitude. An array is formed with the frequency of each peak which is used to differentiate human vocals from background noise such as falling rubble.

An RS-5600 CO₂ sensor is included to detect the build up of carbon dioxide typical of humans in enclosed spaces. The RS-5600 includes a solid electrolyte cell and a heating resistor. Because the heating resistor requires significant power a MOSFET controlled by the microcontroller is included so it can be switched on and off. This circuit is shown in Figure 9.




Figure 9: CO₂ sensor circuit schematic

Qualitatively the CO₂ sensor was found to be sensitive to human breathing. To provide a more quantitative response characteristic a sample of dry ice was placed in a sealed enclosure with the sensor. Measurements were taken of the sensor response versus time as the dry ice sublimated and are shown in Figure 10.




Figure 10: CO₂ sensor response to sublimating dry ice

The recovery time of the CO₂ sensors recovery characteristic was also measured with the robot removed from the sealed enclosure. This is shown in Figure 11.




Figure 11: CO₂ recovery characteristic

3 OTHER DESIGN FEATURES

The focus of this proof of concept is on the development of sensors into an intelligent human detection system. Physical suitability in terms of ability to penetrate rubble and survive rugged environments is outside the scope of this work, but this project does endeavour to provide a mobile platform in order to future test the delivery of these sensors, and to establish the communication protocols with a remotely located mother. Hence a proof concept (not prototype) mobile robot was constructed. The body of this proof-of-concept robot is formed around a double sided PCB board. To provide basic locomotion a dual motor and gear box provides differentially driven rear wheels, with two freewheeling front wheels. Figure 12 shows a photograph of the robot.



Figure 12: Photograph of the Daughter robot proof of concept

Power is provided by two 3.7 V 1100 mAh Polymer Lithium Ion batteries in series. Rechargeable batteries were chosen for ease of use with prototyping. A final design would use single use batteries as the robots will be disposable.

4 EVALUATION

The maximum audio sampling rate achievable while executing no other code was 4 kHz. This was unsatisfactory as a bandwidth of 3.4 kHz was desired as this approximately matched the telephone bandwidth, which is proven to provide vocal intelligibility. It also necessitates a change in the anti-aliasing filter.

Figure 13 shows the time taken to process the FFT versus the number of bins. Preliminary tests show that the data collected using a 64 bin FFT is sufficient to distinguish between vocal noises and background sounds (such as falling rubble). However, because of the microcontroller's processing limitations this system can not be integrated to run in real time with the other programming required for the robot.




Figure 13: FFT processing time versus bin size using the AT89C51AC3 microcontroller

The CO₂ sensor has proven to be sufficiently sensitive to detect the presence of humans however it has a slow response and recovery time. Because the sensor's heating resistor requires a large amount of power this is a significant problem. Due to the difficulty in finding other available sensors, and the fact this sensor has been shown to be usable for detecting humans, the sensor will be persevered with for this stage of testing. Work will continue into trying to find a better alternative.

Both the CO₂ and thermopile sensors are sensitive enough to detect trapped humans however both suffer from drift. This has to be corrected by adjusting gain and offset potentiometers to maximise the usage of the 0 V to 3 V range of the microcontroller's ADC. Testing has shown that this drift is only significant over a length of time greater than a day. Because of the disposable nature of the robots this means only a single calibration, just before deployment, is required.

The effect of ambient light on the motion detector was tested by running detection tests as described above in both a darkened room and with bright light incident on the pyrodetector. Figure 14 shows both results.

The results show that while ambient light conditions do have an effect on the sensor response even bright inci-

dent light does not render the sensor ineffective. In typical USAR environments the Daughter robot is likely to be inside collapsed structures and therefore low light levels can be expected, aiding the sensor's effectiveness.




Figure 14: Motion detector response in dark (top) and bright (bottom) conditions

5 FURTHER WORK

Because the AT89C51AC3 has shown poor performance for audio signal processing a more powerful microcontroller should be selected.

Sensor drift was a significant problem with the CO₂ and thermopile sensors. Using higher quality sensors could mitigate this however it would increase the unit cost. Automatic calibration could be implemented by replacing the potentiometers with digitally controlled equivalents. When the Grandmother arrives at the USAR site the Daughter robots could be calibrated in the reasonably stable environment of the Grandmother robot.

Communication with the Mother and other Daughter robots requires the addition of a ZigBee module, which will provide both the means of communication and localisation. The use of ZigBee communication modules for this purpose has already been tested and shown to work [9]. This needs to be interfaced with the microcontroller. ZigBee modules with included controllers are widely available so interfacing should be a relatively simple task.

Currently locomotion is controlled using a joystick input connected to two ADC channels on the microcontroller. Once a ZigBee module is included, locomotion will be controlled by a combination of search algorithms and commands received from Mother robots (who in return has received the global instructions from the grandmother).

Because the proof of concept was developed with ease of development in mind it uses mainly thru-hole components. The size can easily be reduced by using surface mount devices and more compact spacing. Electronics for programming should be removed as pre-programmed chips would be used.

6 REFERENCES

- [1] M. Rozakis, "The socio-technical organisation of Urban Search and Rescue (USAR) operations: Learning from Earthquakes in Greece", European Sociological Association Conference, 2005
- [2] A. Drenner, I. Burt, B. Kratochvil, B. J. Nelson, N. Papanikolopoulos, K. B. Yesin, "Communication and mobility enhancements to the scout robot", Proceedings of the IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems, Lausanne, Switzerland, Oct. 2002.
- [3] R.R. Murphy and S. Stover, "Rescue robots for mudslides: A descriptive study of the 2005 La Conchita mudslide response: Field Reports", Journal of Field Robotics Volume 25 issue 1-2, 2008.
- [4] D.A. Carnegie, "A Three-Tier Hierarchical Robotic System for Urban Search and Rescue Applications", IEEE International workshop on Safety, Security and Rescue Robotics, 2007
- [5] D.A. Carnegie and J.C. Cordes, "The mechanical design and construction of a mobile outdoor multi-terrain mechatron," Proceedings of the Institution of Mechanical Engineers, Part B, Engineering Manufacture, Vol 218, pp 1563-1575, Nov 2004
- [6] B. Howarth, P. Gaynor and D.A. Carnegie, "Design of a medium sized wheeled robot for urban search and rescue," in Proc. 12th Electronics New Zealand Conference, Christchurch, New Zealand, 2006, pp. 373-378.
- [7] T.M. Roehr, "Control of a hierarchical team of robots for Urban Search and Rescue" MSc. Thesis, Victoria University of Wellington, 2008.
- [8] Internet WWW page at URL: <http://www.audiolabga.com/pdf/W4-654S.pdf> last accessed 28/09/09.
- [9] C. Tsui, D.A. Carnegie, Q.W. Pan, "USAR Robot Communication Using ZigBee Technology", IEEE International workshop on Safety, Security and Rescue Robotics, 2009

The design and implementation of a reliable robotic pipe inspection system

S. J. Winch, D. A. Carnegie
Victoria University of Wellington

School of Engineering and Computer Science
Victoria University of Wellington
P.O. Box 600
Wellington, New Zealand
Stephen.winch@vuw.ac.nz

Abstract: Associated Environmental Services (AES) is the company contracted by the Wellington City, Hutt City and Porirua City councils to inspect all storm water and sewage pipes in the local region. They require robust and reliable units where the maintenance and replacement of components happens in a timely fashion in order to keep up with the demand of their clients. Hence the objective of this Masters project is to design and implement a robust and reliable pipe inspection system based primarily on off the shelf products.

Keywords: Pipe, inspection, robot, robust, reliable, AES, environmental, sewage, drain, storm water, circuit, mechatronics, electronics.

1 INTRODUCTION

Associated Environmental Services (AES) own a number of pipe inspection robots that are non-functional some of which are no longer supported by the manufacturer. Repair options are currently restricted to sending the units overseas to the British company PearPoint^{[1][2]}. This is an expensive and time consuming operation. Also, since they use custom components in all of their repairs, servicing by AES technicians is not an option. Indicatively, a replacement robot would cost AES in the order of NZD\$150,000^[1].

This non-support of older robotic units and a scarcity of repair options provided the motivation of this project. The goal is to create hardware and software solutions for the AES robots and control stations as well as a software interface that can be easily used by the field technicians. The ideal outcome for AES would be for the robots to be able to be repaired by a person with minimal training in electronics. The robots considered in this project are two non-functioning robots from the Germany company Optimes^{[1][3]}.

