ENZ CON 2016

PROCEEDINGS

OF THE 2016 ELECTRONICS NEW ZEALAND CONFERENCE

NOVEMBER 17 & 18, 2016
VICTORIA UNIVERSITY OF WELLINGTON
WELLINGTON, NEW ZEALAND
WWW.ENZCON.ORG.NZ

VICTORIA UNIVERSITY WELLINGTON
TE WHARE WĀNANGA O TE ŪPOKO O TE IKA A MĀUI

SPONSORS

MECHADEPT

RF TEST SOLUTIONS

PUBLISHED BY: Electronics New Zealand Inc. (ENZCon Inc)

ADDRESS: **ENZCon**
School of Engineering
University of Waikato
Private Bag 3105
Hamilton, New Zealand 340
http://enzcon.org.nz

TITLE: Proceedings of the 2016 Electronics New Zealand Conference

ISBN: 978-0-473-37509-6

FORMAT: PDF

PUBLICATION DATE: 11/2016

ORGANISING COMMITEE: Jim Murphy (jim.murphy@vuw.ac.nz), Dale A. Carnegie

REVIEWERS:
Dale A. Carnegie, Michael Hayes, Jason Long, Rick Millane,
Timothy Molteno, Jim Murphy, Jonathan Scott

CREDITS: ENZCon 2016 Logo, cover, and proceedings designed by Jim Murphy
Cover image courtesy Raster Noton (http://www.raster-noton.net/shop/grid-index)

CONFERENCE SPONSORS:

http://www.rftest.co.nz/

http://www.mechadept.com/

http://www.vuw.ac.nz

FOREWARD

ENZCon 2016 represents the 22nd iteration of the Electronics New Zealand Conference. I'm very excited for the opportunity to convene the conference here at Victoria University of Wellington.

ENZCon occupies a crucial niche in New Zealand electronics engineering conferences: its regional nature allows students, academics, and technicians to meet up with their New Zealand colleagues on a yearly basis, sharing ideas, giving and receiving constructive feedback in a relatively low-key environment, and sharing with each other the current state of Electronics Engineering education and research in New Zealand.

As an early-career academic, I very much value the feedback that my work has received at past ENZCon conferences. Further, the chance to share my work in front of knowledgable audiences has helped greatly with my own presentation skills, allowing me to share my work with others both in New Zealand and internationally. It is in this that I see the largest benefit of ENZCon for students: the chance to emerge as researchers in an environment that will prepare them to take their work to an international stage.

A special thanks to our event sponsors: MechAdept and RF Test Solutions. Both of these Wellington-area companies have been eager and generous in their support of this conference.

Thanks also to our wonderful referee teams, who operated on very short notice to generate high-quality feedback for all of the conference papers. Special thanks to Professor Dale A. Carnegie for passing along years of ENZCon organisation wisdom.

And, lastly, I am grateful and thankful for all of the contributions: this year's conference came together under a relatively tight timeframe, and I know that the main deadlines happened to (very conveniently) coincide with what's often the busiest time of the teaching year. For those of you who dropped what you were doing, put together some high-quality presentations, and submitted them for review, I thank you!

Jim Murphy
ENZCon Convener and Organiser
Victoria University of Wellington
November, 2016

Table of Contents

ENZCon 2016 Papers

ENZCon 2016 Programme (17 & 18 November, 2016) *V 1.6*

Thursday, 17 November	
08:30	Registration opens (Location: Outside of Laby LT118) – *Please note that the registration desk will be open during all tea and lunch breaks, allowing those arriving later to register.*
09:00	Poster and tech demo setup
09:30	**Morning Tea, Poster & Demo Session (Outside of Laby LT118)**
10:30 **Oral Presentation Session 1: Software *(With welcome address by Dale Carnegie)*** (Location: Laby LT118 – Session Chair: Dale Carnegie)	
Elliot Varoy, Nasser Giacaman and Gerard Rowe. **Enhancing Students' Conceptual Understanding Using Visualisation Software and Collaboration**	
Michael Hayes. **Semi-automated electrical circuit drawing with Lcapy**	
Jason Long, Jim Murphy, Dale Carnegie and Ajay Kapur. **Improving the Accuracy of Musical Robots with Programmable Logic**	
Blake Johnston, Bridget Johnson and Ajay Kapur. **A New Framework for Interactive Control of Mechatronic Instruments**	
12:00	**Lunch break**
12:00 – 13:00 **ENZCon Annual General Meeting** *(Location: Laby LT118)*	
13:30 **Oral Presentation Session 2: Sensing** (Location: Laby LT118 – Session Chair: Rick Millane)	
Jason Sun, Rainer Künnemeyer, Andrew McGlone and Nathan Tomer**.** **Investigation of Light Transmission in Healthy and Rotten Onions**	
Michael Hayes and Ian Platt. **A preliminary investigation for an electromagnetic ground water flow measuring device**	
Michael Frampton and Michael Hayes. **A system for non-destructively determining variation in acoustic speed within a harvested tree stem.**	
Malcolm Morrison and Colin Fox. **Vector Green's Functions for Electrodynamics Applications**	
15:00	**Afternoon Tea**
15:40 **Oral Presentation Session 3: Control** (Location: Laby LT118 – Session Chair: Michael Hayes)	
Michael Hutchins and Jonathan Scott. **Speed Control of Small Motors Through the Cuk Converter Topology**	
Andrew Martin and Timothy Molteno. **Real-time control using sequential inference and Bayesian decision theory**	
Faiz Rasool, Syed Adeel Ali and Muhammad Waseem Soomro. **A MATLAB ToolBox for the Bit-Stream Implementation**	
Andy Soundy, Daniel Schumayer and Tim Molteno. **GPS Estimates in the Presence of Correlated Noise**	

19:00	**Conference Dinner**: Mac's Brew Bar http://www.macsbrewbar.co.nz/restaurant.html

(continued on next page)

Friday, 18 November	
09:30	**Morning Tea**
10:00 **Oral Presentation Session 4: Short Talks, Imaging, and Astronomy** (Location: Laby LT118 – Session Chair: Jonathan Scott)	
Sponsor Address by RF Test Solutions Limited *(http://www.rftest.co.nz/)*	
Short Talk: Timothy Molteno. **SkyPi: All-sky camera network for meteorite tracking**	
Short Talk: Arunkumar Jayakumar, Tejas Trivedi, Deepak Sharma, Marilou Rebosura and Latha Karthigaa Murugesan. **Real Time Applications of Electrochemical Impedance Spectroscopy - A Technical Assessment**	
James McPherson and Tim Molteno. **Architecture for the Transient Array Radio Telescope version 3 (TART3)**	
Rick Millane. **Imaging with x-ray free-electron lasers**	
Max Scheel, Tim Molteno and Phill Brown. **Transient Array Radio Telescope: Calibration and Aperture Synthesis**	
11:40	**Lunch break**
12:20 **Oral Presentation Session 5: Signals and Systems** (Location: Laby LT118 – Session Chair: Jim Murphy)	
Daniel Schumayer and Timothy Molteno. **On a representation of feedback shift-registers**	
Itay Ben-Dom and Catherine Watson. **Towards Vocal Health Assessment From The Speech Signal**	
Romain Arnal and Rick Millane. **Effects of non-uniform sampling on phase retrieval**	
Sam Dirks and Jonathan Scott. **Precision Crystal Frequency Reference Disciplined Using Mains Power**	
13:45	**Afternoon Tea**
14:30 **Oral Presentation Session 6: Power & Wireless** (Location: Laby LT118 – Session Chair: Timothy Molteno)	
Rahat Hasan and Jonathan Scott. **Fractional Behaviour of Rechargeable Batteries**	
Michael Frampton and Michael Hayes. **Wireless Treetap: a new electronic device for time-of-flight tree-stiffness measurement.**	
Loren Kersey, Phill Brown, Daniel Schumayer and Timothy Molteno. **Aspects of Quartz Operation in Viscous Fluid**	
T Lee, R B Kennedy, R A Bodnar, Jonathan Scott and W Redman-White. **Wireless Power and Network Synchronisation for Agricultural and Industrial Remote Sensors using Low Voltage CMOS Harvesting and Data Demodulator IC**	
Elliot Baptist, Rainer Künnemeyer, Philip Rowe and Paul Martinsen **Scalable Wireless Network for Remote Sensors: Range Testing**	
16:10	*Short Break*
16:20 – 16:40 **Closing Remarks & Prizes** *(Location: Laby LT118)*	

ENZCon 2016: *ABSTRACTS*

(Abstracts are arranged alphabetically by First Author surname)

Effects of non-uniform sampling on phase retrieval

Romain D. Arnal, Rick P. Millane

Diffraction imaging involves reconstruction of an object from measurements of a diffracted field. In many cases, only the amplitude, but not the phase, of the diffracted field can be measured. Reconstruction of the object then requires recovering the phase, which is referred to as phase retrieval. Uniqueness properties of phase retrieval are thus important. Uniqueness properties of the phase problem are considered in cases where Fourier amplitude data are sampled non-uniformly in a way that arises in imaging of 1D and 2D crystals. Simulations of phase retrieval using an iterative projection algorithm show the influence of sampling non-uniformity in terms of the difficulty of reconstruction and the resulting reconstruction error. The results show that the sampling for 1D, and particularly 2D, crystals leads to more convergence difficulties and increased reconstruction error. The results have implications for *ab initio* phasing in imaging 1D and 2D crystals using x-ray free-electron lasers.

Scalable Wireless Network for Remote Sensors: Range Testing

Elliot Baptist, Rainer Künnemeyer, Philip Rowe, Paul Martinsen

Wireless sensors attached to animals can improve the efficiency of farming, but need a cost-effective way of getting data to the internet. A network using low-cost packet radio modules has been proposed but development was hindered by a lack of information on working range of such radio modules. This report characterises and compares the range of RFM69HW and RFM95W radio modules for indoor and outdoor farm-like situations. It was found that RFM95W radio modules have approximately three times longer range in open air than RFM69HW modules. Viable RFM95W radio links were measured through 9 walls over 52m in a building and over 1.85km in open air. Using the RFM95W over the RFM69HW should reduce the cost of installing such a network in most farms.

Towards Vocal Health Assessment From The Speech Signal

Itay Ben-Dom, Catherine Inez Watson

In speech processing, the glottal pulse contains information about the speaker's voice quality. Through glottal pulse processing, the time-domain parameter Open Quotient can be extracted from a glottal pulse. This paper presents a new Open Quotient criterion, called OQsub50. This paper introduces the implementation of an Open Quotient detection method in R. Phonetic analysis of vowel production was carried out over two New Zealand English (NZE) speech corpora, both containing 11 NZE monophthongs. This study finds the decrease of Open Quotient measurements from young to old age. The trend was observed in both male and female speakers, suggesting the Open Quotient measurements are independent of pitch.

Abstracts, Continued

Precision Crystal Frequency Reference Disciplined Using Mains Power

Sam Dirks, Jonathan Scott

We describe a clock that keeps time and provides a reference frequency with better than 1PPM accuracy, through a quartz crystal automatically calibrated via the power line frequency. The long-term precision of line-frequency-based time is used to correct the crystal-based timekeeping system, a process referred to as disciplining. A microcontroller calibrates the internal reference dynamically with no user assistance. Precision is improved to better than 1ppm as measured against a rubidium standard. The precise signal is useful for calibrating frequency- based instrumentation or keeping time without periodic correction and costs much less than comparable alternatives. The algorithm can be implemented in any mains-powered device using a microcontroller, such as thermostats and lighting timers. The prototype clock uses a Nixie-tube display, requires no transformer, is housed a glass tube, and is offered as an open- source design.

A system for non-destructively determining variation in acoustic velocity within a harvested tree stem.

Michael Frampton, Michael Hayes

Acoustic methods are increasingly being employed by the forestry industry to determine wood quality. Commonly, acoustic velocity is used to determine the stiffness of a harvested tree stem, which is useful for grading and sorting purposes. Existing systems provide a single value for the velocity of a stem. However, a stem does not have a uniform velocity; it varies considerably over its length. This paper investigates two different approaches for non-destructively finding the variation in velocity in a stem. An experiment was performed to capture acoustic signals along the length of a complete stem, which was then sawn into segments and further measurements taken. An analysis was performed on the complete-stem data to find the velocity variation, however the proposed algorithms were unsuccessful. Further development of these algorithms is required.

Wireless Treetap: a new electronic device for time-of-flight tree-stiffness measurement.

Michael Frampton, Michael Hayes

Treetap is a time of flight (ToF) acoustic tree- stiffness measurement system developed at the University of Canterbury. Wireless Treetap (WTT) is the next iteration of this project. The goal of the project is to create a smartphone based system for measuring wood stiffness. The implemented system can be described as a portable wireless sensor network (WSN) of accelerometers with an emphasis on high time accuracy. A prototype of the system has been built and a number of measure- ments performed including: the drift rate of the oscillators, the accuracy of the time synchronisation algorithm, the noise power of the measurement front-end and the battery performance of the device.

Abstracts, Continued

Semi-automated electrical circuit drawing with Lcapy

Michael Hayes

Lcapy is an open-source Python package developed by the author for teaching of linear circuits. The analysis is performed symbolically in the Laplace domain using the Python SymPy package. An important part of Lcapy is generating annotated schematics. This is achieved using a semi-automated algorithm with drawing hints specified for each circuit element. The drawing is performed using CircuiTikZ for text-book quality schematics.

A preliminary investigation of an electromagnetic groundwater flow measuring device

Michael P. Hayes, Ian Platt

This paper describes a preliminary simulation of an electromagnetic system for measuring groundwater flow. The simulation solves the flowmeter equation for a 1m prototype system using a circular coil. A parametric model is used to calculate the magnetic field and Jacobi relaxation is used to determine the induced electric field due to the moving water ions. Peak potential differences of the order of a few tens of nanovolts are predicted for water speeds of 10 mm/hour and magnetic field strengths of 10 mT. The potentials are reduced compared to those for a flowmeter due to additional conductive paths outside the region where the magnetic field is strong.

Speed Control of Small Motors Through the C´uk Converter Topology

Michael Hutchins, Jonathan Scott

TheC´ukConverter offers simultaneous buck-boost operation, but requires careful design owing to its having a fourth-order transfer function and numerous practical design constraints. We exploit a serendipitous overlap between the converter circuit and the equivalent circuit of a dc motor to design a motor controller that can operate with supply voltage that is lower than the motor full-speed requirement. We examine the transfer function when such a topology is used to control the speed of a small motor. We conclude that the approach is relatively straightforward owing to the impact of the motor’s inductance. Measurements agree with theory.

Real Time Applications of Electrochemical Impedance Spectroscopy - A Technical Assessment

Arunkumar Jayakumar, Tejas Trivedi, Deepak Sharma, Marilou Rebosura and Latha Karthigaa Murugesan

DC power systems such as solar cell, Fuel cells and batteries are expected to play a remarkable role for a sustainable future of the society. Electrochemical Impedance Spectroscopy (EIS) or Dielectric spectroscopy is a powerful technique for determining the characteristic of these DC electrical systems where impedance is measured over a range of frequencies to have a greater insight on their system behaviour. The application of EIS is numerous and the present paper assess the system characteristics considering solar cells and proton exchange membrane (PEM) fuel cells into consideration.

Abstracts, Continued

A New Framework for Interactive Control of Mechatronic Instruments

Blake Johnston, Bridget Johnson, Ajay Kapur

Over recent years there has been significant developments in the burgeoning field of mechatronic musical instruments. These developments have seen mechatronic instruments greatly increase in complexity, and concomitantly, their expressive capabilities. However, the development of designing control systems for these instruments has been lacking, meaning that the use of these expressive new instruments has largely being left to the builders of the instruments themselves, or other especially skilled users. This paper proposes a new framework for the control of mechatronic musical instruments that is designed to afford an intuitive interaction for all users. This design allows for more aesthetic exploration of mechatronic instruments as well as expanding the potential demographic of users. This framework explores the potential of new iPad applications that offer high level control parameters and gestures, for the real-time interaction with mechatronic instruments in both performance and installation settings without assuming prior mechanical knowledge of the user. This paper discusses the current state of the field and the need for this new framework. It then demonstrates and explains three case studies that have being developed by the authors as examples of this new framework.

Aspects of Quartz Operation in Viscous Fluid

Loren Kersey, Phill Brown, Daniel Schumayer and Timothy Molteno

A model for the oscillation of an AT-cut quartz crystal submerged in a weakly viscous fluid is developed and used to characterize the expected shifts in resonant frequency. These shifts are complex and predict a change in magnitude of around 0.034% in water. A trend for the quality factor deviation is proposed. An extension to include the oscillation amplitude is in progress. These characterizations are essential to predict the feasible limits of QCM operation in liquid and influence the selection of a frequency counter such that the shifts in frequency can be measured.

Wireless Power and Network Synchronisation for Agricultural and Industrial Remote Sensors using Low Voltage CMOS Harvesting and Data Demodulator IC

T Lee, R B Kennedy, R A Bodnar, Jonathan Scott and W Redman-White

This paper presents a wide area medium frequency loosely coupled magnetic energy harvesting system with power delivery and network synchronisation for remote sensors, intended for agricultural and industrial environments. Intended for situations with poor service access, power is supplied from a source via a large area loop. Receiver nodes may use ferrite cored coils for good efficiency with modest volume. Transmission of low bandwidth network synchronisation data permits very low operational duty cycle with the need for real time clocks or wake up receivers and their associated power drain. As a key enabler for the system, a full custom energy harvester and QPSK data demodulator IC has been designed and fabricated in a commercial 180nm CMOS technology. The IC occupies 0.54mm2 and can deliver 10.3µW at 3V to an external battery or capacitor. With standard MOS device thresholds the rectifier can start from cold with only 250mV peak from the antenna loop, and the battery charge output is delivered with 330mV peak input. Results are presented from laboratory evaluation and from preliminary measurements in the field with a 10m x 10m loop driven at 800kHz.

Abstracts, Continued

Improving the Accuracy of Musical Robots with Programmable Logic Circuits

Jason Long, Jim Murphy, Ajay Kapur, Dale A. Carnegie

In order to provide expressive performances, musi- cal robots must be capable of varying their playing parameters in fine-grained and consistent ways. Unfortunately, due to their often high actuator count and signal processing requirements, conventional purely microcontroller-based methods of control- ling their actuators have proven insufficient at providing the accuracy required for precise and reproducible control. This paper provides an outline of the shortcomings of traditional methods of musical robot control with an example robot which outputs pulse widths with up to 100 microseconds of error, and describes a newly created control system which makes use of a hybrid microcontroller-FPGA design to achieve error rates in the order of single-digit nanoseconds. The system's components and performance is then described, and comparisons are drawn between the new system and traditional methods. This novel system leverages both the sequential logic capabilities of a microcontroller with the massively parallel capabilities of an FPGA to provide a fine-grained, reliable, low latency and high resolution control for musical robots.

Architecture for the Transient Array Radio Telescope version 3 (TART3)

J.P.W. McPherson and T.C.A. Molteno

We propose a new version of the Transient Array Radio Telescope (TART) based on the currently functional TART2. The TART3 will utilise a more powerful central single board computer (BeagleBone Green) that will replace two separate computational components (Xilinx Spartan 6 LX9 FPGA and Raspberry Pi 3). Due to the increase in power and replacement of certain components, we hope to produce a more efficient and cost effective aperture synthesis radio telescope with a smaller physical footprint.

Real-time control using sequential inference and Bayesian decision theory

A. D. Martin and T. C. A. Molteno

This paper formulates a framework for real-time control using a Bayesian control-system. A Bayesian controller uses a Bayesian Sensor, which can characterise its own un- certainty. The sensor outputs a posterior distribution, P, and the controller, given this distribution, can calculate the cost of different actions, a. It chooses the action, aP that minimizes the expected cost; this is known as the Bayes act. This framework is demonstrated by an example of detecting mains frequency fluctuation, and shutting down equipment to mitigate the resulting damage. The Bayesian controller is contrasted with a conventional control system, which uses Fourier methods. The Bayes act is calculated for a range of equipment and shutdown cost ratios and the algorithm performance is verified. The Bayesian controller detects frequency changes much faster than the conventional controller, and shuts the equipment down at the optimum time to minimize loss.