Initially the smaller robot of the two is considered. This robot measures 630 mm in length, 120 mm in width and 120 mm in height and weighs approximately 18 kg including the camera head. The robot is four wheel drive and has the ability to inspect 360° of the pipe in close proximity. It is only able to travel in a straight line but the head of the robot, housing the camera, is able to pan 200° and rotate 360° making it a highly useful system.

The main body of the robot is solid brass with 12 mm thick walls, inside of which are two geared motors, a rotary encoder, two motor drives, the main control board, a 24 V DC to DC convertor with external protection circuitry and all necessary wiring. Hence designing compact circuitry is an important consideration.

In the camera head of the robot an analogue Sony camera is housed along with 4 high powered LEDs for lighting the inside of the pipe, a motor to control the panning of the camera and a camera control board. The camera control board handles the transmission of control signals to the camera (such as focus and zoom), the transmission of the video signal back through the robot to the user, and provides the power to the LEDs and camera. This internal section of the head is only able to fit a control board within a 57 mm × 57 mm space, hence the circuit must be compact.

The robot is also at the end of a 200 to 300 metre long cable. When the robot is inserted into a pipe, communication is handled from a laptop computer via this cable. Wireless communication is not required as the robot needs a power supply which is also provided via this cable that can easily support transmission wires. The robot is unable to internally carry its own power supply due to its small dimensions dictated by the size of the pipes it must inspect.

2 PREVIOUS WORK

This project was started by another Victoria University of Wellington student, Aleks Ristich, in 2008. During his

work Mr Ristic was able to reverse engineer or redesign most of the pipe inspection robot's internal circuitry and create a Graphical User Interface (GUI) in Labview [4]. The internal circuitry of the robot involves a main control board which is designed to implement serial communication between the computer and the robot, a camera control board to control the camera and the 4 lighting LEDs and three motors, two of which control the pan and tilt of the camera head with the third controlling the forward movement of the robot. The final part of the internal circuitry is the power module which converts the incoming 300 VDC to the 12 VDC at 10 A maximum output and the power module protection board.

AES field technicians found the Labview GUI to be extremely hard to use and requested for it to be changed. Mr Ristic had started the conversion into Visual Basic with the .net framework setting up the initial serial communication and adding the hexadecimal codes for activating and controlling the robot. It was controlled by clicking on the appropriate command button in the GUI. By the end of his project Mr Ristic was able to control the robot via the Labview code, with the video from the camera displayed on the computer screen and capture to file, but had not yet implemented the Visual Basic GUI. The control was only possible when the robot was directly connected to the computer, and not possible via the cable.

3 CURRENT WORK

As discussed above, there were a number of tasks that were not completed by Mr Ristic, and with further feedback from AES management and field technicians, a substantial redesign was required.

3.1 Internal Control Board

The first change that needed to be implemented was converting the circuitry to handle a 24 V input instead of a 12 V input and to have the DC to DC converter adequately protected. This voltage was changed as the main motor is rated as a 24 V motor and the speed of the robot operating with a 12 V supply was not sufficient for AES. To achieve this, a large number of the components needed to be changed to handle this increased voltage and, in the case of the main driving motor, the increased current drawn.

The first board changed was the power module board that is designed to fit on-top of the 300 VDC to 24 VDC converter (Figure 1). On the initial design of this board there was no protection against voltage spikes, incorrect polarity, or noise. Hence TVS diodes were connected between the positive and negative inputs and outputs to protect against back emf and voltage spikes. Capacitors were also added across these inputs to minimise noise on the lines (Figure 2a). This board was designed to slot

onto the converter to keep the connections robust and the size minimal (Figure 2b).



Figure 1: Vicor 300V to 24V DC/DC converter module, V300B24C250F^[5]

A change to the internal control board was also required. This involved converting the linear regulators to switch mode regulators due to the increased voltage causing the linear regulators to heat up rapidly in normal conditions (motors running freely). Such heat build up is a major concern in the air-tight robot as there can be no air flow to cool the components. Two options for countering this increase in heat are to incorporate heat sinks into the design or to use switch mode regulators. Switch mode regulators have been chosen as the required heat sinks are bulky and size is a major constraint.



Figure 2a: Power module protection board, left-hand side is the 300 VDC input whereas right-hand side is the 24 V output, designed to fit atop of power module.




Figure 2b: PCB layout of power module protection board showing slot-in design and isolation of 300 VDC input

One of the 12 V regulators is shown in Figure 3. This regulator is able to take a maximum input of 40 V and is 93% efficient for the input/output relationship of 24 V to 12 V (Figure 4). There are three regulators used on the circuit board, two 12 V regulators and one 5 V regulator. Two 12V regulators are employed to distribute the high current requirements and reduce the load experienced by any one regulator to less than 3 Amps. The first 12 V regulator feeds the camera motors only, whereas the second regulator feeds the camera board (supplying power

to the camera and LEDs) and the 5 V regulator which in turn supplies power to the onboard ICs. Each of the switch mode power supplies uses capacitors to smooth any ripples. Most of the power is consumed by the two 12 V regulators as the components attached to the 5 V regulator are low current devices. Consequently the majority of the heat dissipation occurs from these 12 V regulators. The maximum heat dissipation can be calculated as the maximum power that needs to be dissipated is 7% of the power input, or 5.04 W due to the rated efficiency of the regulator. The linear regulator that it has replaced was only 50% efficient for this voltage conversion and hence would have dissipated a maximum of 36 W of power as heat.




Figure 3: LM2676T-12, 12V switch mode power supply with schottky diode and shielded inductor.




Figure 4: Efficiency of LM2676T at maximum current through-put for different input and output voltages.

Another major change to this board was to redesign it to slot on top of the DC/DC converter and to plug directly into it instead of using screw terminals to attach the two. This was done to make it easier to put the robot together and to hold the control board firm so wires are less likely to cross or become detached.

Diode packs and other protection circuitry were also incorporated into the control board to protect local and networked components from back emf, voltage spikes, over current, under voltage and noise^[6]. These are vital additions as these robots need to be robust and reliable when operating in rather hostile environments. The final change to the internal control board was to remove the relay that initially turned on the motors. This relay could only handle 3 A whereas the main motor running at 24 V has a stall current of 4.1 A. The relay was replaced with a dual power MOSFET IC (FDS6982AS) which can handle up to 6.3 A continuously. Other components such as the regulators, fuses and diodes were also checked to

see if they could handle the increased voltage and current drawn, and were replaced. A BZG03C27 27 V Zener diode was also inserted in place of the existing 20 V Zener diode (BZG03-C20).

3.2 Motor Driver modification

As the main drive motor is now running at 24 V, the previously selected Sabertooth 2×5 motor driver (Figure 5) can no longer be used and so was replaced with a SyRen10 motor driver (Figure 6) to control the main motor. This new motor driver has a nominal 24 V output and is able to handle 10 A. One of the Sabertooth motor drivers is still used to control the two 12 V camera motors.




Figure 5: Sabertooth 2×5 Motor Driver unit^[7]

The SyRen10 has the dimensions of 35 × 57 × 14 mm compared to the Sabertooth 2×5 with dimensions of 45 × 40 × 13 mm so some of the mounting posts needed to be moved to accommodate this new driver.



Figure 6: SyRen10 Motor Driver unit^[7]

3.3 Recreation of a software control system

A major adjustment to the code was necessary as the technicians at AES were unhappy with using a mouse to click buttons on the GUI to control the robot. It was then found that Visual Basic was not able to complete some of the tasks required by the code (specifically this included video capturing and acquiring a USB game control device) for this robot without an array of external plug-ins. The implementation of video capturing and control of movement and functions with a joystick were implemented after the original code was converted to C# with the Microsoft .net framework installed. For video capturing the Windows DirectShow methods are needed as they are designed to be used to capture, play and record audio or video signals from an attached device. Distinction between multiple devices needed to be made as the controlling laptop contains an inbuilt webcam which has precedence over any externally connected device.

This was accomplished by adding in a device selection method. To implement a USB game control into the control system, the Windows DirectX plug-in needed to be installed such that the computer would recognise the control. The code was tested with a Logitech Extreme 3D Pro (Figure 7) joystick, which is a 3-axis joystick with multiple buttons on the yoke and a Point of View (POV) hat on-top.



Figure 7: Logitech Extreme 3D Pro joystick, all buttons other than the trigger (button 1) are numbered [8].

The code is sensitive to the position, not just the direction, of the yoke. Hence the user has significantly more control over the speed of the motors. This includes the speed the robot is travelling and the speed of pan or tilt of the camera head. This was implemented by determining the input values of the joystick for each extremity and the central position. These were then plotted against the required output values to get a linear conversion equation. Due to slight vibrations on the joystick the bottom 100 input values were ignored and the values in the conversion equation were reduced to three significant figures to make the system more stable.

The x-direction, or forward direction, was the only one where the relationship between input and output was negative, shown in Figure 8. This setup is due to the joystick being designed for aircraft games where a pull back on the yoke would cause the aircraft to rise not dive.




Figure 8: Relationship for forward direction of joystick, relating to driving of robot, output versus input.

The camera motors were controlled by twisting or pushing the yoke sideways causing the camera head to either pan or tilt respectively. The output relationship is posi-

tive and linear with the input from the joystick, as shown by Figure 9.




Figure 9: Relationship for twist and sideways directions of joystick, relating to control of camera motors of robot, output versus input.

The trigger of the joystick has been attached to the emergency stop function which overwrites the current position of the joystick and stops the robot. The button on the side of the joystick is programmed to stop the camera head from moving, which allows the user to drive the robot without accidentally moving the camera. The other buttons, including the ones on top of the yoke and on the base of the joystick were programmed as shown in Table 1.

Table 1: Joystick movement and button functions with their corresponding commands, the commands are sent via RS485 to the robot. X-axis corresponds to forward or backward movement, y-axis to sideways movement and z-axis to twist of yoke, 'X' corresponds to state not cared about and 'V' corresponds to a sent value

Button	Function	Command
Trigger	Stop robot	a077XXXXXX11FF
Yoke x-axis	Drive robot	aVVVV07707711FF
Yoke y-axis	Tilt camera	a077077VVVV11FF
Yoke z-axis	Pan camera	a077VVVV07711FF
Button 2	Stop Camera	aXXX07707711FF
Button 3	Zoom in	8101040702FF
Button 4	Zoom out	8101040703FF
Button 5	-----	
Button 6	-----	
Button 7	Focus far	8101040802FF
Button 8	Focus near	8101040803FF
Button 9	Decrease yoke axis sensitivity	N/A
Button 10	Increase yoke axis sensitivity	N/A
Button 11	Calibrate yoke/joystick	N/A
Button 12	Stop all motors	a07707707711FF
Automatic	Stop zoom	8101040700FF
Automatic	Stop focus	8101040800FF

Table 1 also shows that the functions of stop zoom and stop focus have been setup to happen automatically when the user releases the button, and only happens with its corresponding action; that is, the robot will only be sent “stop zoom” if it has previously been zooming in or out and is not currently doing so.

3.4 Design and implementation of a base station

A board for the base station must be created to facilitate communication between the laptop and the robot. This base station must also be able to handle control of a cable reel which is motorised to prevent the robot running over the cable and breaking the internal wires.

An initial design consideration was the conversion of the USB RS232 signals from the Laptop to the Transistor, Transistor Logic (TTL) levels required by the onboard microcontroller. These TTL levels then need to be converted to RS485 form to communicate with the robot at the end of the cable. Also the reel motor needs to be controlled depending on the actuation value sent to the main driving motor. For the robot accelerating, the reel should be held in a free moving state, for the robot decelerating after forward movement or stopped a small reverse voltage should be applied to the reel motor to keep the cable taut. If the robot is reversing a voltage needs to be applied to the reel motor to cause it to rotate the reel at a greater speed than the robot is moving, this voltage is dependant on the speed of the robot, the gear ratio on the reel and the amount of cable on the reel. The amount of cable on the reel determines the rate at which the cable is leaving the reel hence the angular velocity of the reel must be continuously varied in order to maintain a constant cable velocity.

This board was designed to be generic and accommodate either the 12 V Sabertooth 2×5 (Figure 5) or the new SyRen10 (Figure 6). This provides flexibility in future designs where either voltage can be accommodated.

The reel was initially designed to have a shaft encoder attached to the front, resting on top of the cable. This provides a distance reading to the user which is used to locate the position of problems in the pipe. The reel encoder provides two signals in quadrature which enables both the speed and direction of the robot to be determined. It also of course, provides a measure of the distance the robot has travelled.

The base station provides both a 15 V and a 5 V supply Protection against over voltage, over current and noise was added to each of the inputs before they were used. The 5 V input was used to power all the IC chips on the board including the serial/TTL conversion ICs and the microcontroller. The 15 V supply was used to power the reel encoder and motor.

4 FURTHER WORK

The first task to be completed is to assemble, program and test the current boards, and then verify laptop/robot communication via the base station. Due to the unknown status of the current motor attached to the reel, other

motor options are being considered. The current motor is believed to be a 12 V geared motor, but the gear ratio and state of internal components of this motor are both

To determine a suitable replacement motor (if required) the maximum speed of the robot needed to be determined. This was calculated for the larger, 85 mm diameter, wheels as these would provide the greatest speed. The free running speed of the robot with these wheels is 11.75 m/min which determines the minimum rotational speed of the reel as 12 m/min at the centre with minimal load. The load on the reel at this time will be minimal as for the cable to be exiting from this position a majority of the 300 metre cable must already be deployed, the reverse case is also true, where when the cable has not yet been deployed the load will be at a maximum but the rotational velocity of the reel at the centre does not need to be as high as the cable will be leaving the reel at a point further from the centre hence leaving at a higher velocity.

To make this system more robust and reliable, a significant amount of testing is required. These tests include pressure testing the unit to determine if the system is watertight and stalling the motors when the case is fully sealed to test whether or not the components can handle the continuous current requirements without air flow for cooling. Before pressure testing or stall testing can commence all the seals in the robot need to be replaced as the seals have degraded severely and will currently let in a significant amount of water. An example of the extent of the problem is shown in the camera head model in Figure 11a which is the assembled camera head, whereas Figure 11b is an exploded view of the same camera showing all the separate parts. Each of these parts needs to be attached to its neighbours in a water tight fashion to prevent water from entering. The full extent of the seal replacement required is displayed in Figure 12a and 12b which are the assembled and exploded views of the robot respectively.



Figure 11a: Assembled view of the robots camera head, mid section is connected to a motor so that the camera is able to be panned.



Figure 11b: Robot camera head exploded view displaying the number of parts and extent of sealing required. Model designed in SolidWorks.

The final task in the construction of this system will be to find a replacement cable option as the current cable has degraded over time with sections of the outer casing breached causing the internal wiring to be exposed to the hostile environment. Ideally the replacement cable will be strong, sourced from a local supplier such that if the cable breaks or degrades at some future point then a replacement can be easily be acquired and use military specification connectors as these are the type currently used by the robot. These connectors are designed to hold strong and not allow water to enter. The cable needs to be lightweight such that one person is able to carry the cable and reel to and from the locations as some of these are not on roadways. The cable also needs to be strong in case the robot malfunctions and needs to be pulled out of the pipe.




Figure 12a: Fully assembled robot with military grade connection plug at rear




Figure 12b: Exploded view of the same robot, displaying the full extent of seals that need to be replaced and the overall complexity of the unit.

One manufacturer of cables that is currently being considered is General Cable New Zealand^[9] which is the New Zealand subsidiary of the US-based General Cable with a manufacturing plant in Christchurch. This company is able to provide multi-core cabling options (Figure 13) which are required for this system.



Figure 13: Multi-core wire manufactured by General Cable New Zealand, Production of cable with 2 to 36 cores is possible with each core holding 7 to 16 wires.

5 REFERENCES

- [1] Correspondence with Bryan Pobar, Owner operator of Associated Environmental Services, 27/7/09 – 23/10/09
- [2] Pearpoint Service Centres, Accessed October 28, 2009, from: <http://www.pearpoint.com>
- [3] Optimess information, Accessed October 28, 2009, from: <http://www.optimess.com>
- [4] Ristich, A. 2008. Iterative approach on redesigning and construction of a pipe inspection robot. Masters Thesis, Victoria University of Wellington, New Zealand.
- [5] 24V mini Datasheet. Retrieved September 21, 2009, from: <http://www.vicorpower.com>
- [6] Horowitz & Hill (1989). *The Art of Electronics 2nd Edition*. Cambridge: Cambridge University Press
- [7] SyRen 10 and Sabertooth 2x5 datasheets. Retrieved September 21, 2009, from: <http://www.dimensionengineering.com>
- [8] Logitech® Extreme™ 3D Pro product information. Retrieved October 20, 2009, from: <http://www.logitech.com>
- [9] Multi-core cabling products manufactured by General Cable, Retrieved October 28, 2009, from: <http://www.generalcable.co.nz>

ENVIRONMENTAL MONITORING WITH WIRELESS SENSOR NETWORKS

A. Ghobakhlou¹, T.A.S.A. Perera³, P. Sallis¹, O. Diegel² and S. Zandi¹
¹ Geoinformatics Research Centre, Auckland University of Technology, Auckland, New Zealand
² Creative Industries Research Institute, University of Technology, Auckland New Zealand
³ School of Engineering, Auckland University of Technology, Auckland New Zealand
 Email: akbar@aut.ac.nz

Abstract: A architecture framework is presented for wireless sensor networks (WSN) designed to capture and monitor micro-climates in a crop field. WSN are rapidly improving their place and performance in the automotive industry and for agricultural, industrial and environmental monitoring. Uptake of this technology can also be seen in many other application areas. Recent developments and advances in wireless technology as well as affordability give rise to this emerging field in the realm of Precision Agriculture (PA). Vineyard monitoring is an emerging application field in PA. This paper presents an architecture framework for developing a WSN and gives a brief overview of hardware components needed to build wireless sensor nodes.