Abstracts, Continued

Measurement of Implant Electrode Leads using Time-Domain Reflectometry to Predict the Resonant Length for MRI Heating

Steven McCabe, Jonathan Scott

Magnetic Resonance Imaging (MRI) machines can generate hazardous RF heating of patients with implanted neurostimulation leads. Consequently, most patients with these implants are contraindicated from having MRI scans. The level of RF heating has a strong dependence on lead length and is most severe when the length is close to a specific resonant length. Recent studies have shown that simple modifications to the lead construction and insulating material can alter the resonant length and significantly ameliorate this heating hazard, achieving MRI safety. We propose a technique using time domain reflectometry (TDR) to find the resonant length of an arbitrary lead such to minimise the amount of MRI machine time needed to find the length of highest heating. The results are compared with temperature measurements made in a 3-Tesla MRI machine and with a CW dipole radiator in the lab.

X-ray free-electron lasers: A revolution in biomolecular imaging

Rick P. Millane, David H. Wojtas and Romain D. Arnal

X-ray crystallography is a technique for imaging molecules at atomic resolution that is based on the diffraction of x-rays by a crystalline specimen. X-ray free-electron lasers are a new x-ray source that can replace conventional synchrotron sources used for x-ray crystallography experiments. They produce extremely bright and brief x-ray pulses, and are revolutionising high resolution imaging of biological macromolecules as they circumvent difficulties with specimen preparation and radiation damage. In these experiments, diffraction data are collected before the specimen is destroyed by the intense x-ray pulse. This requires new experimental protocols and new data processing techniques. Key aspects of biomolecular imaging using x-ray free-electron lasers are reviewed and future opportunities discussed.

SkyPi: All-sky camera network for meteorite tracking

Timothy C. A. Molteno

We describe the hardware and software for SkyPi – an all-sky camera network for transient event detection in the night sky. The hardware design is presented as well as the operating software and some preliminary results.

Vector Green's Functions for Electrodynamics Applications

Malcolm Morrison, Colin Fox

The use of scalar Green's functions is commonplace in electrodynamics, but many useful systems require computation of one or more vector quantities. Vector Green's functions have been used in electrodynamics for many decades, but are inconsistent and poorly understood. We recast vector Green's functions in the language of distribution theory to match their scalar counterparts. We use vector Green's functions to calculate the electrodynamic vector potential under physical boundary conditions.

Abstracts, Continued

Fractional Behaviour of Rechargeable Batteries

Rahat Rasan, Jonathan Scott

For decades authors have preferred to model batteries with either Thevenin-style models using RLC, or Randles-style by adding a Warburg element. These are claimed to model accurately. We present convincing empirical evidence suggesting that a fractional-derivative (constant-phase element) model is required. Our data shows that existing state-of-the-art models may be overly complicated, requiring numerical rather than physical considerations to find parameters.

A MATLAB ToolBox for the Bit-Stream Implementation

Faiz Rasool, Syed Adeel Ali and Muhammad Waseem Soomro

The supremacy of bit-streams over dedicated VLSI technology in the implementation of parallel, real time systems, in terms of cost and system structure, has been reported in recent research activities. The aim of this paper is to demonstrate the development and the implementation of a library of bit- stream elements, in a GUI environment like the Simulink This gives the advantage of using the bit- stream elements in conjunction with a large number of existing continuous time Simulink elements. Thus, all sorts of new or existing system models, supported by Simulink can be tested under bit- stream environment. The focus of this study is on the bit-stream implementation of building blocks for Fuzzy Systems and Artificial Neural Networks.

Transient Array Radio Telescope: Calibration and Aperture Synthesis

Max Scheel, Timothy C. A. Molteno and Phill C. Brown

The Transient Array Radio Telescope (TART) is a 24-element aperture synthesis array radio telescope. It is designed as a test bench for imaging algorithm development, as well as a survey-instrument for transient events. We present an update for the telescope hardware, describe software for traditional radio-interferomic aperture-synthesis imaging and also explain the process of calibration. While refining the calibration is under ongoing development, the number of the imaged sources that are consistent with known sources provides some confidence in the telescopes conceptional design and operation.

Abstracts, Continued

On a representation of feedback shift-registers

Daniel Schumayer and Timothy C. A. Molteno

Feedback shift-registers (FSRs) are fundamental building blocks in specialised digital electronics, e.g., cryptographic or counter circuits. FSRs have exponentially large number of possible states and their mathematical description is well developed. Constructing maximal length sequences for linear FSRs can be done algorithmically. However, the analysis of nonlinear FSRs is in its infancy. In the last decade a novel algebraic structure has emerged in the field of Boolean networks, which may provide a useful tool to describe FSRs. Here we explain the main ideas of this new formalism via examples and show its capabilities on a particular NFSR.

GPS Estimates in the Presence of Correlated Noise

Andrew Soundy, Daniel Schumayer, Timothy Molteno

This paper presents an experimental and theoretical investigation of GPS errors. Data from stationary GPS units were gathered and the time dependence of the observed noise is presented. Two quantitative approaches for dealing with time dependent noise are provided.

Investigation of Light Transmission in Healthy and Rotten Onions

Jason Sun, Rainer Künnemeyer, Andrew McGlone, Nathan Tomer

Light transmission through healthy and rotten onions was investigated in support our development of optical sensors for internal defects. Two types of systems were used: the first consisted of a broadband light source and a CCD spectrometer. Here the spectra had two peaks in the wavelength range from 670 to 850 nm. The second system employed a 728 nm laser diode and a rotating table, which enabled us to examine the light transmission spatially. Healthy, botrytis and pseudomonas onions were measured. The laser system could detect a signal attenuated by a factor of $3.25*10^{-8}$ to $1.81*10^{-7}$ due to the onions. The light level decreased exponentially with distance between source and detector. The effective attenuation was estimated for each onion and increased proportional to the degree of rottenness. The results of this study reveal the light transportation in onions and the effect from internal rots, which provide requirements for signal processing and optical systems.

Enhancing Students' Conceptual Understanding Using Visualisation Software and Collaboration

Elliot Varoy, Nasser Giacaman, Gerard B Rowe

Students entering into Engineering education at the tertiary level aren't prepared. They lack fundamental understanding of the requisite subjects, namely physics and calculus, resulting in a snowball effect of misunderstanding in the tertiary sector. In order to remedy this issue we need to focus on improving students' conceptual understanding instead of enforcing rigourous memory-based learning. This paper recognises two viable teaching methodologies, visualisation and collaboration, by proposing, developing and evaluating a mobile and tablet application that utilises these ideas to encourage deeper conceptual understanding in students.

Effects of non-uniform sampling on phase retrieval

Romain D. Arnal
Computational Imaging Group
Department of Electrical and Computer Engineering
University of Canterbury
Christchurch, New Zealand.

Rick P. Millane
Computational Imaging Group
Department of Electrical and Computer Engineering
University of Canterbury
Christchurch, New Zealand.
Email: rick.millane@canterbury.ac.nz

***Abstract*—Diffraction imaging involves reconstruction of an object from measurements of a diffracted field. In many cases, only the amplitude, but not the phase, of the diffracted field can be measured. Reconstruction of the object then requires recovering the phase, which is referred to as phase retrieval. Uniqueness properties of phase retrieval are thus important. Uniqueness properties of the phase problem are considered in cases where Fourier amplitude data are sampled non-uniformly in a way that arises in imaging of 1D and 2D crystals. Simulations of phase retrieval using an iterative projection algorithm show the influence of sampling non-uniformity in terms of the difficulty of reconstruction and the resulting reconstruction error. The results show that the sampling for 1D, and particularly 2D, crystals leads to more convergence difficulties and increased reconstruction error. The results have implications for *ab initio* phasing in imaging 1D and 2D crystals using x-ray free-electron lasers.**

Keywords— uniqueness; crystal; constraint ratio; phase problem

I. Introduction

Reconstruction of an object $f(\mathbf{x})$, where $\mathbf{x}$ is position, from measurement of radiation diffracted from the object is referred to as diffraction imaging. Examples are optical astronomy, x-ray crystallography, and microscopy [1], [2]. Under quite wide conditions, the diffracted field is equal to the Fourier transform, $F(\mathbf{u})$ of $f(\mathbf{x})$, given by

$$F(\mathbf{u}) = \int_{-\infty}^{\infty} f(\mathbf{x}) e^{-2\pi i(\mathbf{x}\cdot\mathbf{u})}\, \mathrm{d}\mathbf{x}. \tag{1}$$

In many cases, it is possible to measure only the Fourier intensity $I(\mathbf{u})$ of $f(\mathbf{x})$ given by

$$I(\mathbf{u}) = |F(\mathbf{u})|^2. \tag{2}$$

Knowledge of the Fourier intensity alone precludes reconstruction of $f(\mathbf{x})$ by the inverse Fourier transform as the Fourier phase is unknown [3]. This is known as the 'phase problem' in diffraction imaging [4]. 'Solving' the phase problem consists of retrieving the phase information from the intensity data.

Uniqueness of the solution to the phase problem in particular cases is important, i.e. does the known Fourier intensity $I(\mathbf{u})$ and any additional information available on the object lead to a unique reconstruction of $f(\mathbf{x})$? To this end, Elser and Millane [7] considered the case of an isolated object for which the continuous Fourier amplitude is available and defined a constraint ratio Ω given by

$$\Omega = \frac{|A|}{2|S|}, \tag{3}$$

where $|A|$ is the volume of the autocorrelation of the object and $|S|$ is the volume of the object (i.e. the number of degrees of freedom in the object). For $\Omega > 1$, it is likely than a unique solution to the phase problem exists, and for $\Omega \leq 1$ many other alternative solutions exist. In practice, $\Omega \gg 1$ leads to easily solvable phase problems and for Ω above but close to 1, the solution can be difficult to find. For a single object, $\Omega = 4$, phase retrieval is thus well-conditioned in this case and a unique solution is expected. We also recently derived the constraint ratio for macromolecular crystals for which the Fourier intensity is available only at the Bragg positions [8]. The constraint ratio in this case is equal to 0.5 unless additional information is available and the solution is not unique.

Interesting cases occur when continuous and Bragg sampled Fourier intensity coexist in a measurement. This occurs for the cases of imaging one-dimensional and two-dimensional crystals, as it is being explored with the application of x-ray free-electron lasers [9]. In the 1D crystal case, as a result of the object being periodic in one direction, the Fourier intensity information is available continuously along planes in Fourier space but the planes are available only at the Bragg spacing. In contrast, for 2D crystals, the data is Bragg sampled on planes and continuously sampled along the remaining dimension. Visualisation of the Fourier space sampling for both cases is shown in Fig.1.

Uniqueness of the phase problem for 1D and 2D crystal, is investigated. In Section 2, the constraint ratio is derived for these cases. In Section 3 we outline the difference map algorithm for phase retrieval that is used in the simulations. Simulation results and their interpretation are presented in Section 4, and concluding remarks are made in Section 5.

II. Theory

The constraint ratio, Ω, the number of independent data to the number of object parameters, as described above, is a useful quantity to characterise uniqueness. Since the problem is non-linear, $\Omega > 1$ does not necessarily guarantee

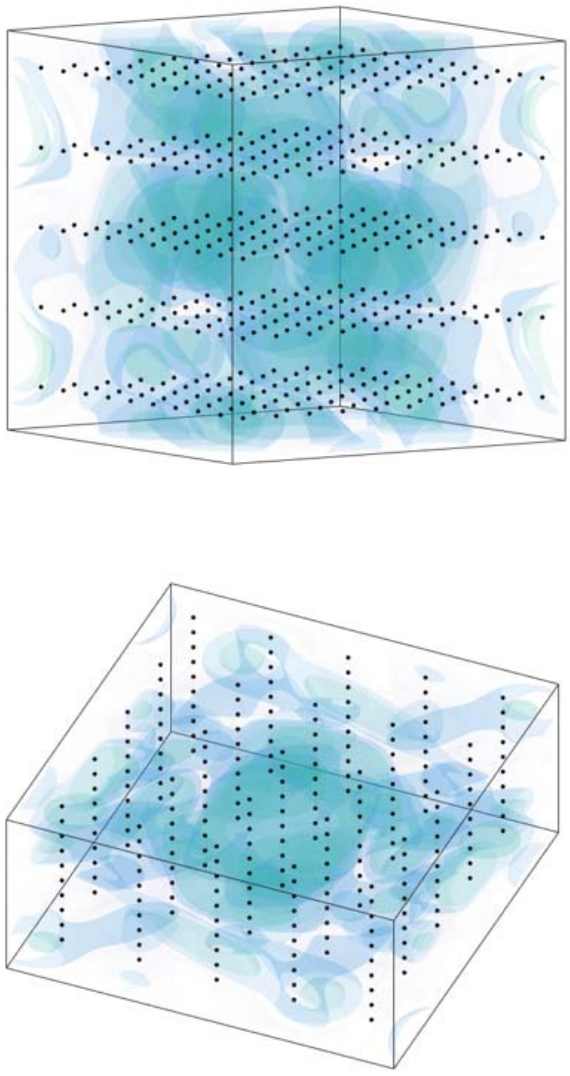

Fig. 1. **Top**: Fourier intensity with black dots indicating the sample positions for a 1D crystal. The samples lie on horizontal planes. **Bottom**: Fourier intensity with black dots indicating the sample positions for a 2D crystal. The samples lie on vertical lines.

uniqueness, but it does ensure that solutions belongs to a low-dimensional set. We therefore evaluate the constraint ratio for 1D and 2D crystals.

Using (3) the constraint ratio for the 1D and 2D crystal cases are derived, for both the continuous and discrete (sampled) cases. We consider only 3D objects. For an isolated 3D object with the continuous Fourier intensity available, the constraint ratio is

$$\Omega = 4. \tag{4}$$

For the discrete case, as the autocorrelation of a discrete $N \times N \times N$ signal has $(2N-1) \times (2N-1) \times (2N-1)$ samples, and half of them are independent, the discrete (denoted by the prime) constraint ratio for a 3D object is

$$\Omega' = \frac{4N^3 - 6N^2 + 3N}{N^3} = 4 - \frac{6}{N} + \frac{3}{N^2}. \tag{5}$$

For a 1D crystal, every second plane of amplitude data is absent from the dataset and the constraint ratio is decreased. To study this effect, we start with a fully sampled data set and remove samples from these planes until only the 1D crystal data remains. Denoting by ρ the proportion of Fourier intensity coefficients removed from these planes, the constraint ratio is given by

$$\Omega'_{1Dc}(\rho) = (4 - \frac{6}{N} + \frac{3}{N^2})(1 - \frac{\rho}{2}), \tag{6}$$

where the $1/2$ corresponds to the proportion of these unknown planes amongst all planes. In the 1D crystal limit, all Fourier intensity coefficients are removed in these planes, $\rho = 1$, the discrete constraint ratio for the 1D crystal is

$$\Omega'_{1Dc}(\rho = 1) = \Omega'_{1Dc} = 2 - \frac{3}{N} + \frac{3}{2N^2} < 2. \tag{7}$$

In the continuous case,

$$\lim_{N\to\infty} \Omega'_{1Dc}(\rho = 1) = \Omega_{1Dc} = 2. \tag{8}$$

For the 2D crystal case, only 1D lines of fully sampled Fourier amplitude data remain, and these lines are spaced by the Bragg spacing in the two orthogonal directions. Similarly to above, we denote by μ the proportion of the Fourier intensity coefficients that are removed from the missing lines. The constraint ratio is then given by

$$\Omega'_{2Dc}(\mu) = (4 - \frac{6}{N} + \frac{3}{N^2})(1 - \frac{3}{4}\mu), \tag{9}$$

where the $3/4$ corresponds to the proportion of the lines removed amongst all lines. In the 2D crystal limit, all the Fourier intensity coefficients are removed from the inter-Bragg lines, $\mu = 1$, and the discrete constraint ratio for a 2D crystal is

$$\Omega'_{2Dc}(\mu = 1) = \Omega'_{2Dc} = 1 - \frac{3}{2N} + \frac{3}{4N^2} < 1. \tag{10}$$

In the continuous case,

$$\lim_{N\to\infty} \Omega'_{2Dc}(\mu = 1) = \Omega_{2Dc} = 1. \tag{11}$$

For the 2D crystal case, we also consider the case where a proportion of the object samples are known. If a fraction s of the object sample values are known, then the constraint ratio is increased and is given by

$$\Omega'_{2Dc}(\mu, s) = \frac{(4N^3 - 6N^2 + 3N)(1 - \frac{3}{4}\mu)}{N^3(1-s)}. \tag{12}$$

III. Iterative projection algorithms

Iterative projection algorithms (IPAs) are constraint satisfaction algorithms that are used to solve phase retrieval problems [10]–[12]. IPAs are used for the simulations described in the next section, and are briefly described here.

Considering a vector $\mathbf{z} \in \mathbb{R}^Q$ as the vectorization of $f(\mathbf{x})$, an object can be represented by a vector in $\mathbb{R}^Q$. In this multidimensional vector space, a constraint set $C_A \subset \mathbb{R}^Q$ is defined as the set of objects that satisfy the constraint A. A projection operator P_A onto a constraint set C_A takes a vector $\mathbf{z} \in \mathbb{R}^Q$ to the closest vector $\mathbf{y} \in C_A$, or

$$\mathbf{y} = P_A\mathbf{z} = \text{argmin}_{\mathbf{y}\in C_A}||\mathbf{y} - \mathbf{z}||, \tag{13}$$

where $\text{argmin}_{\mathbf{y}\in C_A} f(z, y)$ denotes the value of y that minimises $f(z, y)$. The Euclidean norm is generally used. Relaxed projections are a generalization of projections and allow projection further into or away from the constraint set. Relaxed projections are given by

$$F_A\mathbf{z} = P_A\mathbf{z} + \gamma_A(P_A - I)\mathbf{z}, \tag{14}$$

where γ_A is a constant that controls the amount of relaxation. Projections used for *ab initio* phase retrieval impose *a priori* knowledge on the object and the measured Fourier intensities. The current point in the vector space is referred to as the iterate. The Fourier space projection constrains the iterate to have the same Fourier amplitude as that measured and is given by

$$P_M\mathbf{z_i} = \mathcal{F}^{-1}\left(\sqrt{\frac{\mathbf{I}}{|\mathcal{F}(\mathbf{z}_i)|^2}}\mathcal{F}(\mathbf{z}_i)\right), \tag{15}$$

where $\mathbf{I}$ is the corresponding vector of measured Fourier intensities and $\mathbf{z_i}$ is the iterate at iteration i, and $\mathcal{F}^{-1}$ is the inverse Fourier transform. The real space projection generally constrains the support of the object to be the same as the known support region and is given by

$$P_S\mathbf{z} = \mathbf{z} \cdot \mathbf{S}, \tag{16}$$

where $\mathbf{S}$ corresponds to the *a priori* support indicator vector having 1's where the object is non-zero and 0's elsewhere.

Suitable combinations of these projections and their relaxed counterparts form one iteration of an IPA. An IPA is defined by an update rule which is a formula that gives the $i+1$ iterate as a function of the previous iterate.

Here we use the difference map (DM) algorithm whose update rule is [10]

$$\mathbf{z}_{i+1} = \mathbf{z}_i + \beta(P_S F_M \mathbf{z}_i - P_M F_S \mathbf{z}_i), \tag{17}$$

where β is a step size parameter, γ_S and γ_M are the relaxation parameters of F_S and F_M, which are generally chosen as $\gamma_S = -1/\beta$ and $\gamma_M = 1/\beta$.

IV. Simulations and results

Uniqueness properties of the phase problem with the sampling schemes described above are investigated by simulation by reconstructing objects with different sampled datasets. We used a $15 \times 15 \times 15$ real and positive valued random object which was zero-padded to obtain the Nyquist sampled DFT of size $29 \times 29 \times 29$ samples. The DM algorithm is used with positivity and support constraints, and the Fourier intensity projection (15) is applied only at the known Fourier samples.

Fourier amplitude data are removed randomly in the proportions ρ and μ as described above. The final errors in the reconstructed objects and their match with the Fourier intensity data are denoted e and E, respectively, and defined by

$$e = \sqrt{\frac{\sum(\mathbf{z}_i - \mathbf{z}_t)^2}{\sum \mathbf{z}_t{}^2}} \tag{18}$$

and

$$E = \sqrt{\frac{\sum(\sqrt{\mathbf{I}} - \sqrt{\mathbf{I}_i})^2}{\sum \mathbf{I}}}, \tag{19}$$

where $\mathbf{z}_t$ is the vectorised true object, $\mathbf{I}_i$ is the Fourier intensity of the reconstruction at iteration i and $\mathbf{I}$ is the intensity data. The difference map algorithm was run 10 times for each value of ρ and μ. Convergence is defined by $E < 10^{-3}$. Runs are considered unconverged if no convergence is obtained less than $10,000$ iterations.