Keywords: Wireless Sensor Networks (WSN), Precision Agriculture (PA) and Vineyard Monitoring

1 INTRODUCTION

Wireless Sensor Networks (WSN) have been the subject of research in various domains over the past few years and deployed in numerous application areas. WSN is seen as one of the most promising contemporary technologies for bridging the physical and virtual world thus, enabling them to interact. A WSN is composed of a number of sensor nodes, which are usually deployed in a region to observe particular phenomena in a geospatial domain. Sensor nodes are small stand-alone embedded devices that are designed to perform specified simple computation and to send and receive data. They have attached to them a number of sensors gathering data from the local environment that is being monitored. WSNs have been employed in both military and civilian applications such as target tracking, habitat monitoring, environmental contaminant detections and precision agriculture [3][4].

The work described in this paper is a realisation of a concept outlined in Eno-Humanas project [1]. It is a system for gathering (sensing) and analysing climate, atmosphere, plant and soil data. It is specifically designed for micro-climate analysis in vineyards and other agricultural/horticultural environments. This research has produced and is further developing a prototyping in order to demonstrate how state-of-the-art devices could be used in precision viticulture as a management tool to improve crop yield quantity and it is assumed, crop quality.

2 WIRELESS TECHNOLOGY

WSN technology eliminates connectors, provides safe/flexible connectivity, improves resources sharing, easy installation and mobility. The other advantage is that these systems require a low micro power levels, thus

it can last for longer period of time [7]. Two major protocols are used in wireless networking, IEEE 802.15.4 and ZigBee stack. The ZigBee network stack sits on top of IEEE 802.15.4 standard Medium Access Control (MAC) and Physical (PHY) layers (Refer to Figure 1). MAC and PHY layers define the RF and communication components between other devices. ZigBee stack contains the network layer, an application layer and security service provider (SSP) [8].

These protocols have their own limitations and advantages. The main limitation for the both protocols is low data rate (narrower bandwidth). Therefore these protocols are suitable for low data transmission applications.




Figure 1: ZigBee stack layer adapted from [5]

As shown on Figure 2, as the data rate is increased the power consumption, cost and complexity of the system increases geometrically. Since moisture detection system is located remotely, power consumption needs to be

minimal, preferably battery operation for end of product life. Moreover sensor data does not require wide bandwidth. Therefore as illustrated in Figure 2, a suitable protocol for small sensor network is ZigBee [8].

2.1 Wireless Network Topologies

Three main network topologies used in wireless networking are point-to-point, star and mesh networks (refer to Figure 3). Each of these topologies has their own advantages and disadvantages in different applications.




Figure 2: Standard technology map adapted from [7]

Their main difference is the use and behaviour of network components.

2.1.1 Point-to-point

Point-to-point is the simplest topology and mostly used to replace single communication cable. Point-to-point network can work adequately, when two end points are located close to each other [9].

2.1.2 Star topology

Star topology is also known as a point-to-multipoint wireless system. This system is usually based on IEEE 802.11 or Bluetooth communication standards. This system has one main base station, which controls the communication with all the other end nodes. The reliability of this network depends on the quality of the RF link between the coordinator and each of the end nodes. The main problem with this system in industrial applications is the difficulty finding a suitable place for the central controller in order to communicate with each end node [9].

2.1.3 Mesh topology

Mesh topology is also sometimes called a peer-to-peer system. This system is an ad hoc multi-hop system. Mesh networks are based on the ZigBee protocol, which means each node can be used to send and receive data. In this manner, network consists of multiple redundant communication paths, which can be used in event of node failure. Therefore, data will reach its destination

reliably, via the intermediate nodes (Refer to Figure 3) [9].

The mesh network has three important properties: Self-Configuring, Self-Healing and Scalability. A ZigBee mesh network configures itself automatically. The network identifies new nodes and automatically includes them in the network. Moreover, if one node fails the network re-routes the message through an alternative path. According to the ZigBee mesh network standards, it can support up to 65,536 network nodes (clients) [9].




Figure 3: A mesh network topology applied in a vineyard monitoring application

3 MODELING THE EFFECTS OF CLIMATE CHANGE

The WSN ability to simultaneously capture and relay real time data for analyzing the variability in climate change and its effects on plant physiology is significant for different grapevine varieties. This is because modeling the relationships between climate variability requires both data on the cause and effects recorded without any time discrepancies. Complexity in the models increases with spatial information combined with other environment related parametric variables. It is assumed that in combine with one another this variable set will correlate with grape and consequently wine quality. Gaining more insights into natural systems and their functioning including climate change involves many complex dynamic and diverse processes with nonlinear interactions that pose huge challenges to modellers [6].

A wireless network such as the one discussed herein will enable the vineyard management to decide on the kind of measures (i.e., sprinkler system, gas/ turbine heaters/ helicopter and a schedule) required for example, to prevent frost damage to the crops. Further details on the frost prediction and wireless sensor network issues could be found at [10]. Modelling macro-micro climate change effects has begun in this project with the use of WSN data obtained from vineyards located in three continents. This enables observations of the variability in global climate change across the continents using prediction model values provided by NASA and other institutions.

The models and results will be used in a comparative analysis on climate change effects on viticulture, especially its variability and its influence on grapevine “cultivars” or varieties and wine quality, such as aroma, colour and mouth feel as climate change effects on viticulture is described to be dramatic and varying across the globe significantly.

4 THE WSN ARCHITECTURE

The proposed WSN system [2] consists of sensor nodes located in critical locations within vineyards for collecting weather, atmospheric and environmental data as well as plant related data such as leaf wetness and sap flow. Figure 4 shows the system architecture consists of three layers namely, mote layer, server layer and application layer.




Figure 4: A schematic view of WSN architecture

Mote layer: This layer consists of all the wireless sensor nodes and a Base Station (BS). Each node has one or more sensors plugged into the hardware device with a transmitter, power supply (usually a small battery) and microcontroller. The nodes are distributed over an area of interest uniquely arranged as required provided the distance between the sensor devices does not exceed the maximum communication range. Therefore, energy optimized routing becomes essential. Data transmission from sensor nodes to the BS depends on application maybe continuous, event driven, query-driven or hybrid. In continuous approach, data is transmitted to the BS periodically according to predetermined intervals. In query and event driven models, data is transmitted when an event takes place or query is generated from the BS. The Hybrid model uses combinations of these approaches to transmit data from sensor nodes to the BS. Various types of routing protocols such as data-centric, hierarchical and location-based protocols are available [6].

Server layer: Data are sent to the data server from the BS through the internet. Two main tasks performed by data server are to:

- 1) obtain and process data from the BS.
- 2) populate database with WSN data and enabling the application layer to access WSN data.

The server layer also deals with on-time data delivery from the BS and generates alarm when an undesirable event takes place.

Application layer: This layer allows users of the system have remote access to WSN data using web browsers. This provides a powerful tool to visualize real-time WSN data and compare data from various nodes. In addition, the BS can be accessed remotely to modify sensor nodes’ configurations.

5 IMPLEMENTATIONS

To achieve objectives described in Eno-Humanas project [1], a WSN prototype was designed and developed for gathering and monitoring environmental data within vineyards. Both hardware and software were designed and built by researchers at Auckland University Technology (AUT).

5.1 Hardware design

WSN hardware can be divided into two main components namely, coordinator node and router node. For both node types one wireless plug-in module is used. The plug-in board is based on CC2431 wireless micro controller. This microcontroller has on chip 2.4GH wireless radio, 128KB in-system programmable flash and hardware based location awareness. Figure 5 below shows the controller unit, which has Printed Circuit Board (PCB) antenna, on-board power and all the I/O ports are connected to two headers.