A. 1D crystal

The object error for the converged runs and the average number of iterations for the 1D crystal case are plotted as a function of $\Omega'_{1Dc}(\rho)$ in Fig. 2 (top). As ρ approaches 1, or equivalently as $\Omega'_{1Dc}(\rho)$ approaches $\Omega_{1Dc} = 2$, the average number of iterations increases dramatically. Additionally, a small increase in the average reconstruction error is seen. The same metrics for random sampling of all the intensity data are shown in Fig. 2 (bottom). The error and average number of iterations are steady at around 5×10^{-3} and 200 iterations, respectively, in this case.

The convergence difficulties and degraded reconstruction error seen in Fig. 2 (top) compared to Fig. 2 (bottom) are attributed to the particular form of non-uniform sampling, for a 1D crystal. The sampling structure leads to a decrease in performance, in terms of number of convergences obtained, average number of iterations, and final reconstruction error, of the reconstruction algorithm. As the reconstructions are still good ($e < 2 \times 10^{-2}$ even at the 1D crystal sampling limit), uniqueness of the phase problem appears to be unaffected (i.e. only correct solutions exist for $\Omega > 1$).

B. 2D crystal

The object reconstruction error for the converged runs and the average number of iterations for the 2D crystal case are plotted as a function of $\Omega'_{2Dc}(\mu)$ in Fig. 3 (top). For $\Omega'_{2Dc}(\mu) > 1.51$ many runs converged with an error

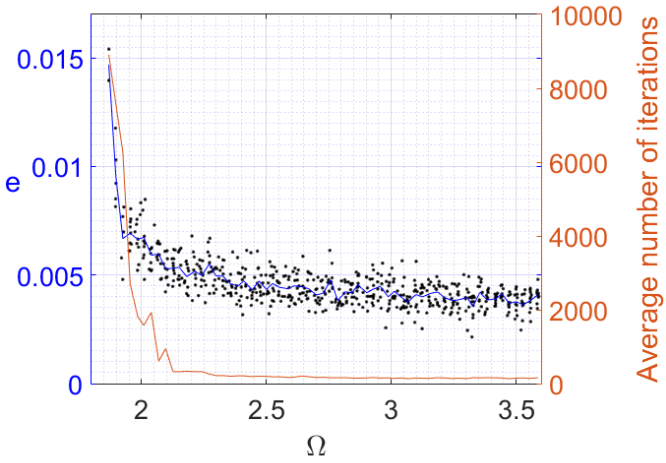

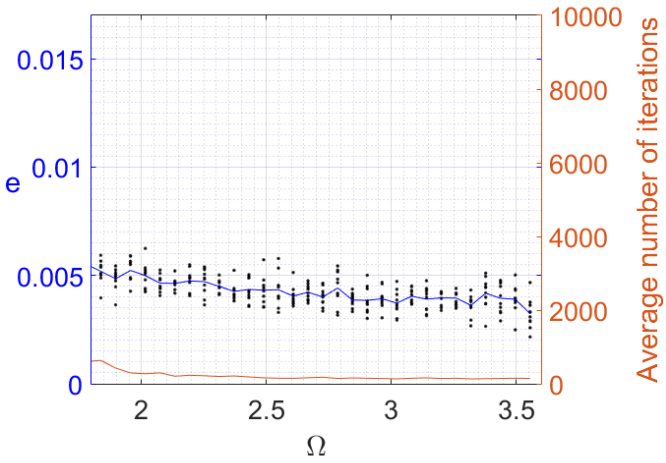

Fig. 2. **Top**: Object error e as a function of the constraint ratio, $\Omega'_{1Dc}(\rho)$, for a 1D crystal (dots). **Bottom**: Object error e as a function of the constraint ratio Ω', for a random sampling (dots). **Both**: The blue curve shows the mean error. The brown curve shows the average number of iterations.

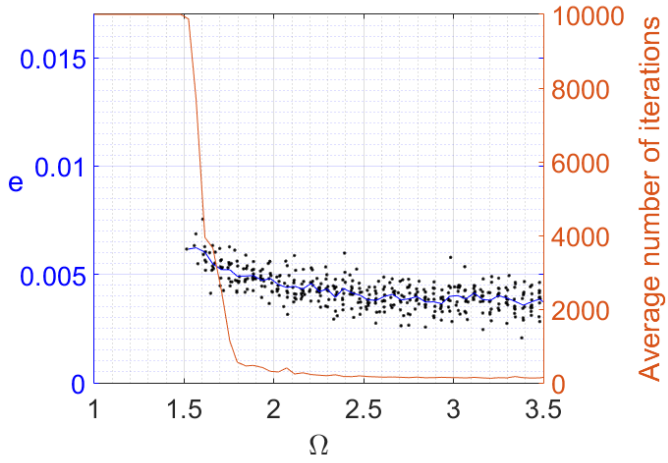

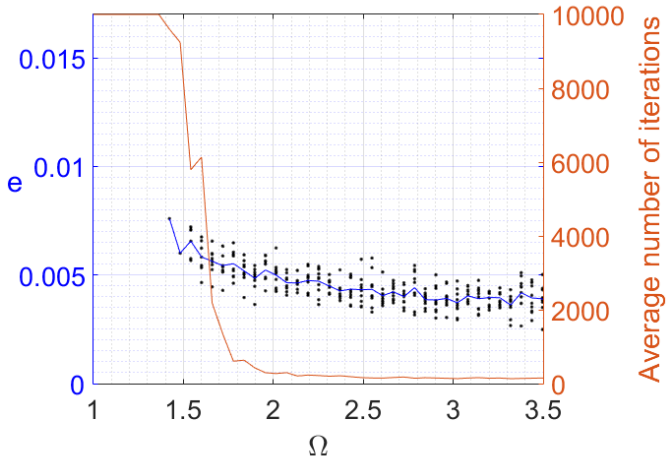

Fig. 3. **Top**: Object error e as a function of the constraint ratio, $\Omega'_{2Dc}(\mu)$, for a 2D crystal (dots). **Bottom**: Object error e as a function of the constraint ratio Ω', for a random sampling (dots). **Both**: The blue curve shows the mean error. The brown curve shows the average number of iterations.

of around 5×10^{-3} but no convergence was obtained for $\Omega'_{2Dc}(\mu) < 1.51$. The same metrics, for random sampling of the intensity data, are plotted in figure Fig. 3 (bottom). Similarly to Fig. 3 (top), no convergence was obtained for $\Omega' < 1.43$

As in the 1D crystal case, 2D crystal non-uniform sampling of the Fourier intensity reduces the performance of the IPA. This can be seen in the respective minimum constraint ratio for convergence (i.e. $1.51 > 1.43$). Moreover, as no convergence was obtained for $\Omega' < 1.43$ for random sampling, this indicates that *ab initio* phasing will be difficult in practice with sparsely sampled Fourier intensities, and either additional data or more Fourier intensity data will be needed.

In order to obtain convergence for smaller values of μ, or equivalently with smaller $\Omega'_{2Dc}(\mu)$, we assume that 13% of the object samples are known. The object error is plotted as a function of $\Omega'_{2Dc}(\mu, s = 0.13)$ using (12), where 13% of the object is known to be zero, in Fig. 4 (top). As $\Omega'_{2Dc}(\mu, s = 0.13)$ approaches 1, the reconstruction error slowly increases, and at $\Omega'_{2Dc}(\mu, s = 0.13) = 1.24$ the average number of iterations increase dramatically indicating a difficult phase problem. The same metrics for random sampling of the intensity data are shown in Fig. 4 (bottom).

Comparison of Fig. 4 (top) and (bottom) shows similar trends but the curves are shifted to the left for the case of random sampling. The phase problem for 2D crystals is difficult and the structured sampling leads to reduced performance of the IPA, also shown by the respective minimum constraint ratio for convergence (i.e. $1.24 > 1.16$). As the reconstructions are still good ($e < 2 \times 10^{-2}$), uniqueness of the structured phase problem is unaffected. Contrary to the 1D crystal phase problem where reconstruction was possible (but difficult) with $\Omega = 2$, the 2D crystal phase problem with $\Omega = 1$ is not soluble on its own and considerable additional information is needed.

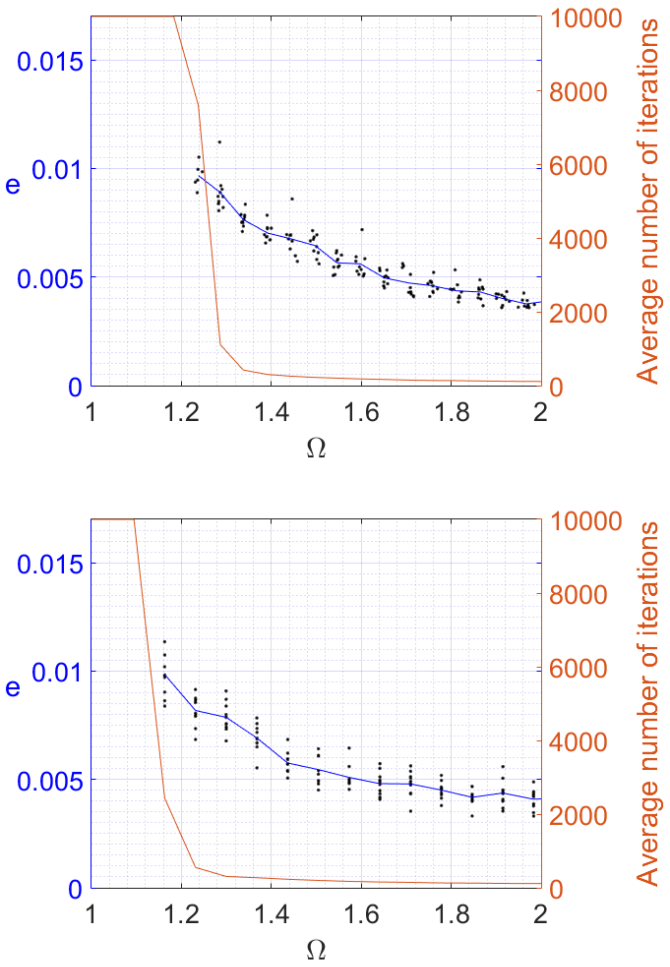

Fig. 4. **Top**: Object error e as a function of the constraint ratio, $\Omega'_{2Dc}(\mu, s = 0.13)$, for a 2D crystal with 13% of the object known (dots). **Bottom**: Object error e as a function of the constraint ratio $\Omega'(s = 0.13)$, for a random sampling (dots). **Both**: The blue curve shows the mean error. The brown curve shows the average number of iterations.

V. Conclusions

A particular non-uniform sampling of the Fourier intensities occurs with imaging of 1D or 2D crystals. Uniqueness properties of the corresponding sampling structures are illustrated by the results presented here. In the case of 1D crystals, the results indicate that the reduced sampling does not have a noticeable effect on uniqueness of the solution but performance of the reconstruction algorithm is reduced. Uniqueness is also preserved in the 2D crystal case but reconstructions can only be achieved if considerable additional information on the object is available.

Acknowledgments

This work was supported by a Marsden grant to RPM and a University of Canterbury College of Engineering Doctoral Scholarship to RDA.

References

[1] Millane, R. P. (1990). Phase retrieval in crystallography and optics. *JOSA A*, 7(3), 394-411.

[2] Bates, R. H. T. (1982). Fourier phase problems are uniquely solvable in more than one dimension. I: Underlying theory. *Optik*, 61, 247-262.

[3] Oppenheim, A. V. and Lim, J. S. (1981). The importance of phase in signals. *Proceedings of the IEEE*, 69(5), 529-541.

[4] Harrison, R. W. (1993). Phase problem in crystallography. *JOSA A*, 10(5), 1046-1055.

[5] Drenth, J. (2007). Principles of protein X-ray crystallography. *Springer Science & Business Media.*

[6] Hauptman, H. A. (1991). The phase problem of x-ray crystallography. *Reports on Progress in Physics*, 54(11), 1427.

[7] Elser, V. and Millane, R. P. (2008). Reconstruction of an object from its symmetry-averaged diffraction pattern. *Acta Cryst. A*, 64(2), 273-279.

[8] Millane, R. P. and Arnal, R. D. (2015). Uniqueness of the macromolecular crystallographic phase problem. *Acta Cryst. A*, 71(6), 592-598.

[9] Frank, M., Carlson, D. B., Hunter, M. S., Williams, G. J., Messerschmidt, M., Zatsepin, N. A., *et al* (2014). Femtosecond X-ray diffraction from two-dimensional protein crystals. *IUCrJ*, 1(2), 95-100.

[10] Elser, V. (2003). Phase retrieval by iterated projections. *JOSA A*, 20(1), 40-55.

[11] Marchesini, S. (2007). Invited article: A unified evaluation of iterative projection algorithms for phase retrieval. *Review of scientific instruments*, 78(1), 011301.

[12] Millane, R. P. and Lo, V. L. (2013). Iterative projection algorithms in protein crystallography. I. Theory. *Acta Cryst. A*, 69(5), 517-527.

Scalable Wireless Network for Remote Sensors: Range Testing

Elliot Baptist and Rainer Künnemeyer
School of Engineering
University of Waikato
Hamilton 3240, New Zealand
elliot.baptist@gmail.com

Philip Rowe
Accelerenz Ltd.
Hamilton, New Zealand
philip.rowe@accelerenz.com

Paul Martinsen
Blue Leaf Software Ltd.
Hamilton, New Zealand

***Abstract*—Wireless sensors attached to animals can improve the efficiency of farming, but need a cost-effective way of getting data to the internet. A network using low-cost packet radio modules has been proposed but development was hindered by a lack of information on working range of such radio modules. This report characterises and compares the range of RFM69HW and RFM95W radio modules for indoor and outdoor farm-like situations. It was found that RFM95W radio modules have approximately three times longer range in open air than RFM69HW modules. Viable RFM95W radio links were measured through 9 walls over 52 m in a building and over 1.85 km in open air. Using the RFM95W over the RFM69HW should reduce the cost of installing such a network in most farms.**

***Index Terms*—Packet radio module, range testing, wireless sensor network, RFM95W, RFM69HW.**

I. Introduction

With sufficiently accurate and timely data from sensors attached to animals, certain behavioural indicators can be identified in real time and used to increase the efficiency or effectiveness of farming activities. Existing examples of this include Cowlar [1] and Connecterra [2]. This report summarises testing to characterise and compare the range of two low-cost radio modules that have been identified as being a cost-effective way of collecting data from sensors mounted on animals and delivering it to the internet via a farm's fixed line internet connection.

A. Network System

It was decided that custom 'relay' devices should be used to receive data from nearby sensors and deliver it to the internet. Animals tend to congregate near or regularly visit key objects or areas such as water troughs and milking sheds. Relay devices can be installed in these locations and gain coverage over many sensors.

These relays are to use a low-cost sub-GHz ISM band radio module to receive sensor data and a similar but longer-range radio module to form a network with other relays. This inter-relay backbone network is to use custom software to enable data to be sent from relay to relay over the potentially large distance to the farm internet connection.

The choice of which long-range radio module to use in the design was made difficult by the lack of information on the operating range of specific hardware implementations in existing literature [3] [4] and product datasheets [5] [6], especially for the lower-cost modules. This report helps rectify this by characterising and comparing the operating range of two low-cost radio module options in two examples of farm-like conditions.

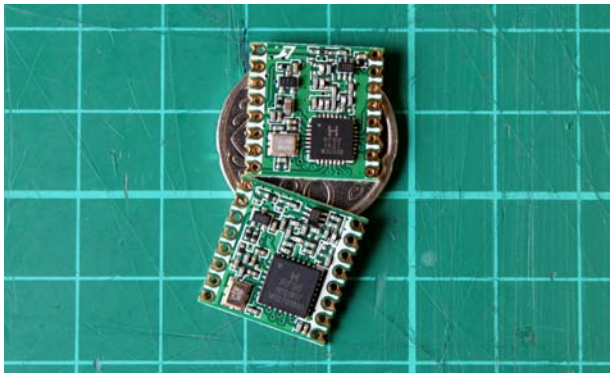

Fig. 1. RFM69HCW (top) and RFM95W (bottom, IC markings shared with RFM96W) radio modules on top of a New Zealand 20 cent coin and a centimetre grid. The RFM69HCW is a compact version of the RFM69HW and is functionally identical

B. Radio Modules

The RFM69HW [7] and RFM95W [8] radio modules (Figure 1) manufactured by Hope Microelectronics [9] were identified as cost-effective options for the backbone network. Prices were quoted in May 2016 as USD$1.89 each for RFM69HW modules and USD$4.20 each for RFM95W modules at quantities of 2000 direct from the manufacturer.

Unfortunately the lack of quantitative information about their range performance made it difficult to decide which was best suited for use on farms. The RFM69HW is a basic narrowband Frequency Shift Key (FSK) modulation transceiver while the RFM95W is a more advanced spread spectrum FSK transceiver. The code spreading technique of spread spectrum transmission used by the RFM95W allows greater transmission power and increased resistance to interference, both of which increase range by varying degrees in most circumstances. The RFM69HW can gain some of these benefits if frequency hopping spread spectrum techniques were implemented, but

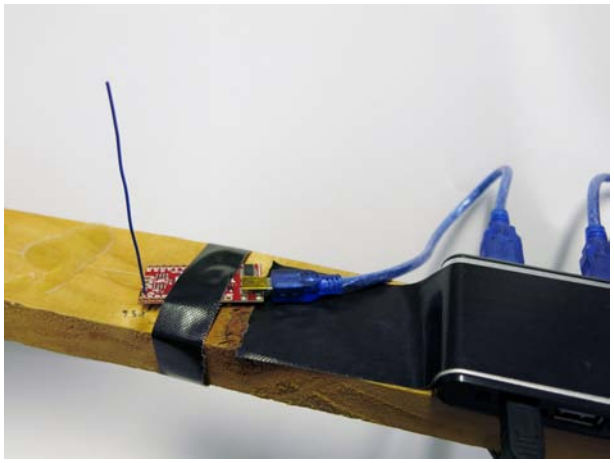

Fig. 2. A Moteino-USB with RFM95W on underside used as a 'receiver'

Fig. 3. Test mounts used, 'receivers' on table on left, 'transmitters' on mobile trolley on right

would require extra software development and increase system complexity.

This report covers efforts to quantify the effective range of each of these modules in two examples of farm-like conditions in order to decide which to use in the relay devices.

II. Methodology and Materials

A. Electronics

Radio modules such as the RFM69HW and RFM95W need supporting electronics to function. In this case they are operated by commands over an SPI bus, most easily implemented by attaching them to a microcontroller. The Moteino-USB (Figure 2) [10], regularly called simply a 'Moteino' within this report, produced by LowPowerLab LLC [11] was used as it combines the popular ATMega328P microcontroller with a Serial to USB bridge IC and one of a selection of radio modules, including the RFM69HW and RFM95W. These come with an appropriate $\lambda/4$ monopole wire antenna and are compatible with the RadioHead Packet Radio library [12].

B. Test Mounts

The testing hardware primarily consisted of two $1260\,\mathrm{mm} \times 52\,\mathrm{mm} \times 24\,\mathrm{mm}$ wooden beams mounted to tripods (Figure 3). Four Moteinos were spaced $350\,\mathrm{mm}$ apart along each in alternating order (69 - 95 - 69 - 95) and secured using cloth tape. One set of these had their USB cables connected to an external power capable USB hub also mounted on the beam while the other set had USB cables hanging free. The USB hub was used to connect one set of Moteinos to a tablet computer while the other set were connected directly to a laptop with four USB ports. These computers provided power and recorded data over USB. The 'transmitters' tripod had to be kept at a low height of $650\,\mathrm{mm}$ to allow the USB cables to reach the laptop.

C. Software

Many different pieces of software were used in testing. Embedded C++ with Arduino [13] and RadioHead libraries was used to program the Moteino-USB's to communicate with each other over radio and report performance over serial in CSV format.

The performance data was received over USB and stored in CSV format by a sketch written in Processing. Data from all test devices was collated, processed and plotted in Matlab. The map of in-building results was made using Inkscape.

D. Transmission Format

To test transmission performance, some payload must be transmitted from one radio module to another. To do this, the Moteinos on one tripod were designated the 'transmitters' and the others the 'receivers'. Each one of the 'transmitters' was assigned a matching 'receiver' to which they would send packets with a $4\,\mathrm{byte}$ RadioHead RHReliableDatagram header and $32\,\mathrm{byte}$ payload and get a $4\,\mathrm{byte}$ header with $1\,\mathrm{octet}$ payload acknowledgement packet in return.