Sensor nodes used are small low-cost, single-chip device which have their own built-in micro-processor. These nodes can automatically set up an ad hoc wireless communication network with other motes that are within range (up to 100 m). These nodes are small battery powered devices allowing wireless communication capabilities (Refer to Figure 7). Sensors may be configured to send data to the BS in following manners:

Periodically: sensor node sends data according to predetermined intervals.

Event-based: sensor node starts sending data when a specific requirement is met. For instance, there may

be no need for sending data to the BS unless temperature drops below a certain threshold in frost prediction task.

In this study, data are sent to the BS periodically at varies intervals for different sensors. In addition, an ultra light minicomputer with wireless communication capability is used to receive sensed data from sensor nodes. The BS also can be accessed remotely via internet to upload sensor data to the main off-site server.

In order to create a wireless sensor network, ZigBee wireless protocol based CC2431 micro controller is used. This microcontroller has on-chip 2.4GH wireless radio, 128KB in-system programmable flash, hardware based location awareness and much more other useful features. Figure 5 shows the main controller unit, which has Printed Circuit Board (PCB) antenna, on-board power and all the I/O ports are connected to two headers.




Figure 5: Wireless controller plug-in board

5.1.1 USB coordinator node

The coordinator is used to collect the information from the network and transfer the data into the computer via USB port. This board is powered through the computer USB power supply. Once data is received from the network, sensor data is transferred to the RS232 to USB converter. The converter sends the RS232 data to the computer via USB port (Refer to Figure 6).




Figure 6: Base station with coordinator node

Every time new node is connected to the network, coordinator will identify the node and issues the new node with an IEEE address. During the communication between these nodes, coordinator identifies each node with the IEEE address.

5.1.2 Sensing router node

The router node can be divided into two main modules; sensor module and wireless module as shown in Figure 7. The CC2431 controller module collects the sensor readings via Analogue to Digital Converters (ADC) and

transmits the data to the USB coordinator dongle via mesh network. The sensor unit is powered via 4x1.5V AA battery pack. Due to limitation on flash memory, raw data is sent to the BS without any processing. The BS converts ADC reading into appropriate values by using conversion factors.

The sensor module comprises six different environmental sensors namely pressure, leaf wetness, sunlight, humidity, temperature, and soil moisture.




Figure 7: Router node with environmental sensors

5.2 Software Implementation

A graphical user interface designed to manage sensor nodes communication, data logging and server upload. It also provides a dashboard for displaying sensor readings and derived parameters such as dew point (Refer to Figure 8). Historical data displayed in grid and graphical formats which can be customized to provide better visualisation of logged data.



Figure 8: Graphical user interface for displaying WSN data

Sensor modules can be added as required. Data that are collected at predetermined intervals or at specific times are transmitted through the network elements to a main computer or controller

4.3 Online web monitoring

A web application was developed enabling users to interactively access the WSN data over the internet. It allows live monitoring and visualization of climate, atmosphere, plants and soil data from each vineyard (Refer to Figure 9).



Figure 9: Live web monitoring of Awarua vineyard's environmental data

6 DISCUSSION

Power and resource management are two main issues that needed to be taken into consideration. In this implementation, one way to achieve this is to incorporate a low level controller that facilitates for more sensors and computational power. The higher level controller will handle the ZigBee mesh networking, and request the data from the low level controller. The main task for the low level controller is to collect data from the sensors and process the data for suitable format and send it the high level controller on demand. The low level controller will be responsible for the sensor power management.

7 CONCLUSIONS AND FUTURE WORK

The paper investigated the recent advances in remote wireless sensor devices, and how WSN of these devices could be combined with the internet and used in vineyard operations, such as management decision making, by monitoring weather, atmospheric, environmental conditions and plant physiology, and also for online display of climate information at larger scales, such as regionally within a state and cities in the Asia Pacific Region.

In order to improve the data management TinyOS real time system can be deployed. TinyOS tiny data base system can be used to temporary store the data within the network. Power management can be improved by introducing renewable energy sources, such as solar and wind energy to charge the batteries. The down side to this approach is that the cost of the node will increase signifi-

cantly. But long term sustainability and life-time without any servicing will increase considerably.

8 ACKNOWLEDGMENTS

We would like to acknowledge Dr John Collins, School of Engineering at AUT for his technical assistance with the embedded software design.

9 REFERENCES

- [1] Sallis, P.J., Shanmuganathan, S., Pavesi, and L., and Jarur, M., "A system architecture for collaborative environmental modelling research", International Symposium on Collaborative Technologies and Systems (CTS 2008), Eds., Waleed W. Samari and William McQuay, A publication of the IEEE, New Jersey, USA. ISBN: 978-1-4244-2248-7, Irvine, California, May 19-23 2008 pp 39-47.
- [2] Ghobakhlou, A., Shanmuganathan, S., and Sallis, P., "Wireless sensor networks for climate data management systems.", In B. Anderssen et al. (eds) /18th IMACS World Congress - MODSIM09 International Congress on Modelling and Simulation/, 13-17 July 2009, Cairns, Australia. ISBN: 978-0-9758400-7-8. pp. 959-965.
- [3] Camilli, A., Cugnasca, C. E., Saraiva, A. M., Hira-kawa, A. R., and Corrêa, P. L. From wireless sensors to field mapping: Anatomy of an application for precision agriculture Computers and Electronics in Agriculture Volume 58 , Issue 1 (August 2007) Pages 25-36, ISSN:0168-1699.
- [4] Xianghui Cao, Jiming Chen, Yan Zhang, and Youxian Sun (2008) Development of an integrated wireless sensor network micro-environmental monitoring system. ISA Transactions, Volume 47, Issue 3, July 2008, pp 247-255.
- [5] Garcia R.R. Understanding the Zigbee stack. pp. 1-2, 02/01/2006 (Freescale Semiconductor Inc.).
- [6] Ankur Suri, S.S. Iyengar, and Eungchun Cho Ecoinformatics using wireless sensor networks: An Overview. Ecological Informatics, Volume 1, Issue 3, November 2006, pp 287-293.
- [7] Jose D., Gutierrez A. IEE Std. 820.15.4 Enabling Pervasive Wireless Sensor Networks. Innovation Centre, pp.2-54,2005.
- [8] Egan D. The emergence of ZigBee in building automation and industrial control. Computing & Control Engineering Journal, pp. 14-19, April-May 2005 16(2).
- [9] Wheeler A. ZigBee Wireless Networks for Industrial Systems - White Paper. 2006 [cited 04/09/2007]; Available from: <http://www.microcontroller.com/Embedded.asp?did=149>
- [10] Sallis, P., Jarur, M., and Trujillo, M. Frost prediction characteristics and classification using computational neural networks. In Australian Journal of Intelligent Information Processing Systems (AJIIPS) volume 10.1, 2008 (ISSN 1321-2133) pp50-58. Also published in M. Kppen et al. (Eds.): ICONIP 2008, Part I, LNCS 5506, 2009. Springer-Verlag Berlin Heidelberg 2009. pp. 1211-1220.

Wireless Soil Moisture Sensor for Vineyard Soil Monitoring

T.A.S.A Perera¹, J.D. Collins²

School of Engineering
Auckland University of Technology
Private Bag 92006, Auckland 1142, NEW ZEALAND.

¹Email: achala.perera@aut.ac.nz

²Email: john.collins@aut.ac.nz

Abstract— The soil moisture level is one of the critical aspects which controls the quality of the grapes grown in vineyards. This research investigates the development of a low cost soil moisture sensor. The two probe soil moisture sensor uses the basic principle of Hartley oscillator frequency shift due to the varying dielectric constant of the soil according to the soil Volumetric Water Content (VWC). When the soil VWC increases, the dielectric constant of soil also increases. This variation can be captured and transmitted via a Wireless Sensor Network (WSN) to a base station for further data analysis. In order to calculate the soil moisture levels accurately, the soil temperature is also measured at the same time. Texas Instrument's (TI) ZigBee protocol based wireless micro controller (CC2431) provides an affordable mesh networking facility and power management for the soil moisture data communication.

Keywords— soil moisture sensor, volumetric water content, wireless sensor network

I. INTRODUCTION

Increasing population and deforestation is causing tremendous pressure on the world's drinking water resources. Water is one of the most precious commodities in the world and it needs to be preserved for future generations. One of the world's main water consuming industries is irrigation. It has been estimated that the world's irrigation efficiency is less than 40 percent [1]. In order to improve irrigation efficiency, it is important to monitor the soil Volumetric Water Content (VWC) and avoid over watering plants.

The main aim of the soil moisture sensor is to optimise the water usage of the wine industry in New Zealand and Chile. By maintaining the optimum water level for grapevines, the grape harvest is improved and water usage is reduced. The sensor developed in this research will measure the VWC at three different depths: 5m, 3m and surface level. Measuring the VWC at different depths gives a better understanding of the VWC up to the grapevine root levels. To understand the VWC at the micro-level, sensors are to be installed in a rectangular pattern at 10m intervals in the vineyard.

One popular water sensing approach is the Time Domain Reflectometry (TDR) method [2]. TDR is highly accurate, but it is very costly to implement on a large scale. Therefore

a capacitance based, fast reacting soil sensor method is used in this research.

Wireless Sensor Networking (WSN) provides a mesh network to deliver the data to a base station computer. WSN has been the subject of research in various domains over the past few years and has been deployed in numerous application areas. WSN is seen as one of the most promising contemporary technologies for bridging the physical and virtual world, thus enabling them to interact. A WSN is composed of a number of sensor nodes, which are usually deployed in a region to observe particular phenomena in a geospatial domain. Sensor nodes are small stand-alone embedded devices that are designed to perform specified computations and to send and receive data. A WSN system includes a number of sensors gathering data from the local environment that is being monitored. WSN's have been employed in both military and civilian applications, such as target tracking, habitat monitoring, environmental contaminant detections and precision agriculture [3][4].

II. SOIL MOISTURE SENSOR DESIGN AND CONSTRUCTION

The sensor system design can be divided into two main parts. These are the soil moisture sensor and the wireless nodes (router and coordinator). Each soil moisture sensor is attached to a router node, and all the router nodes are wirelessly connected to the coordinator node. The coordinator node is connected to the computer via a USB interface (refer to Figure 1). A network may contain thousands of nodes. All the measured data can be stored in the computer for each region of the vineyard and can be used to correlate the wine quality with the soil VWC.




Figure 1: System Block Diagram

A. Soil Moisture Sensor




Figure 2: Capacitive Probe

The capacitance of the probe in soil is

$$C_s = \epsilon_r C_0 \quad (1)$$

where:

- C_0 is the capacitance in air
- C_s is the capacitance in soil
- ϵ_r is the relative permittivity of the soil

The basic principle of operation is to make the capacitive probe part of a Hartley oscillator tank capacitor and to measure the change in oscillator frequency caused by introducing water.

The dielectric constant of soil mineralogical materials varies from 2 to 14, and water has a dielectric constant of approximately 80 [5]. Therefore the dielectric constant of soil is a reasonable indicator of the VWC of soil. Adding water to the soil changes the dielectric constant, resulting in an increase in capacitance of the soil and this causes a decrease in the frequency of the oscillator (refer to formula 2).

$$f = \frac{1}{2\pi\sqrt{LC}} \quad (2)$$

The operating frequency of the oscillator is one of the critical factors which needs to be considered before designing the oscillator. The complex permittivity of the soil (ϵ) can be defined as:

$$\epsilon = \epsilon' + j\epsilon'' \quad (3)$$

The real part (ϵ') and the imaginary part (ϵ'') respectively represent the dielectric permittivity and the conductivity of the soil.

Equation (3) can be written as:

$$\epsilon = k\epsilon_0 - \frac{j\sigma}{\omega} \quad (4)$$

where k is the real part of the dielectric constant, ϵ_0 is the permittivity of free space (8.85×10^{-12} F/m), σ represents the conductivity of soil (mho/m), and ω is the angular frequency (rad/s) at which the measurement is made [6].

It has been found that for frequencies higher than 10MHz, the imaginary part of the complex permittivity has an insignificant effect [7, 8, and 9].

The Hartley oscillator is designed to operate at 16MHz when the surrounding soil is fully dry. The sensor has a full frequency range of 1MHz variance from fully dry to 100% wet. Therefore the effective range of the sensor is 16MHz down to 15MHz as the soil moisture content changes from 0% to 100%.

The output of the oscillator is fed into a frequency down converter. The down converter multiplies the input signal with a 16MHz reference signal from the microcontroller crystal, producing a signal with two frequency components. The higher component is filtered out using a low pass filter and the lower frequency component is fed into a comparator and converted into a square wave. The square wave is then fed into the Input Capture interrupt input on the ATmega16 microcontroller. From the timing measurement on this signal, the microcontroller calculates the input frequency and this is then sent to the wireless microcontroller for transmitting (refer to Figure 3).




Figure 3: Soil Moisture Sensor Block Diagram

The soil moisture sensor needs to be calibrated. This is achieved by adjusting the variable inductor core to obtain an oscillator frequency of 16MHz in a dry soil sample. The microcontroller soil calibration is performed by comparing with an accurate commercial soil moisture sensor (Delta-T Devices HH2) (refer to Figure 4). Both the Delta-T HH2 soil moisture sensor and the capacitive soil moisture sensor were placed into same soil sample and the frequency from the capacitive sensor was recorded against the measured soil moisture level from Delta-T HH2 soil moisture sensor. Once data is recorded, the relationship between the frequency output and soil moisture level can be formulated and programmed into the ATmega 16 micro controller.



Figure 4: Delta-T HH2 Commercial Soil Moisture Sensor

B. Wireless Sensor Network

In this application, a mesh topology is used to increase the reliability and scalability of the network. This system is an ad hoc multi-hop system. The mesh network is based on the ZigBee protocol, which means each router node can be used to send and receive data. In this manner, the network consists of multiple redundant communication paths, which can be used in the event of a node failure. Therefore, data will always reach its destination reliably, via the intermediate nodes (refer to Figure 5) [10].

The ZigBee mesh network has three important properties: it is self-configuring, self-healing and scalable. A ZigBee mesh network configures itself automatically. The network identifies new nodes and automatically includes them in the network. According to the ZigBee mesh network standard, it can support up to 65,536 network nodes (clients) [10].




Figure 5: Mesh Network Topology Used in the Soil Moisture Experiment Setup

In order to create a wireless sensor network, the ZigBee wireless protocol based CC2431 micro controller is used. The key features of this microcontroller include: on-chip 2.4GHz wireless radio, 128KB in-system programmable flash and a hardware based location awareness engine. The main controller unit has a printed circuit board (PCB) antenna and onboard power regulation. All the input/output ports are connected to two headers.

1) Router Node

The router node can be divided into a sensor module and a wireless module. The ATmega16 microcontroller collects the sensor reading via its Input Capture interrupt and sends the

data to the CC2431 wireless microcontroller via RS232. The CC2431 wirelessly transmits the data to the USB coordinator dongle via the ZigBee mesh network. The sensor unit is powered from a 12V battery pack. Due to flash memory limitations, raw data is sent to the base station coordinator without any processing. The base station converts the frequency reading into appropriate values using conversion factors.




Figure 6: Soil Moisture Sensor

2) Coordinator Node

The coordinator node is used to collect the information from the network and transfer the data to the computer via a USB port. This board is powered through the computer USB power supply. Data received from the network sensors are transferred to an RS232 to USB converter. The converter sends the RS232 data to the computer via the USB port (refer to Figure 7).



Figure 7: Base Station with Coordinator Node

Each time a new node is connected to the network, the coordinator identifies the node and issues it with an IEEE address. When communicating with the nodes, the coordinator identifies each node using its allocated IEEE address.


III. RESULTS

The first prototype soil moisture sensor units have been developed at Auckland University of Technology, New Zealand.

A test mesh network was established with four router nodes and one coordinator node. This was used to test the system communications. Due to the micro scale soil moisture detection typically required, a 10m grid network was used. The maximum communication range of the wireless unit is 100m in an open area. According to the user's requirements, data was stored at 10 second intervals. This data sampling rate can be adjusted according to the user's requirements. The communications test was successful.

An initial experiment has been conducted to test the operation of the soil moisture sensor. Using the capacitive

sensor, the frequency was measured in a soil samples and compared with soil moisture measurements made using the Delta-T HH2 soil moisture sensor. The raw data is plotted on Graph 1.



Graph 1: Soil Moisture Content vs Sensor Frequency

The relationship between soil moisture and capacitance can be formulated as follows. Assume that dry soil consists of solid material with air gaps, and that the proportion of space filled by soil is p . When water is added, the water displaces some air from the gaps. The water content M is the proportion of air that has been displaced. It is assumed that no solid material dissolves in the water. Then the capacitance of wet soil is given by

$$C_s = (pk_s + (1-p)Mk_w + (1-p)(1-M)) C_0 \quad (5)$$

where:

- C_0 is the probe capacitance in air
- C_s is the probe capacitance in wet soil
- k_s is the dielectric constant of the soil matter
- k_w is the dielectric constant of water
- p is the proportion of soil matter
- M is moisture content of the soil

When $M = 0$, the capacitance of dry soil is:

$$C_s = (pk_s + (1-p)) C_0$$

This includes terms due to the soil matter and air gaps.