Each payload packet contained a unique identifier. For every payload packet sent, this identifier and two performance values were output over serial to USB. These performance values were whether an acknowledgement packet was received from its paired 'receiver' and the Received Signal Strength Indicator (RSSI, see below) in dBm reported by the radio module if an acknowledgement was received.

If a 'receiver' got a payload packet, it replied with a brief acknowledgement packet and output the payload's identifier and the RSSI of the packet received over serial to USB. These were later matched up to their corresponding 'transmitter' measurements.

When enabled to transmit, 'transmitters' sent a new payload packet as soon as an acknowledgement for the previous payload was received or after a fixed time limit, whichever comes first.

E. Performance Indicators

The two performance indicators used were RSSI and the ratio of acknowledged packets to non-acknowledged packets.

RSSI is a measurement of the received signal power as seen by the radio module. It can indicate the remaining headroom in a fully functioning system when compared to the receiver sensitivity for those particular transmission settings. As the RSSI of packets not received inherently cannot be measured by the radio module itself, the performance of a system that is missing packets is better indicated by the ratio of payload packets that got a response to those that did not.

The RFM69HW and RFM95W calculate RSSI differently. The former samples during the preamble while the latter samples during the reception of the payload and presents an averaged value. As the 'receivers' and 'transmitters' each receive different length payloads, the RFM95W RSSI measurements reported by each for a given payload/acknowledgement cycle differ. As such, only the 'receivers' RSSI measurements were used as results as they are averaged over a greater time period.

F. Radio Configuration

Many parameters of both radio modules can be configured to optimal values. For these tests, all devices used a carrier frequency of $918\,\mathrm{MHz}$, a data rate close to $2.4\,\mathrm{kbit\,s^{-1}}$ and no transmission retries.

The RFM69HW devices were configured with a transmission power of $18\,\mathrm{dBm}$ ($2\,\mathrm{dB}$ lower than maximum), a preamble length of 8, an acknowledgement timeout of $200\,\mathrm{ms}$ and configuration registers set as: reg02=0x00, reg03=0x34, reg04=0x15, reg05=0x00, reg06=0xAD, reg19=0xF4, reg1A=0xF4, reg37=0xD0.

The RFM95W devices were configured with a transmission power of $21\,\mathrm{dBm}$ ($2\,\mathrm{dB}$ lower than maximum), a preamble length of 8, an acknowledgement timeout of $250\,\mathrm{ms}$ and configuration registers set as: reg1D=0x58, reg1E=0x74, reg26=0x04.

Most of these values were obtained through trial and error optimisation of failure range at lower output power levels.

It was also found that even if carrier frequencies $4\,\mathrm{MHz}$ apart were used, performance was worse if multiple radio pairs were run simultaneously than if run alone. Thus all tests were done with only one 'transmitter' transmitting at any one time.

G. Test Locations

Two farm-like locations were used for testing. To characterise operating range within and into farm buildings, a large barn-like workshop (Figure 3) was used to house the 'receivers' and the surrounding grounds and the connected building with sheet metal exterior walls (Figure 8) were used as obstacles. To characterise operating range across flat farmland, a $6\,\mathrm{km}$ long straight stretch of country road was used (Figure 4). Traffic was on the order of 8 vehicles per minute. 'Receivers' location coordinates: -37.75960 175.34415

H. Procedure

For the most part, the same test procedure was used for both open air and in-building tests. At the start of each test the 'receivers' were set up in a fixed location and communication with the 'transmitters' verified. Then data capture was enabled on the 'receivers' computer and the 'transmitters' moved to various test points. For the in-building test, the order of these points was randomised and the 'transmitters' equipment was moved on a resin trolley (Figure 3). For the open air test, the range was far less certain and each test point was chosen based on the performance at the previous points. The 'transmitters' equipment was packed in and out of a car to cover these greater distances.

Fig. 4. 'Transmitters' during open air, on road test

When at a test location, the test rig was orientated so that the transmitters were lined up approximately perpendicular to the direct path to the 'receivers'. One at a time each 'transmitter' had its data capture enabled on the connected computer and transmission initiated. The test operator would then step back approximately $2\,\mathrm{m}$ behind the test assembly to avoid interfering with the test until over 200 transmissions had been made over about $1\,\mathrm{min}$. The operator would then approach the computer, deactivate the current 'transmitter' and activate the next one until all had been run, then move on to the next test point.

III. Results and Discussion

A. Equipment Test

A quick test was conducted to evaluate the impact of powering the 'receivers' through the USB hub used and one other, compared to powering directly from a USB port. An RFM69HW RSSI drop of approximately $2\,\mathrm{dB}$ was observed from the external power capable USB hub eventually used and a drop of approximately $10\,\mathrm{dB}$ from the other hub. This difference between hubs was thought likely caused by the external power capable USB hub having better power delivery capability, potentially from lower power line resistance and/or greater capacitance. The other hub had a much smaller diameter USB cable, indicating a high resistance connection. The drop in signal may also indicate that the Moteinos contain insufficient power buffering for the radio module.

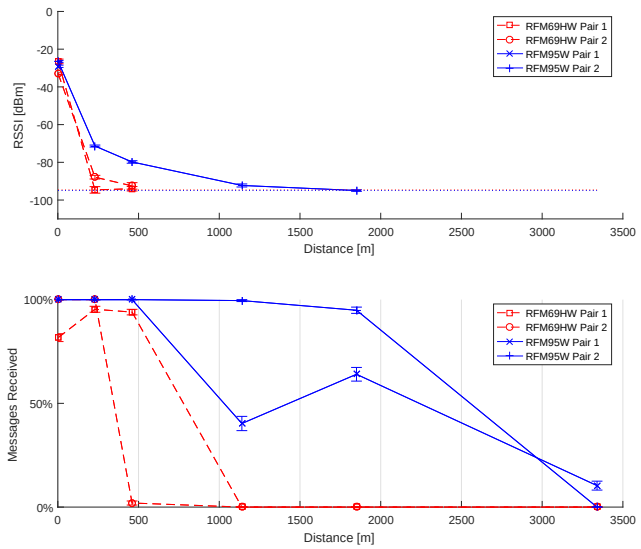

Fig. 5. Radio performance over the open air, on road distance tested. Floor is minimum value observed. Error bars are one standard deviation

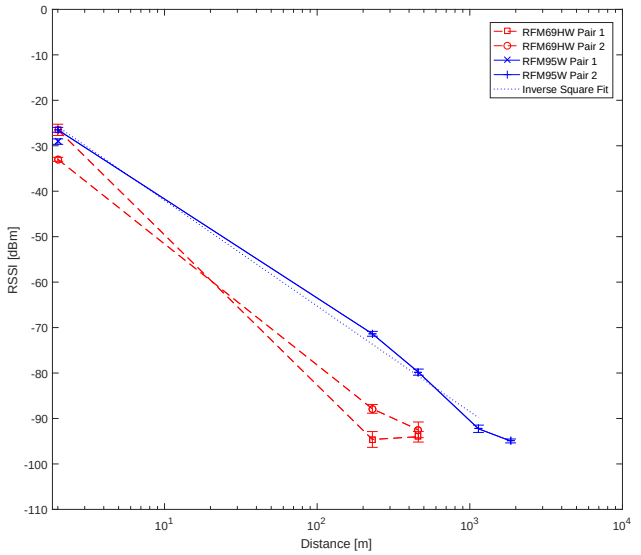

Fig. 6. Logarithmic plot showing fit to theoretical inverse square relationship between signal power and distance for open air signal power

B. Open Air Test

First discussed here are the open air results. The test locations are plotted over distance on both linear (Figure 5) and logarithmic (Figure 6) axes. They show a drop off in both signal power (RSSI) and packets delivered with distance as expected. Unfortunately, the desktop program recording the data from the RFM95W Pair 1 'receiver' failed after the first measurement location, resulting in a lack of RSSI data for that pair (Figure 5).

The signal power drop off indicated by the RSSI from the fully recorded RFM95W pair appears to follow the theoretical inverse square reduction (Figure 6). Insufficient data was recorded to determine if the RFM69HW pairs followed the same trend, though it is thought likely that they would.

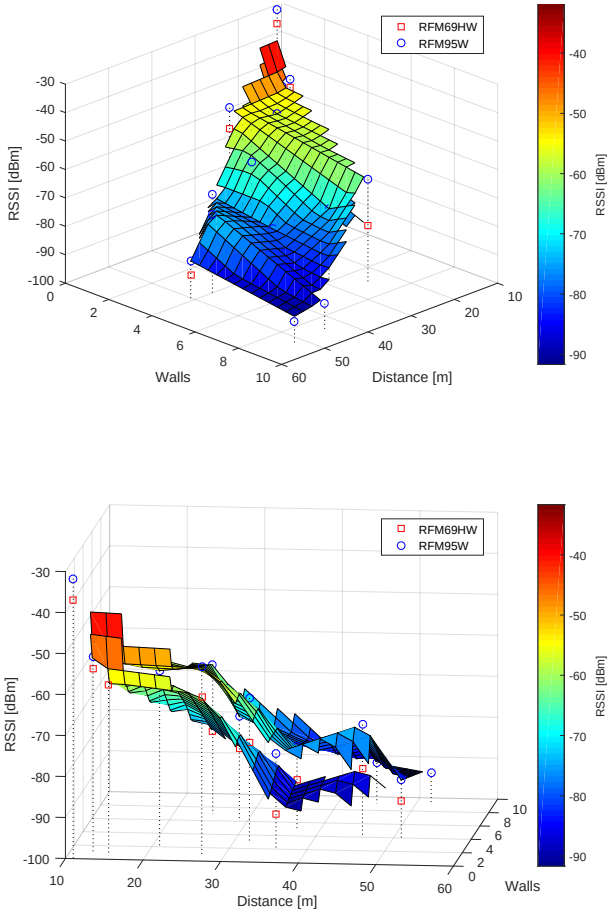

Fig. 7. Interpolated 3D surfaces depicting how RSSI varied with distance and number of walls in direct path and the difference between the two technologies for the in-building test

One RFM95W pair managed to get replies to 10% of packets sent to a 'receiver' 3.34 km away. Whereas the RFM69HW pairs failed somewhere between 460 m and 1.14 km, less than a third of the distance (Figure 5). In terms of a connection with an acceptably low packet loss, both RFM95W pairs had over 50% of packets conveyed successfully at 1.85 km, while the RFM69HW pairs managed the same at only 230 m, an eighth of the distance. Time constraints limited the number of test locations to those shown (Figure 5). These points show that the RFM95W obtained between three and eight times greater range than the RFM69HW.

Overall, this test shows that RFM95W radio modules also have significantly greater range than RFM69HW radio modules in open air.

C. In-Building Test

Unlike the open air test, distance is not the only major factor affecting radio performance in the in-building results. The building layout and construction affected each test location differently. Thus the RSSI results have been plotted as a function of walls as well as distance (Figure 7) in order

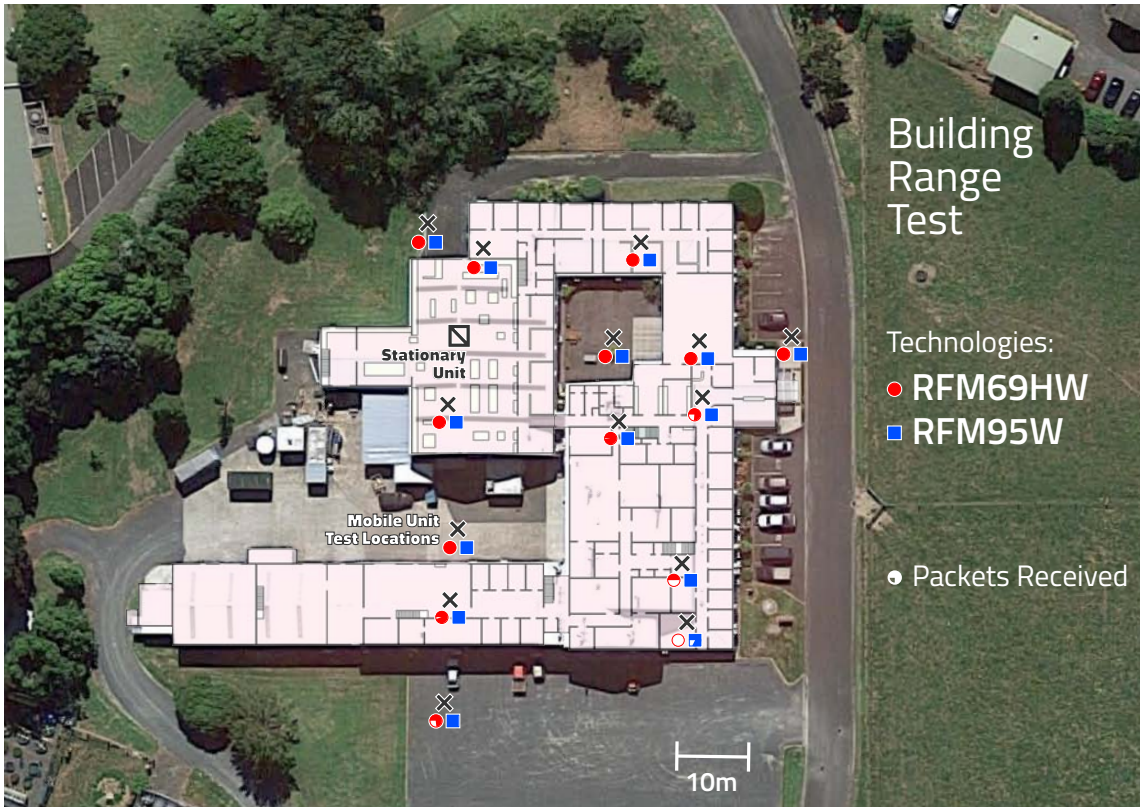

Fig. 8. Overhead map of in-building test area with ground floor internal diagram and successful packet reception results overlaid

to include some representation of the amount of building included in each test path. Similarly the packet loss results have been overlaid on an overhead map including the interior layout of the ground floor (Figure 8).

The signal attenuation with distance in this test (Figure 7) is significantly more than in the open air test (Figure 5). This indicates that the additional building and building contents that comes with distance is reducing the signal power more than the fundamental inverse square power reduction over distance. Thus the effect of distance could vary greatly depending on the building, making predicting the range in other buildings based on these results difficult. However, as the building tested has metal exterior walls and contains much equipment, it is likely somewhat comparable to many metal farm buildings containing farming equipment.

Between the number of walls on the direct path and the distance of that path distance, it appears distance had the greatest effect on signal power (Figure 7). The locations with the most walls took a diagonal path that offered more opportunities for alternate signal paths, adding additional variation to measurements with a similar wall count. Further, the distinction is not made between the differing signal attenuation presented by different wall materials. Despite this, the reduction in RSSI due to distance is approximately three times greater than the reduction due to a greater number of walls for a given distance.

Over the fourteen test locations, one or both of the RFM69HW pairs missed some or all packets at six locations. This compared to the RFM95W pairs which at worst only missed 17% of packets at one location (Figure 8), at which both RFM69HW pairs failed completely. This was the furthest corner of the building with nine walls in the direct path over a distance of $52\,\mathrm{m}$.

Overall, this test shows that RFM95W radio modules have noticeably greater range in and around buildings than RFM69HW radio modules, managing over $52\,\mathrm{m}$ through 9 walls.

D. Combined

Taking into account the outcomes of both tests gives a more complete view of the difference between these radio modules. For a little over twice the price, RFM95W modules get around three times the open air range and better building penetration than RFM69HW modules. This greater range allows the distance between the monitoring location and the internet connection be serviced by fewer relay devices. Considering the notable cost of the supporting components of each relay, this could significantly reduce cost or increase redundancy for long range or building dense installations. As most farms are spread out over large distances and have many metal walled structures, it is likely that these advantages from using the RFM95W over the RFM69HW would be realised in most cases.

IV. Conclusion

Range testing was conducted to determine which of RFM95W or RFM69HW radio modules would be best for use in a scalable wireless network for transporting sensor data. Testing showed that RFM95W modules have approximately three times greater range than RFM69HW modules in open air and have greater building penetration. Viable RFM95W radio links were observed up to 9 walls and $52\,\mathrm{m}$ through buildings and over $1.85\,\mathrm{km}$ in open air. Using the RFM95W instead of the RFM69HW in the proposed wireless network should reduce the cost of installation for most farms.

References

[1] Cowlar, "Cowlar homepage," Available: http://cowlar.com/ [Online; Accessed 12 October 2016].

[2] Connecterra, "Connecterra homepage," Available: http://www.connecterra.io/ [Online; Accessed 12 October 2016].

[3] A. Hardy, F. Bouhafs, and M. Merabti, "A survey of communication and sensing for energy management of appliances," *International Journal of Advanced Engineering Sciences and Technologies*, vol. 3, no. 2, pp. 61–77, 2011.

[4] P. Rawat, K. D. Singh, H. Chaouchi, and J. M. Bonnin, "Wireless sensor networks: a survey on recent developments and potential synergies," *The Journal of Supercomputing*, vol. 68, no. 1, pp. 1–48, 2013.

[5] Hope Microelectronics Co., Ltd., "RFM69HW ISM transceiver module V1.3 datasheet," Available: http://www.hoperf.com/upload/rf/RFM69HW-V1.3.pdf [Online; Accessed 22 July 2016].

[6] ——, "RFM95/96/97/98(W) - low power long range transceiver module V1.0 datasheet," Available: http://www.hoperf.com/upload/rf/RFM95_96_97_98W.pdf [Online; Accessed 21 July 2016].

[7] ——, "RFM69HW product page," Available: http://www.hoperf.com/rf_transceiver/modules/RFM69HW.html [Online; Accessed 12 October 2016].

[8] ——, "RFM95W product page," Available: http://www.hoperf.com/rf_transceiver/lora/RFM95W.html [Online; Accessed 12 October 2016].

[9] ——, "HopeRF homepage," Available: http://www.hoperf.com/ [Online; Accessed 12 October 2016].

[10] LowPowerLab LLC, "Moteino-USB product page," Available: https://lowpowerlab.com/shop/product/100 [Online; Accessed 12 October 2016].

[11] ——, "LowPowerLab homepage," Available: https://lowpowerlab.com/ [Online; Accessed 12 October 2016].

[12] M. McCauley, "RadioHead packet radio library for embedded microprocessors," Available: http://www.airspayce.com/mikem/arduino/RadioHead/ [Online; Accessed 12 October 2016].

[13] Arduino, "Arduino homepage," Available: https://www.arduino.cc/ [Online; Accessed 12 October 2016].

Towards Vocal Health Assessment From The Speech Signal

Itay Ben-Dom
Department of Electrical and Computer Engineering
University of Auckland
Auckland, New Zealand
Iben350@aucklanduni.ac.nz

Catherine Inez Watson
Department of Electrical and Computer Engineering
University of Auckland
Auckland, New Zealand
c.watson@auckland.ac.nz

I. ABSTRACT

In speech processing, the glottal pulse contains information about the speaker's voice quality. Through glottal pulse processing, the time-domain parameter Open Quotient can be extracted from a glottal pulse. This paper presents a new Open Quotient criterion, called OQsub50. This paper introduces the implementation of an Open Quotient detection method in *R*. Phonetic analysis of vowel production was carried out over two New Zealand English (NZE) speech corpora, both containing 11 NZE monophthongs. This study finds the decrease of Open Quotient measurements from young to old age. The trend was observed in both male and female speakers, suggesting the Open Quotient measurements are independent of pitch.

Key Words: *Voice quality, glottal pulse, Open Quotient (OQ), glottal opening instance (GOI), glottal closing instant (GCI), emuR.*

II. INTRODUCTION

Extraction of parameters from a speech signal can provide invaluable information about the speaker's voice quality. During voiced phonation, air expelled from the lungs flows up the larynx and causes quasi-periodic vibration of the vocal folds. The airflow exiting the vocal folds opening (glottis) results in glottal flow. Glottal flow is known as the excitation signal for voiced speech. Extraction of the glottal waveform from the speech signal is of significant importance, as it relates to pitch, volume, voice quality, emotional speech, etc.

There are several methods to obtain information regarding glottal behaviour. One of those methods is Electroglottography [1-3]. Electroglottography (EGG) provides a measure of the change in impedance across the vocal folds as they vibrate; the larger the contact region of the folds, the larger the impedance. Thus EGG provides a measure of the changing contact of the vocal folds. However, during the glottal pulse open phase the folds are not in contact, so no information about the folds during the open phase can be deduced from the EGG signal. Therefore, even though EGG provides information about the vocal folds, it does not provide information about the glottal flow waveform. Moreover, the EGG method requires the use of an apparatus which is inconvenient and invasive. An alternative, non-invasive method is the estimation of the glottal waveform via inverse filtering of the speech waveform. This method requires only the recording of the speech signal. This method is widely popular and is the one implemented in this study [2-5].