When $M = 1$, capacitance of saturated soil is:

$$C_s = (pk_s + (1-p)k_w) C_0$$

This includes terms due to the soil matter and gaps completely filled with water.

Overall, (5) can be written as

$$C_s = ((pk_s + (1-p)) + M(1-p)(k_w - 1)) C_0 \quad (6)$$

The variation of C_s is linear in M .


$$C_s = (a + bM) C_0 \quad (6)$$

where a and b are calibration constants.

From equation 2 we can derive the relationship between the measured frequency f and soil moisture content as

$$4\pi^2 f^2 LC_s = 1 \quad (7)$$

We expect a graph of $1/f^2$ versus M to be a straight line. This is plotted in Graph 2 for the measured results.



Graph 2: 1/Frequency² vs Soil Moisture Content

The relationship does not fit the model well and further investigation of these results is required. It has been found that the commercial soil moisture meter has a limited range (10% to 50%). Therefore in future testing, the experiment method needs to be changed. The proposed testing method is to weigh the soil sample and add a known weight of water so the actual moisture content can be reliably estimated over a wide range of values. The resulting data will be used to investigate the relationship between soil and capacitance in more detail.

IV. CONCLUSION AND FUTURE WORK

This paper describes the design and construction of a wireless soil moisture sensor. The design is based on the Hartley oscillator frequency variation caused by the soil moisture content.

The total cost of the sensor unit is approximately US\$70. The cost of the electronics is approximately US\$30 but the most expensive part is the aluminium case used for underground shielding, which costs approximately US\$40 to construct.

Power management could be improved by introducing renewable energy sources, such as solar or wind energy to charge the batteries, but this may increase the total cost of the nodes. However use of renewable energy sources would support long term sustainability and reduce maintenance costs.

ACKNOWLEDGMENT

We would like to acknowledge Brett Holden, AUT School of Engineering, for his technical assistance with the design of soil sensor hardware.

REFERENCES

- [1] Agronomy journal, vol. 93, March-April 2001.J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.
- [2] F. N. Dalton and M. Van Genuchten. “The Time Domain Reflectometry Method for Measuring Soil Water Content and Salinity,” *Geoderma* 38, 1986, pp. 237-250.
- [3] Camilli, A., Cugnasca, C. E., Saraiva, A. M., Hira-kawa, A. R., and Corrêa, P. L. From wireless sensors to field mapping: Anatomy of an application for precision agriculture *Computers and Electronics in Agri-culture* Volume 58, Issue 1 (August 2007) Pages 25-36, ISSN:0168-1699.
- [4] Xianghui Cao, Jiming Chen, Yan Zhang, and Youxian Sun (2008) Development of an integrated wireless sensor network micro-environmental monitoring system. *ISA Transactions*, Volume 47, Issue 3, July 2008, pp 247-255.
- [5] Jeffrey E. Campbell, “Dielectric Properties and Influence of Conductivity in Soils at One to Fifty Megahertz,” *Soil Sci. Soc. Am. J.*, Vol. 54, 1990
- [6] E. T. Selig and S. Mansukhani, "Relationship of soil moisture to dielectric properties," *J. G. T., ASCE*, vol. 101, 1975, pp. 760-770.
- [7] P. Boyapati and K. Shenai “A Novel Wireless Soil Moisture Sensor” *International Conference On Energy And Environment* March 19-21, 2009 ISSN: 2070-3740
- [8] Peter M.Johnson, David V.Thiel, and Daniel A. James, “Contributions to the measured capacitance by the dielectric properties of water insulated electrode soil moisture sensors.”*Proc. of IEEE, sensors*, vol. 1, Nov., 2002, pp. 495–498.
- [9] Darold Wobschall, “A Frequency Shift Dielectric Soil Moisture Sensor,”*IEEE Trans, Geoscience Electronics*, Volume 16, Issue 2, April 1978, pp.112 – 118.
- [10] Wheeler A. ZigBee Wireless Networks for Industrial Systems - White Paper. 2006 [cited 04/09/2007]; Available from:<http://www.microcontroller.com/Embedded.asp?did=149>

The Slot Car Stig: Performance and Consistency of a Slot Car Driven by a Heuristic Algorithm in an Embedded Microcontroller

Steven Kane, Jonathan Scott

The University of Waikato

Gate 1 Knighton Road Hamilton 3240

srk13@waikato.ac.nz

jbs@waikato.ac.nz

Abstract:

We present theory and measured performance of an autonomous slot car driven by a heuristic algorithm on a typical track. The hardware consists of a PIC 8-bit single-chip microcontroller with various sensors driving a standard permanent-magnet (PM) brushed dc (BDC) motor in a mechanically-standard Scalextric platform. We present some interesting results concerning the relative difficulty of apparently-balanced lanes on a track. The car achieves optimum lap times with high consistency. Measured performance agrees with theoretical expectation. The consistency of performance allows the impact of experimental changes to be reliably assessed.

1 Introduction

Slot cars are popular all over the world. One might think that because one does not have to steer a slot car they do not represent a test of skill. However this is not true, as evidenced by many national championships, student competitions [1], many popular accounts of races, and a selection of electronic designs centred around the sport [2], [3], [4]. Recently digitally controlled versions that can change lanes have appeared. There are at least three manufacturers of such systems [5], [6], [7].

The authors wanted to find out whether a simple 8-bit microcontroller could drive a slot car better than a person, and if that would allow them to use slot cars to test real automotive innovations. This manuscript reports the outcome of that study.

2 Track Theory

Most slot car enthusiasts realise that a track needs to be carefully designed in order not to give an advantage to a driver in one particular lane. For example, a simple oval circuit gives a considerable advantage to the driver in the inside lane. Consider the track diagram shown in figure 1. At first glance one might expect this to represent a fair layout because there are an equal number of left and right hand turns of the same radii presented to the driver in each lane. This does not prove to be the case.

Let us assume that there is an optimum speed at which a standard car will traverse any given corner. It is possible to construct a diagram, once the optimum corner speeds are known, such as that in figure 2. This figure shows the maximum possible speed as a function of distance

around the track. The square waveform assumes infinite acceleration and deceleration, the more “triangular” waveform accounts for finite acceleration and deceleration. The heuristic algorithm programmed into the car delivers the thrust required to approach the triangular trajectory as closely as possible.




Figure 1 Scale drawing of an example slot car layout with equal left and right turns in each lane.

Figure 3 differs from figure 2 only in the lane of the track. The interesting observation is that mathematically we predict that a car travelling at optimum achievable speed will traverse one lane more quickly than the other. It is not immediately obvious but this arises because of the order in which each section of track is presented to the driver in each lane is different. Will our robot car perceive this small difference?



Figure 2 Plot of ideal (joined dots) and achievable (light continuous lines) velocity of a car as a function of distance around the left lane of the track shown in figure 1.




Figure 3 Plot of ideal and achievable velocity of a car as a function of distance around the right lane of the track shown in figure 1.

3 Vehicle Parameters

In order to program the car we need to measure optimum cornering speeds, maximum acceleration and braking, etc. These are parameters of the vehicle and track system. The parameter measurement methods are described in this section. The values are stored in non-volatile memory in the microcontroller. Eventually these parameters can be fine-tuned by the microcontroller, as a driver would tune his driving to suit his vehicle and a particular track.

Loops of identical curves, such as the example shown in figure 4, were used to determine the optimum corner speed for each radius. The car was driven around each loop at increasing thrust and the speed measured. As expected there was an optimum speed, above which the car lost traction and increased power resulted in reduced speed. These speeds were recorded. The results for the prototype car's maximum speeds are given in Table 1. Video analysis of the car moving along a straight test track for varied levels of thrust allowed determination of the maximum acceleration and deceleration that could be achieved. Figure 5 plots the car's speed as a function of time, and also shows a linear fit to the measured data. Analysis of the data for the various thrust levels gave the optimum values.



Figure 4 Example set up for measuring optimum cornering speed on a particular track section.




Figure 5 Optimum deceleration performance measured on the prototype car through frame-by-frame analysis of video recording, and straight line fit to the measured data.