Physiological changes affect the way we speak. Being produced via vibration of the vocal fold, aging impacts the voice quality, i.e. breathiness. The aged voice is discernible by ear, and in the acoustics of the speech waveform. The link between the vocal physiology and the acoustic of the speech waveform is well known and the impact of aging on the vocal system affects both the vocal tract [6], and the vocal folds [7-10].

This study extends earlier research investigating the impact that aging has on the glottal waveform. The effects of aging is tested through the glottal pulse time-domain parameter Open Quotient (OQ). OQ is a percentile representation of the glottal pulse open phase. Although Open Quotient is a widely used parameter, its definition lacks consistency among researchers. This lack of consistency, alongside seldom mentions of the implementation techniques, makes it difficult to compare results across studies. After analysis of speech corpora using the established OQ methods, it was decided that a more clearly defined OQ criterion was required in order to allow the comparison of results across studies. This led to the definition of a new Open Quotient criterion definition, called OQsub50.

Glottal source processing is an established research area. However, there are currently no digital implementations allowing for computationally-fast and user-friendly big data analysis. Even though there are voice analysis packages available in *MATLAB* (*COVAREP* [11], *VOICE TOOL BOX*), they do not allow for big data processing and are not aimed for extraction of time-domain parameters. It is the aim of this paper to present an analytical package allowing for robust extraction of features and parameters from speech waveforms.

The suggested OQsub50 parameterization method was tested on vowels extracted from databases containing speakers of different age and gender. Testing was carried out to identify trends correlating Open Quotient measurements a speaker's age and/or gender. The results of this analysis are presented and discussed.

III. Glottal Source Processing

A. Voiced Speech Production

Speech production is a physiological procedure. The human body phonetic system is comprised of three main organs: the lungs, the larynx, and the vocal tract. The lungs are a respiratory organ and are the source of energy, evicting airflow arising from the trachea. The larynx, holds the vocal folds/cords. The opening between the vocal folds is called glottis. When airflow passes through the larynx, it causes the vocal folds to vibrate quasi-periodically. The vocal folds vibrate at a fundamental frequency known as pitch. During this vibration, the vocal folds open and close, thus modulating the airflow as it passes though the glottis, resulting in glottal flow. The vocal tract "colours" the speech signal by spectral shaping of the glottal flow. The modulated and coloured signal is then radiated by the lips and perceived by the listener as spoken speech.

B. Glottal Source Extraction

The glottal volume velocity waveform can be extracted from the speech signal via source-filter deconvolution. In inverse filtering, the glottal pulse is obtained by cancelling the effects of the formants from the speech signal. To estimate the effects of the formants, the vocal tract is first modelled and the effects of the formants is then cancelled by filtering the speech signal through the inverse of the vocal tract model [12].

In the iterative adaptive inverse filtering method (IAIF), the glottal pulse and vocal tract filter function are modelled [13]. The initial estimation of the glottal excitation is obtained by cancelling the effects of the preliminary vocal tract and lip radiation by inverse filtering and integration, respectively. The procedure is then repeated using the glottal wave estimate from the first iteration in order to obtain a more accurate model for the glottal pulse. Alku suggested an improvement to his IAIF method by performing linear prediction analysis pitch synchronously, thus eliminating the influence of the pitch period by computing the IAIF analysis one fundamental period at a time using frames that span between 2 consecutive maximal glottal openings [5].

C. Glottal Flow Model

Speech production can be described according to the source-filter model [14]. This model provides an approximation for the glottal signal and the vocal tract filter function. The model's main advantage is its simplicity. Many time-domain models have been proposed: Rosenberg, Liljencrants-Fant (LF) Fujisaki-Ljungqvist (FL) models [15-17]. The LF glottal flow model is the most widely used model and is shown in Fig. 1 [18]. The glottal flow model allows for parameterisation of the glottal flow. In this study, the Open Phase is of importance. The Open Phase is a time-domain parameter of the glottal cycle, during which the vocal folds are separated and air flows through the glottis. It consists of both the Opening Phase (from initial glottis opening to maximum flow) and the Closed Phase (from maximum flow to glottis closure). The ratio of duration between the Open Phase and the glottal pulse cycle period is referred to as Open Quotient.

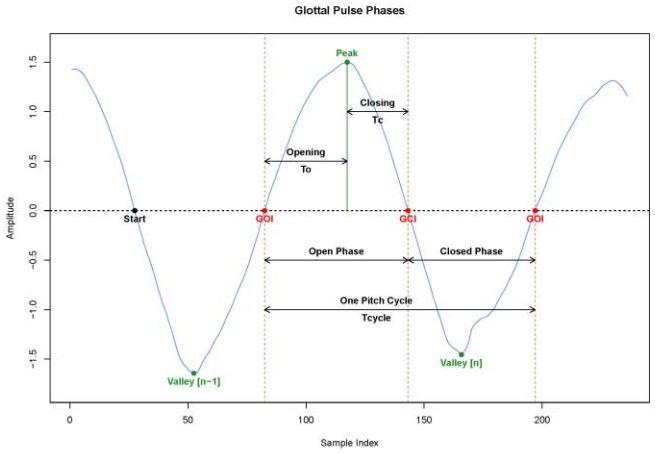

Fig. 1. LF-Model Glottal Pulse Phases

D. Open Quotient

The Open Quotient is a time-domain feature derived from the glottal pulse model. It is the time ratio between the glottis open phase duration and the glottal pulse period. Even though many researchers cite the use of the Open Quotient parameters, the wide range of Open Quotient criteria has made it difficult to compare results across investigations. Moreover, the lack of criteria specifications make it difficult to compare results even when analysis is carried out on the same database [2]. The three main analysis criteria that are often failed to be mentioned are: the definition of the open and closed cycle, interpolation of GOI and GCI points, and whether the analysis was carried out pitch synchronously. Determination of Open Quotient measurements requires the definitions of GOI and GCI points. Two popular Open Quotient criteria are the OQ50 and OQsub. It is important to note that both criteria have different definitions for the GOI and GCI points, as well as definition of the pulse period [3].

1) OQ50

OQ50 uses a 50% criterion level, where the open phase is defined as the time where the glottal flow is greater than a threshold amplitude level. This threshold is defined as 50% of the AC component (AC equals peak flow minus minimum flow). The GOI and GCI points correspond to the temporal points in the glottal waveform at the amplitude threshold level. The glottal pulse period is defined as the duration between two successive GOI points [3].

2) Subjective Airflow Open Quotient (OQsub)

OQsub (subjective airflow Open Quotient) is the ratio between the open phase and the period of vibration [2]. The OQsub is less prominent in research due to its imprecise definition of GOI and GCI points. The GOI point is defined as the point of initial increase from minimum flow. Although easy to implement, this definition is prone to errors if minima flow points exhibit small ripples due to noise coupling. The GCI point is defined to occur at the end of a sharp decline in glottal flow, followed by a sudden increase in flow. This definition is case-specific, as not all waveforms exhibit an increase in flow during the closing phase. Even though the OQsub

measurements are performed pitch synchronously, no mention of results interpolation had been noted in any other research papers.

E. Proposed Open Quotient (OQsub50)

In this paper we suggest the use of a new Open Quotient criterion called OQsub50, computed via Equations 1 and 2. As the name suggest, this criterion is a hybrid of OQ50 and OQsub. The GOI point is found via an amplitude threshold as in OQ50. The GCI point is found via the glottal waveform derivative. The glottal pulse period is defined as the time between two consecutive glottal pulse valleys, as in accordance to OQsub, thus the period is computed pitch synchronously to achieve accurate results. A breakdown of GOI and GCI detection methods follows.

$$OQsub50 = \frac{Open\ Phase}{T_{cycle}} \quad (1)$$

$$OQsub50 = \frac{GOI-GCI}{Val[n]-Val[n-1]} \quad (2)$$

1) GOI Detection

As per OQ50, the glottal opening instance is defined to occur at the first time instant where the flow reaches 50% (half) the peak-to-peak amplitude of the glottal waveform. The glottal opening instance is known to occur in the opening phase; between the beginning of a cycle and its peak point. The GOI detection is a two-fold process. First, the GOI location is roughly estimated. Given an amplitude threshold value, the GOI corresponds to first the glottal waveform sample point that exceeds the threshold, as shown in Fig. 2. Second, the glottal waveform is segmented to an 11 sample points range about the rough GOI point (five points preceding and five points following). An 8th order polynomial linear model is fitted to the segmented wave and its coefficients are determined. By finding the intersection between the polynomial curve and the amplitude threshold, the GOI temporal point can be determined. Solving the rearranged polynomial equation results in root values that correspond to intersection of the fitted polynomial curve with the threshold, as shown in Fig. 3. The result is the time-domain instance corresponding to the GOI point.

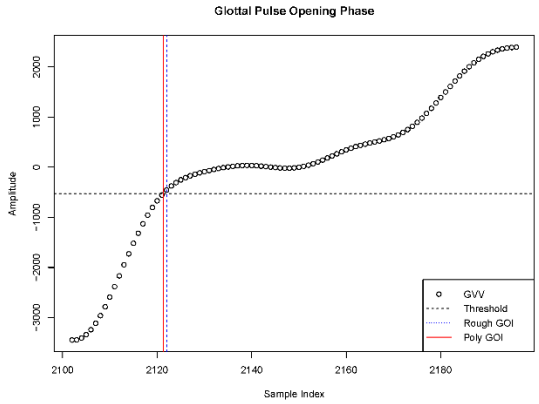

Fig. 2. GOI Detection Diagram

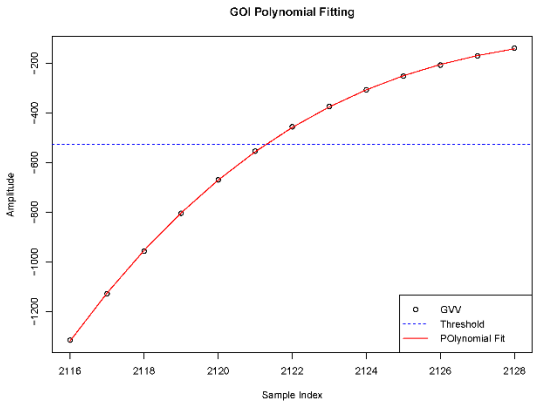

Fig. 3. GOI Polynomial Fitting

2) GCI Detection

In speech processing, there are various methods for the detection of GCI points from the speech signal [19]. This detection method is based on the LF model, where the glottal closing instant correspond to the point of minimum flow in the glottal waveform derivative [4, 16, 18]. The GCI occurs in the closing phase, between a maximum flow and a successive minima. The glottal waveform derivative produces a valley, as shown in Fig. 4. The valley corresponds to the GCI temporal point, with the exact time-instant found via quadratic interpolation.

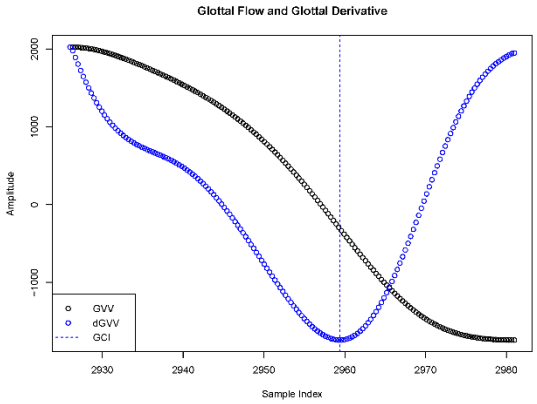

Fig. 4. GCI Detection Diagram

IV. IMPLEMENTATION

A. R Software

R is a free software environment for statistical computing and graphics [20]. Similar to *MATLAB* and its matrices compatibility, *R* has its own datatypes (list, data-frames), which allow for fast computation. Although not as popular as *MATLAB*, *R* allows for large database analysis. It has in-built functions (i.e. *apply*, *lapply*), which allow for fast processing of its datatypes. The analysis was implemented and carried out in *R* due to its large data processing capabilities and its simple integration with speech databases via its *emuR* package.

B. EMU R Package

The *emuR* package is a Speech Database Management System (EMU-SDMS). The package is a collection of software tools for management, preparation, extraction and analysis of speech databases [21]. The *EMU R* software, works with the *emuDB* database format, rather than *emu Legacy*. EMU-SDMS is a powerful tool which allows researchers to label, find, and extract speech segments (e.g. vowels) based on sequential and hierarchical structure of the utterances in which they occur [21]. The extraction of speech segments results in segment lists, which contain phonetic and temporal information about the extracted segment. Specialised commands then allow a method to be executed on the segment list in its entirety without the user's manual involvement.

C. Analysis in R

The time-domain glottal source parameterisation method for OQsub50 was implemented in *R*. Analysis was carried out on speech segments extracted via *emuR*. Each vowel's onset and offset sample speech segments were extracted. To minimize potential errors in the detection of vowel boundaries, the first and last 5% of those extracted vowels were removed before further processing, thus eliminating potential interferences from neighbouring parts of the speech signal (consonants).

The truncated speech waveform's glottal volume velocity waveform is extracted using Alku's pitch-synchronous IAIF (PSIAIF) algorithm [5]. The glottal flow is filtered through a low-pass Butterworth filter, with its cut-off frequency correlated to the speech signal's mean pitch. The glottal flow was segmented into cycles, with a cycle period defined as the duration between two successive flow minima. The flow minima, or valley, was found via a peak detection algorithm, with the exact time-instant found via quadratic interpolation to reduce quantisation error.

As mentioned previously, the amplitude threshold is required in order to determine location of the GOIs. The threshold was computed on a cycle-by-cycle basis, thus removing any DC offset bias which might have occurred if an average threshold was computed pitch asynchronously. Both GOIs and GCIS were computed using the aforementioned processes. When differentiated to produce the glottal flow derivative, the resulting wave was a crude approximation to the theoretical derivative. Thus, up-sampling interpolation was used to generate a smoother glottal flow waveform, which in turn resulted in a coherent glottal flow derivative approximation. This glottal derivative waveform's valley was detected using a peak detection algorithm. Following the detection of GOI and GCI instances, the Open Quotient was calculated.

D. MATLAB-to-R Conversion

MATLAB, as its name suggest, is designed for powerful and fast computation of matrices. However, *R* is computationally fast when processing vectors/data-frames. Two methods used in this analysis were adapted from a *MATLAB* implementation: pitch detection and IAIF method.

The pitch detection algorithm was adapted from an open-source pitch detection method in the COVAREP project [19]. The Fundamental Frequency (F0) is estimated from the residual signal through summation of residual harmonics (SRH). The method proved robust in adverse noise conditions. The *MATLAB* script contained a method to compute the SRH. This method proved to be the most time consuming in *R*, with execution time taking over 0.5 seconds, depending on speech waveform's length. Conversion of matrices into lists operation resulted in the lower computation time of .03 seconds.

The IAIF method was adapted from an open-source *MATLAB* function in the *COVAREP* project [5, 11]. As IAIF is implemented pitch synchronously, the function contains a FOR-loop. FOR-loops are notoriously slow in R. By converting the datatype to data-frames and replacing the FOR loop with the data-frame function *sapply*, the computation time was reduced by more than 80% of the initial implementation time.

V. Results

The purpose of this present study is to investigate correlation between Open Quotient and gender/age. Analysis was first carried out using the existing OQ50 algorithm. The results generated showed no statistically significant difference in OQ values between the age groups, which led to the examining of the OQ algorithm, resulting in the OQsub50 criterion. OQsub50 analysis was carried out for two speech corpora databases. Both databases include New Zealand English (NZE) speech. The results reflect vowel production comparison along the different age/gender groups. For each speaker, citation was collected in the form of hVd words for the 11 NZE monophthongs (both short and long vowels). Each speaker was asked to read out loud five word lists, with each list containing 12 hVd words in random order. The final hVd word was removed from each recorded speech sentence to avoid list effects [23, 24]. Sentences were automatically labelled at word and phonetic level by the Munich Automatic Web Segmentation System, webMAUS [26]. Labelled sentences were converted into an EMU database with vowel's onset and offset time extracted [27].

A. Speech Corpus A: Male Aging

The study was carried out using a database consisting of 30 male speakers. Participants were comprised of two age groups: 15 young speakers (20-26 years old, mean age = 23.3, SD = 1.8) and 15 old speakers (56-71 years old, mean age = 62.5, SD = 5.6) [1]. All speakers were male New Zealand English speakers. Per questionnaire, none of the participants reported specific vocal health issues at the time of recording. The recording task was performed in a whisper room using a laryngograph EGG equipment. Recording were stored with a 16 kHz sampling frequency as 16 bit numbers. The speech corpus contains 1604 vowel tokens, with 796 and 808 tokens for young and old speakers, respectively. Results are presented in Fig. 5 and Table 1.

TABLE I. Speech Corpus A Results

	Young	***Old***
Median	0.6570	0.5669
SD	0.0866	0.0583
Mean Pitch (Hz)	115.2	125.3

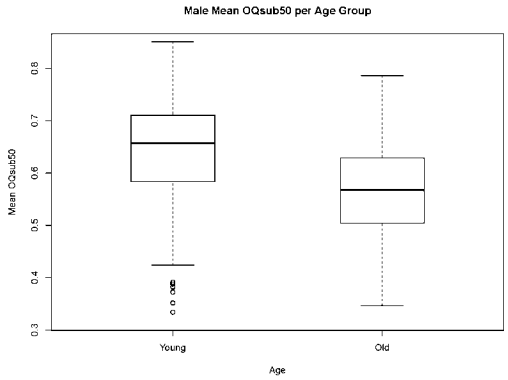

Fig. 5. Male Aging Mean OQ per Age Group

B. Speech Corpus B: Male and Female Aging

The study was carried out using a database consisting of 26 speakers of New Zealand English. It comprises of 15 male speakers and 11 female speakers. Participants' were categorised into three age groups: young (18-25), mature (40-50), old (65-80), see Table 2. Speech was recorded in a sound booth (Whisper Room MLD8484E) directly on to a Marantz PMD670 Solid State Recorder at a sampling rate of 16 kHz, using a Shure SM58 Microphone [23, 24]. The speech corpus contains 1247 vowel tokens, with 729 and 518 tokens for male and female speakers, respectively. Results are presented in Fig. 6, Fig. 7, and Table 3.

TABLE II. Number of Speakers per Gender/Age

Gender/Age	**Young**	**Mature**	**Old**
Male	7	6	2
Female	4	4	3

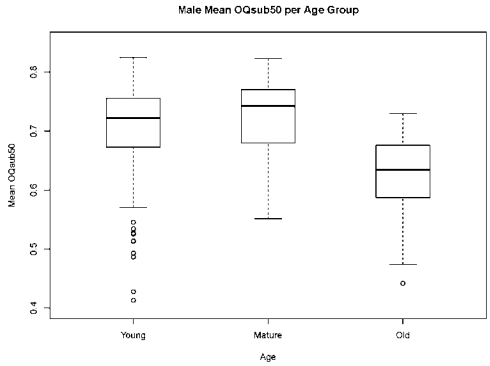

Fig. 6. Male Aging Mean OQ per Age Group

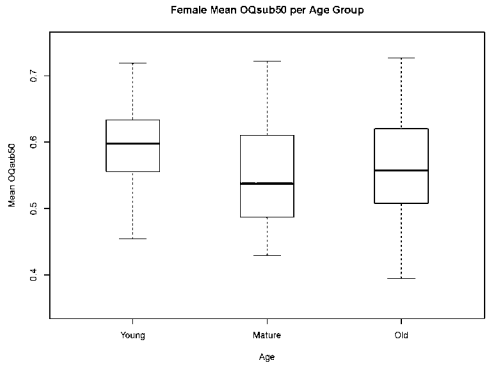

Fig. 7. Female Aging Mean OQ per Age Group

TABLE III. Speech Corpus B Results

	Male		
	Young	*Mature*	*Old*
Median	0.7218	0.7420	0.6338
SD	0.0627	0.0582	0.0617
Mean Pitch (Hz)	115.2	114.8	134.7
	Female		
	Young	*Mature*	*Old*
Median	0.5975	0.5372	0.5571
SD	0.0546	0.0800	0.0777
Mean Pitch (Hz)	178.3	183.0	169.4

VI. Discussion

Correlation between age and voice quality is a well-established idea in speech analysis. Voice production is a physiological process. With aging, the speaker's vocal organs exhibit degradation. Advancing age produces physiological changes that alter the voice and affect the voice quality. With aging, vocal folds deteriorate, losing collagen fibres (consisting of bone, cartilage, and other connective tissues). As a result, the vocal folds grow thinner, losing their elasticity and becoming stiffer, resulting in vocal changes associated with aging such as breathiness, thinning of the voice, and loss of vocal efficiency. Breathiness results of vocal cord bowing. Bowing occurs when one or both of the vocal folds become weak due to decrease of muscle fibre. Vocal bowing are in the form of a gap between the vocal folds, which prevents complete closure, resulting in hoarseness, strained voice, and breathiness.

The idea of change in pitch due to aging was first established in the late 50s [26]. For men, results suggest an increase in the mean fundamental vocal frequency (pitch) as a function of age. A trend of progressive decrease in pitch is observed from adulthood to middle-age, with a reversal of this trend into elder age. This trend is evident in the results obtained from both speech corpora. For female speakers, the fundamental vocal frequency was found to be remain consistent throughout adulthood and middle-age, with a drop in pitch occurring with older age due to menopause. This trend is evident in the results from speech corpus B.

Open Quotient in aging voice has been documented for both glottal flow and EGG studies [18]. Results show a progressive trend in OQ values with increasing age [27, 28]. However, the results obtained in this paper suggest a reversed trend, with OQ decreasing from young to elder speakers. Note that similar trends were observed for both OQ50 and OQsub50, with OQsub50 resulting in a statistically significant difference. Although the trend reported in this paper is inversed to the one previously reported, it can be concluded that the OQ trend is correlated to aging regardless of the speaker's gender. OQ measurements show a significant difference between the mean

OQ measurements for male and female speakers. As a results, we suggest that Open Quotient is independent of pitch.

The findings of this paper are in contrast to the initial assumptions. Elder speakers exhibit a deterioration in voice quality, such as breathiness, resulting in an incomplete closure of the glottis and a longer open phase. Furthermore, the results obtained from Speech Corpus A are in contrast to the EGG contact quotient findings by Bier [29]. Bier notes a decrease in contact quotient with age, which suggests breathiness of voice, with vocal folds remaining open for a longer duration due to incomplete closure.

VII. Future Work

This paper presented the OQsub50 parameter, which is the first of many other glottal pulse time-domain parameters to be implemented as part of a voice quality analysis package in *R*. It is the aim of the author to implement most other time-domain features such as Speed Quotient, Closing Quotient, Normalized Amplitude Quotient, etc. This analysis package will be made freely available as an open source application. Since the *R* voice analysis package is aimed at big data processing, future improvements will look towards reducing the algorithm's computation time using parallel computing in *R* (via the *snow* package).

VIII. Conclusion

This paper presented a new approach to glottal flow parameterization in the form of OQsub50. A hybrid of two well-established Open Quotient criteria, the OQsub50 method is introduced with detailed description on GOI and GCI extraction methods. The OQ method was performed pitch synchronously to eliminate DC bias in the glottal waveform. Quantization error were accounted for by interpolating the results. This paper presented the first of many of glottal pulse parameters to be implemented as a speech analysis package in *R*. Two speech corpora of New Zealand English were used in this study, investigating the 11 NZE monophthongs. The results suggest that OQ values are correlated to the speaker's age. We conclude that OQ values decrease with age, indicating a glottal pulse's open phase is shorter in duration for elder speakers. Although OQ values were significantly different between male and female speakers, the aforementioned trend was present in both male and female speakers, suggesting Open Quotient is independent of pitch.

References

[1] Bier, S. D., Watson, C. I., McCann, C. M., "Using the perturbation of the contact quotient of the EGG waveform to analyse age differences in adult speech", Journal of Voice, Vol. 28(3), 267-273, 2014.

[2] Sapienza, C. M., Stathopoulos, E. T., and Dromey, C., "Approximations of open quotient and speed quotient from glottal airflow and EGG waveforms: Effects of measurement criteria and sound pressure level", Journal of Voice, Vol. 12(1), 31-43, 1998.

[3] Dromey, C., Stathopoulos, E. T., Sapienza, C. M., "Glottal airflow and Electroglottographic measures of vocal function at multiple intensities", Journal of Voice, Vol. 6(1), 44-54, 1992.

[4] Drugman, T., Alku, P., Alwan, A., and Yegnanarayana, B., "Glottal source processing: From analysis to applications", Computer Speech & Language, Vol. 28(5), 1117-1138, 2014.

[5] Alku, P., "Glottal wave analysis with pitch synchronous iterative adaptive inverse filtering", Speech communication, Vol. 11(2), 109-118, 1992.

[6] Watson, C. I., Thorpe, C. W., Lu, X. B., "A comparison of two techniques that measure vocal tract shape", Acoustics Australia, Vol. 37(1), 7-11, 2009.

[7] Hollien, H., Ship, T., "Speaking fundamental frequency and chronological age in males," J. Speech Hear. Res., Vol. 15, 155-159, 1972.

[8] Sataloff, R. T., Caputo Rosen, D., Hawkshaw, M., Spiegel, J., "The three ages of voice: the aging adult voice", Journal of Voice, Vol. 11, 156-160, 1997.

[9] Pontes, P., Brasolotto, A., Behlau, M., "Glottic characteristics and voice complaint in the elderly", Journal of Voice, Vol. 19(1), 84-94, 2005.

[10] Paulsen, F., Kimpel, M., Lockemann, U. and Tillmann, B., "Effects of ageing on the insertion zones of the human vocal fold", Journal of Anatomy, Vol. 196, 41-54, 2000.

[11] Degottex, G., Kane, J., Drugman, T., Raitio, T., Scerer, S., "COVAREP - A Collaborative Voice Analysis Repository for Speech Technologies". URL: http://covarep.github.io/covarep

[12] [13] Rabiner, L., Juang, B., Fundamentals of Speech Recognition, Prentice Hall, 1993.

[13] Alku, P., Vilkman, E., and Laine, U. K., "Analysis of glottal waveform in different phonation types using the new IAIF-method", Proc. 12th Int. Congress Phonetic Sciences, Vol. 4, 362–365, 1991.

[14] Fant, G., "Acoustic Theory of Speech Production", Mouton, 1960.

[15] Rosenberg, A., "Effects of the glottal pulse shape on the quality of natural vowels", J. Acoust. Soc. Am., Vol. 49, 583-590, 1971.

[16] Fant, G., Liljencrants, J., Lin, Q., "A four-parameter model of glottal flow". STL-QPSR, Vol. 26 (4), 1-13, 1985.

[17] Fujisaki, H., Ljungqvist, M., "Proposal and evaluation of models for the glottal source waveform", ICASSP, 1605-1608, 1986.

[18] Doval, B., d'Alessandro, C., "The Spectrum of Glottal Flow Models", ACTA ACUSTICA UNITED WITH ACUSTICA, Vol. 92, 1026-1046, 2006.

[19] Drugman, T., Alwan, A., "Joint Robust Voicing Detection and Pitch Estimation Based on Residual Harmonics", Interspeech11, Firenze, Italy, 2011.

[20] R Development Core Team, "R: A language and environment for statistical computing. R Foundation for Statistical Computing", Vienna, Austria. ISBN 3-900051-07-0, URL http://www.Rproject.org/, 2011.

[21] Winkelmann, R., "The EMU Speech Database Management System (EMU-SDMS)", August 2016. URL https://ips-lmu.github.io/EMU.html

[22] Watson, C. I., "Mappings between vocal tract area functions, vocal tract resonances and speech formants for multiple speakers", Interspeech, 2014.

[23] Watson, C. I., Hui, C. T. J., "Two Short Studies in Vocal Tract Measurements", SST, 2010.

[24] Kisler, T., Schiel, F., Sloetjes, H., "Signal processing via web services: the use case WebMAUS", Proceedings Digital Humanities, Hamburg, Germany, Hamburg, 30-34, 2012.

[25] Cassidy, S., Harrington, J., "Multi-level annotation in the Emu speech database management system", Speech Communication, Vol. 33, 61-77, 2001.

[26] Mysek, E. D., "Pitch and duration characteristics of the older males", Journal of Speech and Hearing Research, Vol. 2, 46-54, 1959.

[27] Winkler, R., Sendlmeier, W., "EGG open quotient in aging voices - changes with increasing chronological age and its perception". Logoped Phoniatr Vocol., Vol. 31, 51-56, 2006.

[28] Mendoza, L. A. F., Cataldo, E., Vellasco, M., Silva, M. A., Canón, A. D. O., De Seixas, J. M., "Classification of voice aging using ANN and glottal signal parameters", Journal of Voice, Vol. 28, 532-537, 2014.

[29] Bier, S. D., Watson, C. I., McCann, C. M., "Dynamic measures of voice stability in young and old adults", Logopedics, phoniatrics, vocology, 1-11, 2016.

Precision Crystal Frequency Reference Disciplined Using Mains Power

Sam Dirks & Jonathan Scott
School of Engineering, The University of Waikato
Hamilton, New Zealand. Email: jonathanscott@ieee.org

***Abstract*—We describe a clock that keeps time and provides a reference frequency with better than 1PPM accuracy, through a quartz crystal automatically calibrated via the power line frequency. The long-term precision of line-frequency-based time is used to correct the crystal-based timekeeping system, a process referred to as disciplining. A microcontroller calibrates the internal reference dynamically with no user assistance. Precision is improved to better than 1ppm as measured against a rubidium standard. The precise signal is useful for calibrating frequency-based instrumentation or keeping time without periodic correction and costs much less than comparable alternatives. The algorithm can be implemented in any mains-powered device using a microcontroller, such as thermostats and lighting timers. The prototype clock uses a Nixie-tube display, requires no transformer, is housed a glass tube, and is offered as an open-source design.**

***Index Terms*—Clocks, frequency measurement, counting circuits, power system measurements, crystal disciplining**

I. Introduction

The properties of quartz crystals were a hot topic in the 1920s, and by 1930 the quartz crystal emerged as a technology to realise frequency-stable oscillators [1]. The second world war saw their application to radio and microwave systems, and by the 1950s quartz crystals were found at the heart of national frequency standards [2]. As the size, cost and power consumption of quartz crystal systems improved through the 20^{th} century, domestic clocks moved from using synchronous motors locked to line frequency to keep time to using internal quartz-crystal ("xtal") circuits, and wristwatches moved from escapements to low-frequency crystal oscillators. In 1950, virtually all wall clocks used mains power. In 2000, most wall clocks used crystals. The annoyance of a power cord and power cuts was traded for the annoyance of changing batteries every year or so.

Clocks in the 21^{st} century typically use one of three sources to keep time: the line frequency, an internal quartz-crystal, or a remote source such as the internet, cellular-phone system, or positioning satellites. The latter remote alternative has the advantage that the time is collected from the remote site and thus the clock never needs to be set, but the disadvantage of depending upon access to an external signal, and usually a need for rechargeable batteries. As of the second decade of this century the crystal is the dominant alternative, although increasing numbers of people, especially teenagers, eschew watches as they carry smart phones. All of these offer better timekeeping than most users need, but ironically the crystal is horologically the weakest, as it needs individual calibration in manufacture.

Quartz crystals are manufactured with something like 100ppm accuracy, or 20ppm accuracy if used with exactly the prescribed loading capacitance. This corresponds to one or more minutes/month or 10 minutes to an hour per year. Wristwatches achieve about 1ppm accuracy with factory calibration, typically via firmware or pulse-based correction circuits, and with careful design expect to achieve 2–5ppm over the life of the product [3]. Oven-controlled quartz-crystal standards are today available with 1ppm precision, but these tend to be quite expensive and power hungry, more than the low-cost instruments with which you might use them.

This manuscript describes a simple, low-cost scheme to "discipline" a quartz-crystal oscillator to automatically calibrate it to better than 1ppm accuracy using the mains power line frequency as the higher-precision reference. The attraction of this approach is that it comes at virtually no extra cost to any system that uses mains power, has a crystal for timekeeping and that is controlled with a microprocessor, as do many building systems and appliances. It would be equally possible to use an external, radio-based, time reference, if that is already available, but these are less common. The idea of using an external time reference to gradually refine a dynamic calibration for a crystal oscillator is not new [4]. Nevertheless, this particular work employs that principle to simultaneously produce an elegant timepiece and a laboratory-grade frequency reference in a very low-cost, robust, open-source, and rather elegant, design.

II. Precision

Line power cumulative frequency time error is tightly controlled in New Zealand, with the Electricity Commission "ensuring the cumulative frequency time error does not exceed 5 seconds and that the error is eliminated at least once per day" [6]. The precision with which a frequency can be determined, and a calibration factor calculated, increases with the length of the observation. Five seconds is one-millionth part of 58 days, so the cumulative-frequency error of a mains power observation will approach 1ppm only after a two-month observation. Mathematically

$$M = T + \delta \tag{1}$$

where M is time according to the mains frequency, T is the true time, and δ is the maximum error, 5 seconds in our case, while

$$X = T \times (1 + \epsilon) \tag{2}$$

where X is time according to the crystal to be calibrated, and ϵ is the error to be determined and corrected. We seek an estimate of ϵ given bounds on δ. Some algebra leads to

$$\epsilon = \frac{X - M + \delta}{T} \approx \frac{X - M}{M} \tag{3}$$

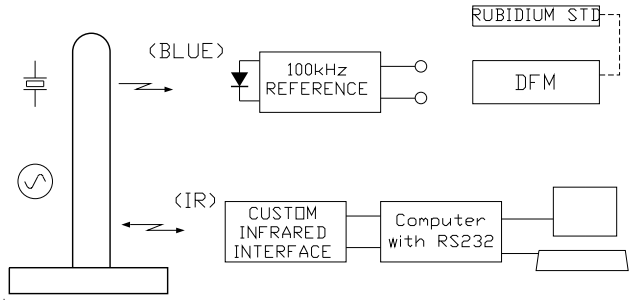

Fig. 1. Block diagram of the system as implemented. The rubidium standard is used to verify performance and is eventually replaced by the device itself.

if δ is small, M large, and temperature variations are negligible. This provides a simple path to obtain an estimate of ϵ. Knowing that $|\delta| < 5$s, the uncertainty in an estimate is δ/M, corresponding to about 1ppm when $M = 5 \times 10^6$ or the 58 days noted above.

We know also that $M = T$ at least once per day. It may thus be possible to reduce the uncertainty. If the diurnal variation, δ_{max} is less than 5 seconds, then $\delta \leq \delta_{max}$ in that period since it must have been zero somewhere in the day, and the uncertainty is reduced. It was observed from clock performance that the convergence rate is greater than expected from theory.

III. Hardware

Figure 1 depicts the system arrangement. The basic unit incorporates a crystal and mains-power monitoring in a free-standing device that functions as a clock. Control of the clock is provided through a custom infrared (IR) serial data interface to a Raspberry Pi computer, although it can function perfectly well without this interface. A reference signal is provided for the calibration of instrumentation. For our work here a rubidium standard frequency generator is interfaced with a digital synthesizer to verify performance, but a Digital Frequency Meter (DFM) will also work provided it has sufficient digits of resolution.

The clock hardware requires a microcontroller with both a crystal oscillator and access to line power frequency. We have chosen to use a PIC16F1709 because it offers an in-built, current-based, zero-crossing detector peripheral, serial communications support, plenty of memory, crystal oscillator, temperature measurement, non-volatile storage, and all the usual features of modern 8-bit microcontrollers. It retails locally for around $2 in small quantities. Power consumption is around 350μA when clocking at 1 Mips and with all peripherals shut down to conserve power. We have based our design on a very elegant clock designed by Dekker and de Graaf [5]. Their clock operates directly on the mains voltage, requiring no transformer, and displays time using a single Nixie-style neon display tube. For safety and novelty the original is mounted inside a one-inch-diameter (25mm) glass test-tube.

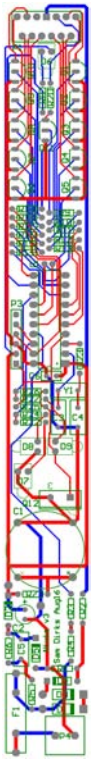

Fig. 2. The PCB layout of the hardware

For our design, a supercapacitor to allow operation without mains power for approximately 2.5 hours, and optical interfaces for exchange of data and for the delivery of the frequency standard signal. The circuit diagram is reproduced in figure 3 with the PCB layout in figure 2. We also house the device in a glass tube extended to 300mm height and 30mm outside diameter.

IV. Firmware

The PIC16F1709 carries out the following functions:

- Count mains cycles since the first zero crossing of a contiguous sequence
- Count processor clock cycles since the first zero crossing of a contiguous sequence
- Detect missed mains cycles/loss of mains power
- Compute correction factors for crystal from cycle counts
- Calculate current time given clock cycles, correction factor, and uncertainty
- Display time and key data on display tube sequentially by fading digits in and out
- Receive commands and supply data via IR serial link
- Generate reference frequency in blue light

Correction of the crystal clock frequency occurs in increments of approximately 1 part in 60 million, or 0.017ppm. This is achieved in the firmware by dynamically adjusting the number of nominal-1μs computer clock cycles in each minute. To

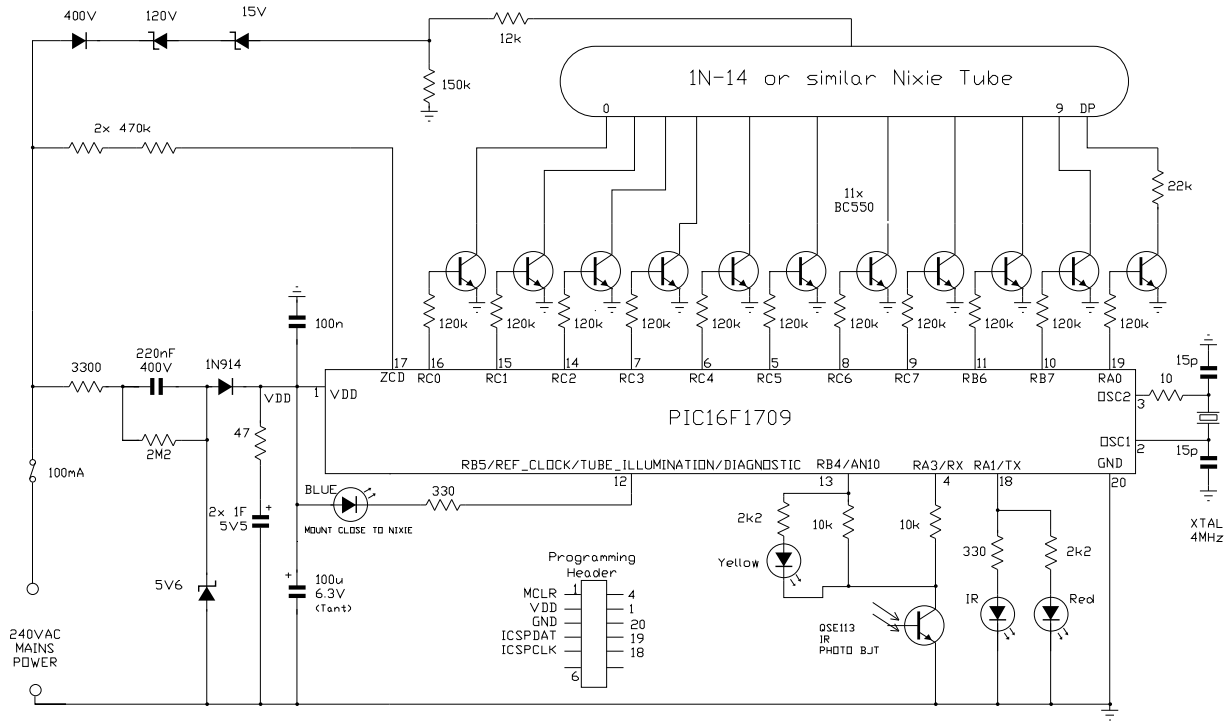

Fig. 3. The circuit schematic of the hardware.

Fig. 4. The clock/frequency reference inside its glass-tube housing. The tube OD is 30mm and height is 300mm including the base. The single IN-14 display tube is at the top.

allow counting of μs over the time frames required the code implements custom 64-bit integers as they are not supported my the Microchip compiler.