Track Pieces and Lengths	Velocity cm/s
Straight section about 70 cm long	400
Straight section about 35 cm long	400
Straight section about 18 cm long	400
Straight section about 8 cm long	400
The Starting Grid, about 35 cm long	400
Half Std curve, turning left, LH lane	211
Half Std curve, turning right, RH lane	208
Half Std curve, turning left, RH lane	208
Half Std curve, turning right, LH lane	186
Hairpin curve, turning left, LH lane	167
Hairpin curve, turning right, RH lane	148
Hairpin curve, turning left, RH lane	132
Hairpin curve, turning right, LH lane	120
Std curve, turning left, LH lane	211
Std curve, turning right, RH lane	208
Std curve, turning left, RH lane	208
Std curve, turning right, LH lane	186
Outside curve, turning left, LH lane	250
Outside curve, turning right, RH lane	219
Outside curve, turning left, RH lane	256
Outside curve, turning right, LH lane	219

Table 1: The maximum velocity measured for each track piece.

4 Hardware

Figure 6 shows a photograph of the prototype slot car. Figure 7 shows the circuit diagram.

We built a barcode reader using an opto-sensor IC. We defined a custom barcode encoding of which an example appears in figure 8. Each barcode has a stop bit and a start bit and 5 data bits in between. The known spacing from start to stop bit allowed the car to verify its velocity as it passed each barcode. The 5 bits encoded the length, radius, and handedness of each track section by means of a lookup table in electrically-erasable non-volatile RAM. The motor was driven with a PWM pulse train. The electronics were arranged so the motor saw a low impedance in both the mark and space parts of the waveform. This was accomplished using a half bridge constructed with MOSFETs.

The track has power constantly applied. When first put on the track, the car drives slowly around until it has memorised and verified the sequence of tracks. Once memorised successfully, it puts the track sequence into non-volatile memory, proceeds to the starting straight and halts. When the photo-sensor detects the flash of a flash gun the car starts racing.

The intention was to have a learning mode, activated by the press of a button rather than the flash of light. In this mode the car will gradually increase speed on each individual track section, one at a time. When the optimum speed on each section of track is identified, the working parameters in memory are updated. This will constitute an artificially intelligent tuning process. Unfortunately time did not permit us to get to this part of the project.



Figure 6 Photograph of the prototype slot car showing the printed circuit board. The barcode scanner is situated adjacent to the right-front wheel. A second photo sensor looks upward. The microprocessor is the 14-pin IC near the center. There is an ICSP connector for programming.




Figure 7 Complete circuit diagram of the car based around a microchip PIC16F684. Apart from a power supply regulator there are no other ICs in the entire design.




Figure 8 Example of a track identifying barcode. This code shows 10011 with start and stop bits equal to 1 giving a complete sequence of 1100111. Ones are white, and are always spaced apart by a black bar, so the 5 white sections in the diagram correspond to the 5 ones in the sequence 1100111. A black region appears before the start and following the stop bit. All codes are equally long, allowing a passing car to simultaneously confirm its speed.

5 Results

Figure 9 shows a cumulative frequency chart as a function of lap time measured on the track of figure 1 in left and right hand lanes for both the prototype car and a skilled volunteer known to the authors. It is immediately clear from the data in the figure that the car is faster than the most skilled human available to us. This did not surprise us.

The next observation is that the car was much more consistent in its performance than any person we tested. This serves to build our confidence that the robot slot car platform will be able to resolve the impact of any small experimental change. Indeed we were reliably able to check that tyres were properly “run in” or sense if they were worn out, or if pickups or track joints were not in order. To a user watching lap times, it feels as if you have a superhuman benchmarking everything, much like “The Stig” of Top Gear fame. Hence the title of this manuscript.

The car could readily distinguish the difference between the left and right lanes of our test track, identifying one as “easier”, precisely as we had predicted. One lane recorded a track time of 3.0 s and the other 3.1 s, with a standard deviation in each case of around 0.05 seconds. This did (pleasantly) surprise us.




Figure 9: Plot of the frequency of lap times achieved by car and expert on left and right lanes of the test track. The consistency of the slot car compared with a skilled human clearly reveals the difference between the lanes.

4 CONCLUSIONS

When controlled electronically a slot car can perform very consistently. Heuristic driving algorithms that are readily programmed into truly low-cost, US\$1, microcontrollers can deliver performance better than a practiced human. The platform allows small design innovations to be assessed with great sensitivity because of the high consistency and optimal performance. The prototype could quantify the relative difficulty of lanes in a track that might otherwise have been considered equal and balanced.

The application of scientific rigour and microcontroller precision to slot car racing has the potential to teach many lessons. We do not think it unreasonable to use the slot car model to test principles and ideas for use with real cars.

5 REFERENCES

- [1] M. Brejl, and J. Necesany, "Student's Contest: Self-Driven Slot Car Racing", International Multiconference on Computer Science and Information Technology, 20-22 Oct., 2008, pp:589 – 592.
- [2] Beverley, P., Keeping track of slot car racers, Electronics & Computing Monthly, v 5, n 2, 34-6, Feb. 1985
- [3] Scott, J., Slot car controller, Electronics Today International, v 11, n 5, p 79-82, May 1982
- [4] McKerrow, P.J., Microcomputers, slot cars, and education, IEEE Micro, v 3, n 1, 62-5, Feb. 1983
- [5] The Incredible Car Racing Game, Scalextric, 2009, <http://www.scalextric.com>, retrieved 12/10/09
- [6] The Digital System, SCX, 2009 <http://www.scxdigitalsystem.com>, retrieved 5/11/09
- [7] Model Road Racing Company Slot cars, MRRC, 2009, <http://www.slotmonsters.com>, retrieved 5/11/09
- [8] Microchip Corporation, 2009, <http://www.microchip.com>, retrieved 14/10/2009



Steven Kane expects to graduate at the end of 2009 with a Bachelor's degree in Engineering from The University of Waikato.

In 2008 he worked on a project for Novel Ways where he used electronics built into a sign to capture people's attention. He is now working in R&D at Gallaghers,

researching ways to improve the electric fence. His interests include mechatronics, power electronics and computer aided engineering in finite element analyses.



Jonathan Scott (S'82-M'97-SM'98)

was born in Brisbane, Australia in 1956. He received the B.Sc., B.E., M.Eng.Sc., and Ph.D. degrees from The University of Sydney, Sydney, Australia, in 1977, 1979, 1983 and 1997 respectively.


He is currently the Foundation Professor of Electronics at the University of Waikato,

Hamilton, New Zealand. From 1998 to 2006, he was with the Microwave Technology Center, Hewlett-Packard (now Agilent Technologies), Santa Rosa, CA, where he was responsible for advanced measurement systems. In 1997 and 1998, he was Chief Engineer with RF Technology, Sydney, N.S.W., Australia. Prior to 1997, he was with the Department of Electrical Engineering, University of Sydney. He is a Professorial Fellow of Macquarie University. He has served on committees of the Standards Association of Australia and the NRC Review Panel for the National Institute of Standards and Technology (NIST). He has authored over 80 refereed publications, several book chapters, and a textbook. He holds nine patents.

Professor Scott's interests include RF and microwave systems, compound semiconductor devices, electronic circuit linearity and nonlinear measurement, sensors and instrumentation, power electronics, and mechatronics.

Index of Authors

- Bainbridge-Smith, A., 87
Benton, C.D., 1
Bonnett, B.C., 75
- Carnegie, D.A., 1, 125, 153, 159
Cavers, J., 61
Chalmers, A.N., 19
Collins, J.D., 37, 171
Cree, M.J., 95
Currie, J., 7
- Diegel, O., 165
Dmochowski, P.A., 143
Dorrington, A.A., 95
Drayton, B.M.M., 153
Dudley Ward, N., 119
- Fisk, R.P., 43
Fletcher, V.A., 125
Fox, C., 49, 55, 119
- Gaynor, P., 1
- Ghobakhlou, A., 165
- Hasan, S.M.R., 25, 43, 131, 137, 149
Hayes, M.P., 61, 75, 87
Hollitt, C.P., 107, 113
- Johnston-Hollitt, M., 107, 113
Jones, H.W.H., 143
Jones, N., 107
- Kamp, W., 87
Kane, S., 177
Kelly, C.D.B., 95
Khurramand, M., 149
Kitaev, V., 81, 101, 107
- Li, J., 131
Lohani, P., 25
- Ma, E., 49
McClure, S.H., 95
Molteno, T.C.A., 13, 67, 71, 107
- Nalli, S., 31
- Opie, A., 61
- Payne, A.D., 95
Perera, T.A.S.A., 81, 165, 171
Prince-Pike, A., 37
- Sallis, P., 165
Scott, J., 177
Sergieiev, O., 101
Soltic, S., 19
Soudlenkov, G., 81
Suggate, P., 67
- Teal, P.D., 143
- Ward, R., 13
Williams, P.J., 71
Wilson, D.I., 7
Winch, S.J., 159
- Zandi, S., 165



Electronics Research
Department of Physics
University of Otago
elec.otago.ac.nz