V. Measurement Method

Several prototypes were constructed so that work could continue while longer-duration experiments were carried out. The microcontrollers were programmed in C using Microchip's XC8 compiler. For frequency reference out from the clocks a blue LED was flashed using hardware PWM. This was configured to such that the PWM frequency is precisely $1/40$ of the crystal frequency driving the system clock. Therefore if the crystal is exactly $4\,\mathrm{MHz}$ the led will flash at $100\,\mathrm{kHz}$. For measurement an optical reader was build to produce a square wave from the flashing of the LED. A photograph of a prototype with the enclosing tube mounted in a solid nylon base appears in figure 4.

Frequency was measured by phase matching the frequency reference LED on the clock and the output of a 12-digit function generator using a $10.000\,000\,00\,\mathrm{MHz}$ rubidium frequency standard as reference. This was done by manually adjusting the frequency of the function generator until 9-digits of precision had been reached minimising phase shift. This method was used as an available digital frequency meter was only able to achieve 8 digits of precision. Typically 8 digits would be enough to calibrate instrumentation against these clocks.

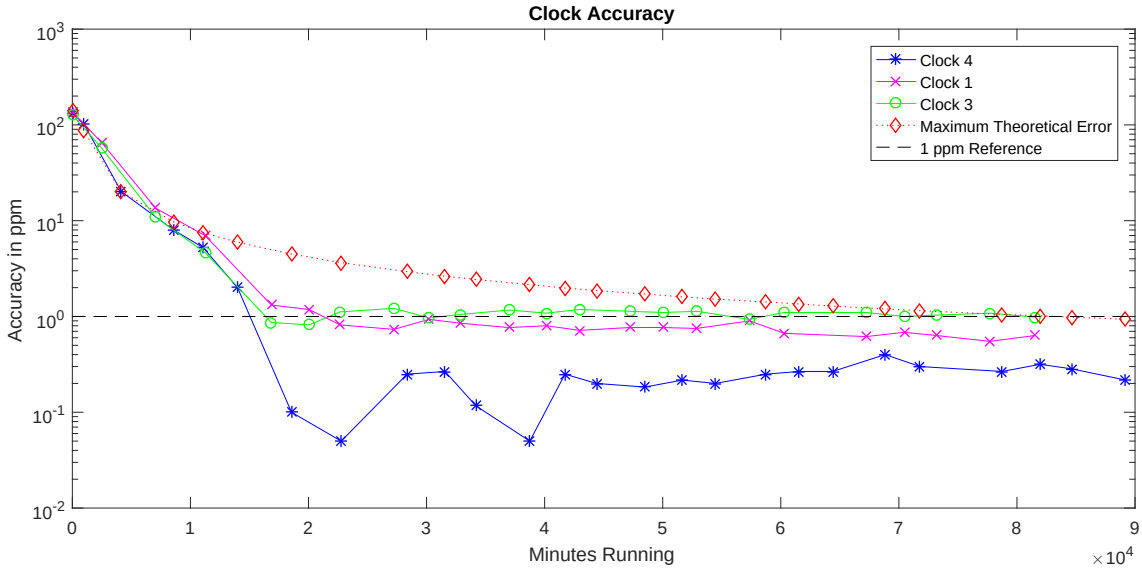

Fig. 5. The accuracy of the clocks with respect to the time running in minutes. For reference the maximum theoretical error and 1 ppm lines are included.

VI. Measured Results

The accuracy calculation was made by comparing the crystal frequency calculated from LED frequency measured and the crystal frequency as calculated by the clock. The frequency of the LED for the clocks were 100.013 653 kHz, 100.012 747 kHz and 100.012 351 kHz respectively for clocks 1, 3 and 4 respectively.

Figure 5 shows the accuracy of the clocks with respect to time. This clearly illustrates the clock 4 achieving an accuracy of 1 ppm after only 16,000 minutes and reaching an equilibrium of approximately 0.3 ppm. Clocks 1 and 3 reached approximately 1 ppm in 17,000 minutes. However they did not reach the same level of accuracy as clock 4.

Clock 4 was run on an open air bench in the laboratory against an internal wall with the only building ventilation controlling room temperature. Clocks 1 and 3 were placed in a windowsill in another room and received direct sunlight for part of the day. As they both had a higher noise floor of similar magnitude we conclude that environmental factors are the limiting case.

The time taken to achieve 1 ppm for clock 4 of 16,000 minutes or 11 days also shows that the regulation of mains frequency is much better than the minimum requirements which would require 58 days.

VII. Conclusion

The precision mandated for power-line frequency has been used to provide a very accurate frequency reference through firmware discipline of a crystal oscillator. This is achieved on a very low-cost platform.

A GPS or GSM time fix could provide a 5x reduction in the time required to achieve maximum accuracy in exchange for access to an RF signal, more circuitry, and cost. Nevertheless, the design presented here can provide a laboratory-grade signal of well-known frequency, or a low-frequency signal of precisely-known period. Automatic correction for ageing and environment could be included in a future development. The devices look attractive and tell the time without either primary or rechargeable batteries or resetting of the time with modest-duration power outages.

Acknowledgement

The authors wish to thank dos4ever.com for their elegant open-source hardware design that was used as a starting point for the clock design. We thank Alan Smith, Michael Hoogeven and Martin Gore in the Waikato School of Engineering Workshop for mechanical support, Nicole Varstraten for initial PCB work, Benson Chang for excellent assistance with hardware, William Redman-White for useful discussions and Steve Newcombe for custom glassware.

References

[1] Lucas, H.J., "Some developments of the piezo-electric crystal as a frequency standard", *Journal of the Institution of Electrical Engineers*, Vol. 68, no. 403, 1930, pp855 –867.

[2] Shaull, J.M., and J.H. Shoaf, "Precision Quartz Resonator Frequency Standards" *Proceedings of the IRE*, Vol 42, no. 8, 1954, pp1300–1306.

[3] Lanfranchi, D., E. Dijkstra, and D. Aebischer, "Microprocessor-based analog wristwatch chip with 3 seconds/year accuracy", *Digest of Technical Papers of the 41st ISSCC, IEEE International Solid-State Circuits Conference*, 1994, pp92–93.

[4] Helsby, N.C., "GPS disciplined offset-frequency quartz oscillator", *Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum*, pp435–439, 2003.

[5] Ronald Dekker, and Pascal de Graaf, "Tube-in-a-Tube Clock", see http://www.dos4ever.com/TiT/TiT.html, retrieved December 2014.

[6] Electricity Commission of New Zealand, "Current Arrangements for Frequency, Voltage, Reliability and Security", Chapter 3, 461406-2, accessed December 2014.

Wireless Treetap: a new electronic device for time-of-flight tree-stiffness measurement.

Michael Frampton
School of Electrical and
Computer Engineering
University of Canterbury
New Zealand, 8025
Email: michael.frampton@pg.canterbury.ac.nz

Michael Hayes
School of Electrical and
Computer Engineering
University of Canterbury
New Zealand, 8025
Email: michael.hayes@canterbury.ac.nz

***Abstract*—Treetap is a time of flight (ToF) acoustic tree-stiffness measurement system developed at the University of Canterbury. Wireless Treetap (WTT) is the next iteration of this project. The goal of the project is to create a smartphone based system for measuring wood stiffness. The implemented system can be described as a portable wireless sensor network (WSN) of accelerometers with an emphasis on high time accuracy. A prototype of the system has been built and a number of measurements performed including: the drift rate of the oscillators, the accuracy of the time synchronisation algorithm, the noise power of the measurement front-end and the battery performance of the device.**

I. Introduction

Foresters are interested in quantifying wood quality. One useful metric of quality is the Modulus of Elasticity (MoE) (E), also known as stiffness. Stiffness is generally defined as the ratio of stress (σ) to strain (ϵ) along an axis, $E = \sigma/\epsilon$. Acoustic tools are well established for the purpose of estimating the stiffness of wood [1], [2], [3]. More specifically, Time of flight (ToF) acoustic devices have been used for a number of years to estimate the stiffness of standing trees [4], [5], [6]. These devices estimate the stiffness by measuring acoustic wave speed within a tree. The relationship between stiffness and speed for a plane-wave propagating in an isotropic lossless-medium is given by

$$E = c^2\rho, \tag{1}$$

where c is the wave speed and ρ is the density of the medium.

The ToF method measures the time taken for an acoustic wave to travel some distance z through a tree stem (typically 1m). There are two similar variants of the ToF method. The first uses separate driving and receiving transducers [7], [8]. The propagation time τ is the time from wave transmission to reception. The second variant uses two receiving probes, with the wave being excited by striking a metal spike in the tree with a hammer [9], [10]. τ is then the time it takes for the wave to propagate between the two transducers. In either case, the speed $c = z/\tau$ is calculated. The density of the wood is then measured or estimated, which gives the MoE using (1).

Treetap is the name of a range of ToF devices developed at the University of Canterbury as a collaboration between Electrical and Computer Engineering (ECE) and the New Zealand School of Forestry (NZSoF). NZSoF desires to have access to a high quality, robust ToF measurement platform for its research purposes. Treetap systems have been in development for over 10 years [9]; the current system suffers from several problems. The currently used system is fragile, particularly where the wired probes connect to the device. The device uses a threshold algorithm, which is prone to errors if signal strength is weak. The old device only has a minimal LCD display which is not intuitive to use. In addition, it is desired to update the device to make use of smartphone technology. A ToF device using smartphone technology could make use of a number of bundled features, including: a high-resolution display, touchscreen input, geo-tagging of data via GPS, time-stamping of data and networked upload of data.

The initial work on this system was conducted as part of a final-year undergraduate engineering project at the University of Canterbury [11]. This paper describes the continuation of the undergraduate work, the design and manufacture of a new wireless ToF system. The project is summarised in the following sections. Section II describes the specifications of this device. Section III describes the high-level system design of the new tool. Section IV outlines the hardware and algorithm employed for time-synchronising multiple probes, a necessary feature for operation of the system. Section V provides results of the measurement performance of the device. Section VI describes the development of a smartphone application for use in the system. Section VII provides measured data on the battery performance of the device. Finally, conclusions are provided in Section VIII.

II. Specifications

Based on knowledge gained from previous Treetap versions (outlined in Section I), the following specifications for the new device were devised.

- The system will use the two receiving-probes topology outlined in Section I.
- The system will consist of two separate wireless probe units. Each containing an acoustic transducer, measurement electronics, wireless communication and a Li-Po battery.

- The two probes will connect to a smartphone application via Bluetooth.
- Each probe will have an independent clock. Thus the probes must periodically synchronise to account for clock-drift. This is necessary to assign a common time reference to measurements. The system should be capable of measuring an acoustic delay between the two probes to an accuracy of 1 μs.
- Each probe should be capable of running for full working day (8 hours) without being recharged. The probes will enter a low power mode after being idle for some time.

III. System design

Several aspects influenced the choice of components used to achieve the specifications described in Section II. The previous Treetap iteration used an Atmel SAM7s ARM microcontroller. The next iteration of the processor, the SAM4s was selected. Many of the peripherals on the SAM4s are similar to the SAM7s, allowing code reuse. The analogue front-end was adapted from that used in the previous Treetap iteration. The previous version included two receiver channels and one transmit channel. The new system only requires a single receive channel per probe. The Analogue to Digital Converter (ADC) and Programmable Gain Amplifier (PGA) components were updated to use more modern variants. Initially a Global Positioning System (GPS) receiver was considered for use in synchronising multiple probes. This would make use of the highly accurate pulse-per-second output for synchronisation which are derived from atomic clocks in the satellite. This feature is commonly included in GPS receivers. However, initial prototyping revealed that the selected receiver could not obtain a satellite fix under forest canopy. A previously established fix would usually be maintained, however this was deemed unacceptable for field work. Instead, a low power radio was included in the design specifically for time-synchronising devices. A low cost Bluetooth transceiver module was included for communication with the smartphone. Lithium-polymer (LiPo) batteries were selected as the best option. An integrated LiPo battery charger was included in the system for charging the device via Universal Serial Bus (USB). The overall system topology is shown in Figure 1. A block diagram of the subsystems implemented in each probe is shown in Figure 2.

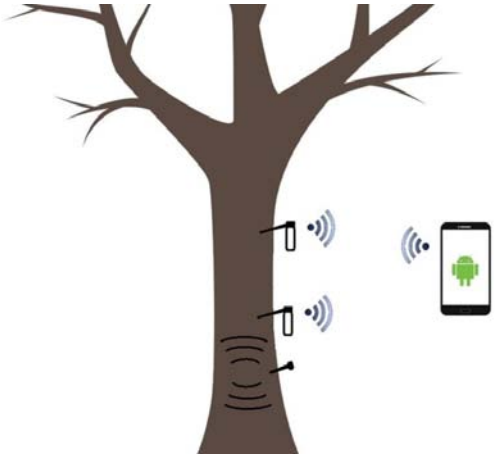

Fig. 1. Wireless Treetap consists of two independent receiver probes and a smartphone, which controls the system and provides an interface for the operator.

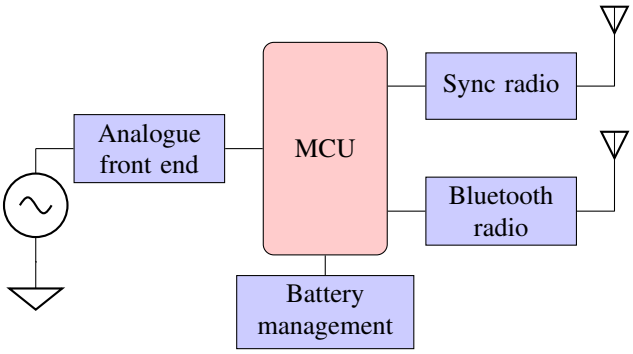

Fig. 2. Each probe device consists of several different subsystems.

IV. Clock synchronisation

In contrast to previous Treetap versions, each Wireless Treetap (WTT) unit is now an independent measurement device. When capturing a measurement, each device independently samples its transducer until a stress-wave is detected. When a signal is detected the device references the capture time using its own internal clock. Each device has an independent clock, thus it is necessary to project the captured times to a common clock.

A. Wireless clock-synchronisation

Clock synchronisation has been extensively studied in Wireless sensor network (WSN) applications [12]. In a WSN each sensor node has its own independent clock. The clock's value is expressed as $C(t)$, where t is the independent (and unknown) reference time. Typically the clock in a low cost electronic sensor is derived from a quartz crystal-oscillator tuned to a particular frequency. The manufacturing error in the tuned frequency is specified in parts per million (ppm) and is typically approximately 40 ppm [13].

A simple model for clocks in a WSN assumes that each clock will run at a different rate, but the rate that each runs at is constant over a period of time. So the $i\text{th}$ node's clock is given by

$$C_i(t) = \theta_i + f_i t, \tag{2}$$

where θ_i is the clock's offset and f_i is the clock's rate (skew). In practice, the clock rate is not constant and is influenced by environmental factors including temperature, humidity, pressure, vibration, and electromagnetic fields [14], [15]. To compensate for the changing clock rate, the sensor network must periodically synchronise.

Suppose the goal is to synchronise two nodes A and B. This requires the two nodes to periodically record their clock values simultaneously. Node A records its clock value, then sends a message to B, which then records its clock value. There will be some delay in transmitting the message to the receiving node.

The delay can be decomposed into the following categories [12].

- *Send time* - Time spent constructing the radio message in the transmitter.
- *Access time* - Time waiting for accessing the channel, depends on multiple factors including the protocol used and the congestion of the channel.
- *Transmission time* - The time to transmit the message, depends on the bit rate and the protocol being used.
- *Propagation time* - Time for the electromagnetic (EM) wave to propagate to the receiver.
- *Reception time* - Time for the message to be received. This is the same as the transmission time.
- *Receive time* - Time to process the received message and signal the application processor.

These components can be classified into deterministic and non-deterministic delays. The majority of these delays are deterministic [12]. The greatest uncertainty arises in the access time, which depends on other devices operating on the channel. Another non-deterministic delay is present in the receive time. This is due to the unknown bit-alignment of the message. The message contains a synchronisation word at the start. The receiver must then align the received bytes to match the expected pattern.

B. The flooding time synchronisation protocol

Tackett et al. [16] demonstrated success in implementing a WSN using the Nordic nRF24L01+ radio. Their implementation attempted to synchronise a network of clocks running at a relatively low rate of 32.768 kHz. It was decided to attempt to emulate this approach.

Tackett et al. used the Flooding Time Synchronisation Protocol (FTSP) [17] for synchronising a network of nodes in a traffic monitoring application. FTSP is an example of a one-way message dissemination protocol. In a one-way messaging system a single master node periodically broadcasts its time to all slave nodes in the network. The clock value of the $k\text{th}$ broadcast message received at the $i\text{th}$ slave can be modelled as [12],

$$C_{i,k} = f_i \left(C_{1,k} + \tau_{\text{rf}} + X_k \right) + \theta_i, \tag{3}$$

where $C_{1,k}$ denotes the clock value of the master at the $k\text{th}$ broadcast time, τ_{rf} is the deterministic component of the radio's send-receive delay and X is the random component of the radio's delay.

As discussed in Section IV-A, quartz crystal oscillators are accurate to just a few 10s of ppm, $f_i \approx 1$. (3) demonstrates that it is not possible to differentiate between the clock offset (θ) and the radio delay ($\tau_{\text{rf}} + X_k$). In Tackett et al.'s case, the radio delay was not a significant component in the overall delay so this component was discounted, this is not the case for WTT.

Tackett et al. use a modified version of the FTSP protocol to calculate clock skew and offset. In this protocol, the master periodically broadcasts synchronisation messages containing the previous transmit time. The nRF24L01+ radio employed has an Interrupt Request (IRQ) pin which can be configured to assert on certain events. The IRQ pin is used to latch the value of a timer when the master transmits. It is necessary for the master to send the time of the previous message because the timer is latched immediately after the message is sent. Messages contain a message ID in addition to the previous master-message time. The message ID enables the slave device to determine whether any messages were dropped in transit. On the slave side, the nRF24L01+ is configured to latch the timer when a message is received. In this way, the slave constructs a table of master-slave times, with an entry for each transmit-receive event.

In Tackett et al.'s implementation; the slave performs linear regression on the slave's table to calculate the relative skew and offset f and θ for each slave.

C. Radio delay experiment

As described in Section IV-B, previous implementations of the FTSP disregard the radio transmit-receive delay component. In an attempt to improve performance over previous implementations, an experiment was performed to measure the delay of the nRF24L01+ radio. Two prototype WTT boards were used, one configured as the master, transmitting. The other as the slave, receiving. The master was configured with two timer-counter (TC) modules running synchronously, i.e. each TC has the same counter value at any time. The first TC was configured to latch its counter when the IRQ line from the master's radio asserts. This occurs when the transmit is complete. The second TC module was set to latch when the IRQ line on the slave's radio asserts, on reception. The slave's IRQ line was connected to the master board via a short wire. By subtracting one latched timer from the other, the transmit-to-receive delay was measured. The timer values were sent to a PC via USB. The experiment was then repeated, modifying the radio's data rate.

The results of the experiment are shown in Figure 3. At each data rate, the delay appears to be split into two distributions. The separation between these two distributions corresponds to the period of one bit at that rate i.e. $\Delta\tau_{1\text{mbps}} = 1\mu s$, $\Delta\tau_{2\text{mbps}} = 0.5\mu s$. This one-period delay may be caused by the preamble bit-alignment delay described in Section IV-A. The mean and standard-deviations of the data is shown in Table I.

TABLE I
MEAN AND STANDARD-DEVIATION OF THE TRANSMIT-RECEIVE DELAYS OF FIGURE 3, ROUNDED TO 2 SIGNIFICANT FIGURES.

	Mean (μs)	Std-dev (μs)
1 mbps	6.32	0.47
2 mbps	5.57	0.22

D. Clock rate experiment

Section IV-A discusses the importance of clock accuracy in WSN applications. To address this requirement, an experiment

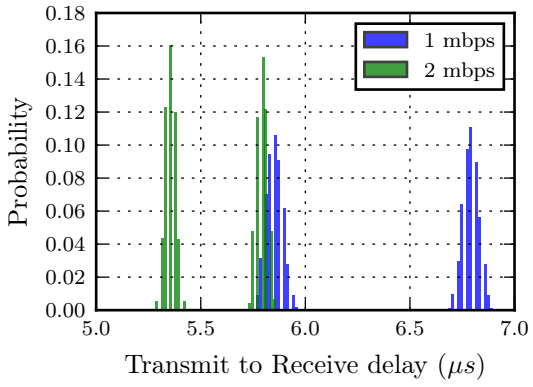

Fig. 3. Histograms showing the transmit-to-receive delay of the nRF24L01+ radio. The two WTT boards were placed parallel to each other with a separation of 10 cm. The delay is dependent on the data rate selected. The bin width of these plots is equal to the period of the clock $t = 1/48$ MHz.

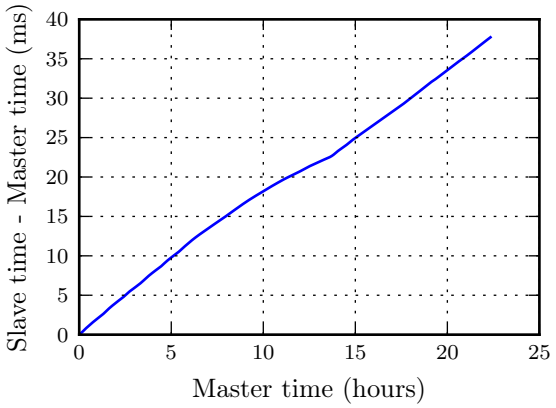

Fig. 4. The difference between the two devices plotted over the length of the experiment. The gradient indicates the relative clock drift, in this case $f_{1,2} = 0.469\,\text{ppm}$.

was performed to measure the relative clock-rates of two oscillators on prototype WTT devices.

One of the devices was configured as the master node, the other configured as the slave. Both the master and slave were set up with a free running timer which was set to latch its value when a pin is asserted. The master was set up to periodically assert a pin which was connected to the timer input pins on both the master and the slave. When this event occurred, both the master and slave timer was printed via USB to a console on a PC for later analysis.

The experiment was performed over a period of 22 hours. The timer was captured once per second. Figure 4 illustrates the difference between the two clocks over time. The measured clock rate between the two devices is significantly lower than that predicted by the oscillator specifications [18].

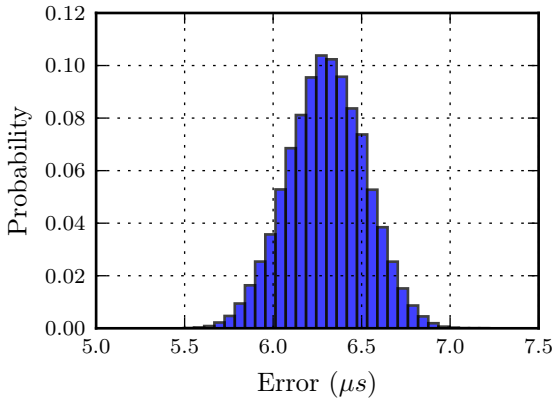

Fig. 5. The error between the slave's estimated time and measured time. The synchronisation algorithm employed a table size $N = 20$, captured at synchronisation interval $\tau = 2\,s$. The mean and std-deviation rounded to two S.F. are 6.30 μs and 0.21 μs, respectively.

E. Synchronisation parameters experiment

Using prototype devices, a trial implementation of the synchronisation algorithm described in Section IV-B was built. The devices were configured according to the synchronisation algorithm described in Section IV-B. One device was configured as the master, the other as the slave. As in the previous experiments, a wire connecting the two devices was used to periodically latch the actual clock values. This allows the slave time as estimated by the synchronisation algorithm to be compared against the actual time i.e. the synchronisation error. The synchronisation table and the latched timer values were dumped to a PC via USB. The regression calculation was performed on the PC. This experiment was performed over a period of 20 hours. Figure 5 shows the error between the master time and the time estimated using the synchronisation algorithm. Note that the transmit-receive delay (Table I) has not been subtracted from the data in this experiment. If this were conducted, the distribution would have zero-mean.

The interval in seconds between synchronisation messages it denoted τ. The size of the table is denoted N. Thus the time required to fully populate the table is given by $T = N\tau$. It is of interest to minimise T because a long synchronisation time causes unnecessary delays for the user. Also, a longer time synchronising will cause a greater drain on the system's battery. The effect of different table sizes was examined, the results of which are shown in Figure 6.

V. Measurement Performance

An experiment was performed to evaluate the performance of the analogue front-end and ADC. The device was set up to continuously capture data with no signal at its input. The Power Spectral Density (PSD) of the captured data is shown in Figure 7. The PSD contains a significant peak at 100 kHz, the origin of this signal is unknown at the time of writing. Figure 8 shows the histogram of the same captured data. The histogram shows that the captured signal was rarely more than

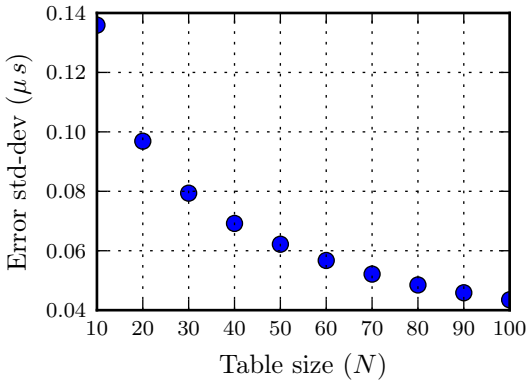

Fig. 6. The Standard deviation of the error plotted against table size N. The synchronisation was performed with an interval of $\tau = 100\,ms$

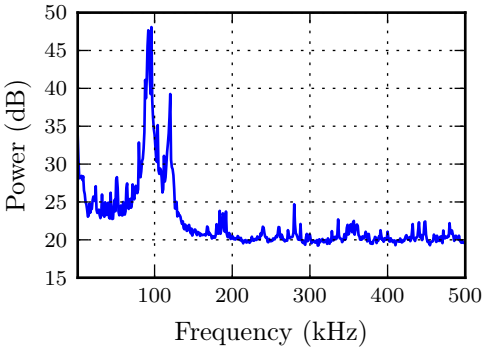

Fig. 7. PSD of signal captured with nothing connected at the input of the ADC.

two ADC counts from the mean, indicating that the amplitude of the noise is small.

VI. Smartphone application development

An application (app) has been developed to run on an Android device. The app performs the following functions:

- Manage Bluetooth connections to probe devices.
- Control the process of initiating, recording and synchronising a measurement.
- Provide an interface for managing recorded data.

The Android Application Programming Interface (API) [19] uses Socket objects to handle Bluetooth connections. In the app these Sockets persist even if the probe has been switched off. This allows a connection to the probe to be reopened when the operator wakes it up again.

The app is responsible for converting the slave's signal into master time using the synchronisation table described in Section IV-B. The regression component of the synchronisation algorithm described in Section IV is implemented in the app. The app can then calculate the time delay between the two signals. The app does not calculate the stiffness as in (1), only the time delay τ is returned to the user.

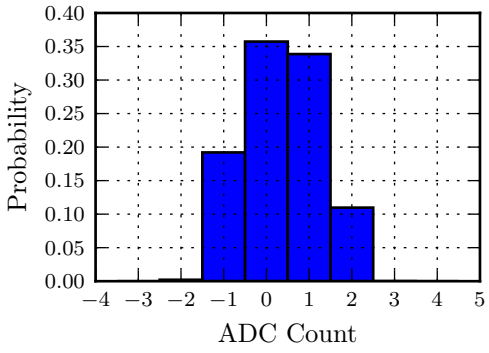

Fig. 8. Histogram of the signal captured with nothing connected at the input of the ADC.

VII. Battery performance

The battery voltage was monitored while the battery was discharged. This gives an idea of the expected run time of the device before a recharge is required. The discharge curve for three battery instances are shown in Figure 9. The device was discharged at its normal operating current while sampling, which is approximately 60 mA. The three batteries follow a nearly identical discharge curve at this discharge current. The batteries used were Turnigy 300 mAh Li-Po.

The probe will operate down to a minimum battery voltage of 3.3 V, thus the estimated run time of the device is approximately 3.5 hours when operating full time. The measured capacity in milliamp-hours (mAh) is given by

$$\begin{aligned} \text{capacity} &= \text{rate} \times \text{hours} \\ &= 60 \times 3.5 \\ &= 210\ \text{mAh}. \end{aligned}$$

It is evident that at a discharge rate of 60 mA the measured capacity is significantly lower than the rated capacity. The device was designed to allow a second battery to be connected in parallel to increase the capacity.

The device is designed to switch to a low-power sleep mode if it receives no commands for a certain amount of time. The device may be awoken by striking the transducer, which triggers an interrupt on the microcontroller. In lower power mode all peripheral devices on the printed circuit board (PCB) are disabled and the microcontroller is placed in a low power mode. The current draw of the device in low power mode was measured at approximately 150 μA. It is expected that the device will spend a significant amount of time in low power mode, allowing the device to meet the battery specification. The actual battery performance will be determined through field trials at a later date.

VIII. Conclusion

This paper describes the implementation of a new wireless time-of-flight tree-stiffness measurement device. A prototype device has been developed and several measurements performed to evaluate its performance.

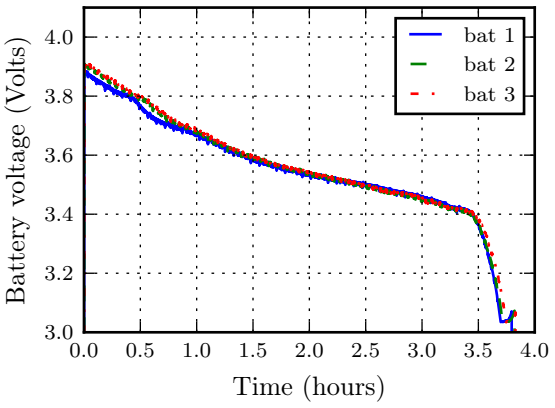

Fig. 9. The voltage is plotted as the battery is discharged, for two different batteries of the same model. The battery tested was a Turnigy 300 mAh. The device was discharged at approximately 60 mA.

The time-synchronisation algorithm implemented in the system is capable of determining the clock skew and offset between two devices to an accuracy of approximately 1 μs. One weakness of the analysis is that no absolute time reference was used in determining the synchronisation accuracy. It is possible that both oscillators have a significant common-mode error. However, this is unlikely as the maximum frequency accuracy specified in the datasheet is $|\Delta f/f_0| = 2$ ppm. There may be room for further improving the synchronisation algorithm. The algorithm subtracts the mean transmit-receive delay (Table I). It may be possible to achieve a more effective synchronisation system by better utilising established knowledge of the radio's delay distribution (Figure 3).

An Android smartphone app was developed for operating the system. The current version of the application rounds the calculated delay to the nearest microsecond. This adds an additional error to the final calculated delay. A future version may perform subsample interpolation to achieve a more accurate delay.

The battery performance of the system was measured, though it has not been shown conclusively that the devices can operate for an 8 hour working day. Field testing will be conducted in the near future to further evaluate the device.

Acknowledgment

The authors would like to thank the NZ School of Forestry for funding this research. Technical assistance provided by Michael Cusdin and Nigel Pink has been invaluable in the execution of this project. Finally, thanks are extended to Thomas Harrison, Daniel Lower, Thomas Potter and Matthew Edwards for conducting the initial research and design work for the device.

References

[1] M. Legg and S. Bradley, "Measurement of stiffness of standing trees and felled logs using acoustics: A review," *The Journal of the Acoustical Society of America*, vol. 139, no. 2, pp. 588–604, 2016.

[2] V. Bucur, *Acoustics of Wood*. Berlin, Heidelberg: Springer-Verlag, 2006.

[3] F. C. Beall, "Overview of the use of ultrasonic technologies in research on wood properties," *Wood Science and Technology*, vol. 36, no. 3, pp. 197–212, 2002.

[4] C. Wessels, F. Malan, and T. Rypstra, "A review of measurement methods used on standing trees for the prediction of some mechanical properties of timber." *European Journal of Forestry Research*, vol. 130, no. 6, pp. 881–893, 2011.

[5] X. Wang, R. Ross, M. McClellan, J. R. Barbour, J. R. Erickson, J. W. Forsman, and G. D. McGinnis, *Strength and stiffness assessment of standing trees using a nondestructive stress wave technique*. United States Department of Agriculture, Forest Service, Forest Products Laboratory, 2000.

[6] X. Wang, R. J. Ross, M. McClellan, R. J. Barbour, J. R. Erickson, J. W. Forsman, and G. D. McGinnis, "Nondestructive evaluation of standing trees with a stress wave method," *Wood and fiber science*, vol. 33, no. 4, pp. 522–533, 2007.

[7] X. Wang, R. J. Ross, and P. Carter, "Acoustic evaluation of wood quality in standing trees. part i. acoustic wave behavior," *Wood and Fiber Science*, vol. 39, no. 1, pp. 28–38, 2007.

[8] P. Carter, S. Chauhan, and J. Walker, "Sorting logs and lumber for stiffness using director hm200," *Wood and Fiber Science*, vol. 38, no. 1, pp. 49–54, 2007.

[9] S. Woods, "Acoustic inspection of timber," Master's thesis, University of Canterbury, 2006.

[10] M. J. Toulmin and C. Raymond, "Developing a sampling strategy for measuring acoustic velocity in standing pinus radiata using the treetap time of flight tool," *New Zealand Journal of Forestry Science*, vol. 37, no. 1, p. 96, 2007.

[11] T. Harrison, D. Lower, and T. Potter, "Android based tree stiffness measurement system: Wireless treetap," University of Canterbury, Tech. Rep., 2014.

[12] Y.-C. Wu, Q. Chaudhari, and E. Serpedin, "Clock synchronization of wireless sensor networks," *Signal Processing Magazine, IEEE*, vol. 28, no. 1, pp. 124–138, 2011.

[13] P. Horowitz and W. Hill, *The Art of Electronics*, 3rd ed. 32 Avenue of the Americas, New York, NY 10013-2473, USA: Cambridge University Press, 2015.

[14] F. L. Walls, "Environmental sensitivities of quartz crystal oscillators," in *Proc. 22nd Ann. Precise Time and Time Interval (PTTI) Applications and Planning Meeting*. DTIC Document, 1990, pp. 465–477.

[15] D. W. Allan, N. Ashby, and C. C. Hodge, "The science of timekeeping," Hewlett Packard Corporation, Tech. Rep., 1997, application Note 1289.

[16] N. Tackett, E. Jovanov, and A. Milenković, "An implementation of time synchronization in low-power wireless sensor networks," in *2011 IEEE 43rd Southeastern Symposium on System Theory*. IEEE, 2011, pp. 61–66.

[17] M. Maróti, B. Kusy, G. Simon, and Á. Lédeczi, "The flooding time synchronization protocol," in *Proceedings of the 2nd international conference on Embedded networked sensor systems*. ACM, 2004, pp. 39–49.

[18] *ASTX H11 TCXO Datasheet*, Abracon LLC, 5 2015.

[19] G. Inc, "Android developers," September 2016. [Online]. Available: https://developer.android.com

Semi-automated electrical circuit drawing with Lcapy

M. P. Hayes
Department of Electrical and
Computer Engineering
University of Canterbury
Christchurch, New Zealand
Email: michael.hayes@canterbury.ac.nz

***Abstract*—Lcapy is an open-source Python package developed by the author for teaching of linear circuits. The analysis is performed symbolically in the Laplace domain using the Python SymPy package. An important part of Lcapy is generating annotated schematics. This is achieved using a semi-automated algorithm with drawing hints specified for each circuit element. The drawing is performed using CircuiTikZ for text-book quality schematics.**

I. Introduction

The goal of the Lcapy project is to provide an open-source Python [1] package for linear circuit analysis. This package was designed and implemented by the author to allow students to experiment with linear circuits and, as a bonus, check their homework! Python was chosen since it is platform independent, can be used with a command-line interface or an IDE, and has excellent packages for plotting (matplotlib), symbolic analysis (SymPy [2]), and numerical analysis (NumPy, SciPy).

Lcapy analyses networks using symbolic analysis in the s-domain for transient analysis and with phasors for AC analysis [3]. Networks are created in Lcapy by connecting components in two ways: using parallel and series operators; and by netlists using a SPICE-like syntax [4]. Transformations can also be applied to networks to provide simplifications. Since Lcapy was designed primarily as a teaching tool, to demonstrate the results of a transformation it was essential that Lcapy had a way of displaying a network as a schematic. A secondary goal was to have a tool that would generate textbook quality schematics for use in lecture notes and exam questions.

This paper describes how Lcapy generates schematics. It starts with a brief review of automated schematic drawing in Section II. It then describes the method used by Lcapyin Section III followed by some examples in Section IV. Finally, Section V contains a discussion and conclusions are provided in Section VI.

II. Background

There are many computer tools for generating a schematic and many are open-source [1]. However, drawing an elegant schematic given a netlist is a difficult problem. There have been a number of algorithms reported for automatically generating logic schematics for VLSI systems but few published algorithms for analog circuits. Beg [5] provides a review of previous works in this area. The most promising approach for automated analog circuit generation is to recognise common constructs [6] and then for each construct have a set component placement rules [7], [8].

Tools have also been developed for automatic analog schematic generation from a library of topologies [5]. The purpose of these tools are for tutor programs or problem generation for students. However, these tools solve a simpler problem since the circuit topology is known in advance.

Lcapy uses a semi-automated approach to the drawing problem where the netlist is annotated with drawing hints. These hints specify the component direction and optionally the component size. In addition, other attributes can be specified such as the component colour, label names, etc. Lcapy also requires the netlist to be augmented with explicit wires. Automated wire drawing is a simpler problem and may be implemented in the future.

III. Implementation

Lcapy describes networks comprised from parallel and series connections with an abstract syntax tree. More complex topologies are described using netlists with a syntax similar to that of SPICE, see [3]. Lcapy uses the netlist representation to draw a schematic and so the abstract syntax tree representation is converted to a netlist for drawing.

The simplified process of converting a netlist to a schematic consists of the following steps:

1) Parse netlist and generate a list of component objects.
2) For each component, use node constraints to generate graph.
3) Traverse graph to determine component placements.
4) For each component, output drawing instructions.

These steps are elaborated in the following subsections. Note, the x and y positions of each component node are calculated independently. A component placement graph for the x-positions is created and solved. This is followed by a creating and traversing a similar graph for the y-positions.

A. Netlist parsing

Lcapy uses two independent component class hierarchies; one for simulation and one for drawing. Both are generated

[1] en.wikipedia.org/wiki/Wikipedia:WikiProject_Electronics/Programs

from the same grammar and parser although not all the components that can be drawn can be simulated, for example, non-linear devices such as transistors. A hand-written parser was implemented to yield more meaningful error messages compared to a lex-yacc approach.

Components are defined as a hierarchy of Python classes. Each class defines the normalised coordinates of each of its nodes and has a number of attributes defining whether the component can be scaled, rotated, mirrored, or stretched. The node coordinates are transformed by scaling, rotating, mirroring, etc., and used to provide constraints for the component placement graphs.

B. Component placement graph creation

The first step is to traverse the list of components to create a component placement graph. The nodes of this graph are the positions of the component nodes and the edges are the connections between the component nodes. Each edge has an attribute specifying a distance and whether the distance is fixed or can be stretched. For example, the nodes of a resistor can be stretched since it has implicit connecting wires but the nodes corresponding to the pins of an integrated circuit are fixed.

To reduce the size of the component placement graphs, component nodes with shared coordinates are grouped together. For example, consider a horizontal resistor. The two nodes must share the same y coordinate. Similarly, the nodes of a vertical capacitor must share the same x coordinate.

C. Component placement graph traversal

There are four main steps to this algorithm:

1) Once a component placement graph is created it is augmented with special start and end nodes. Nodes without a predecessor are linked to the start node and nodes without a successor are linked to the end node.
2) The graph is then traversed to find the longest path using Dÿkstra's longest path algorithm[2]. This longest path fixes the size of the schematic. Nodes on this longest path are assigned a position based on the sum of the edge distance attributes from the start node.
3) Nodes that are not on the longest path are then assigned to a worklist. If they are a fixed distance from a known node they are assigned a position and are removed from the worklist. This iterates until the worklist does not change.
4) The worklist is then reconsidered. This time the graph is traversed from an unassigned node to find the closest assigned nodes in both the forward and backward directions. The shortest path between these bracketing nodes is then calculated along with the number of stretchy edges. From these measurements, the amount of stretch per stretchy edge can be calculated on a pro rata basis and the positions of the nodes along the shortest path between the bracketing nodes can be assigned. For efficiency, step 3 is repeated before considering the next node on the worklist.

[2]en.wikipedia.org/wiki/Dijkstra's_algorithm

D. Component drawing

Once the locations of the component nodes have been determined, the components can be drawn. Each component object has a draw method and this is invoked for each component to be drawn in turn.

The design of Lcapy allows multiple backends to be used for component drawing. Currently, the only implemented backend uses CircuiTikZ[3]. This is set of LaTeX macros built using Tikz/PGF macros [9] for drawing electrical circuit components. The draw method for each component writes CircuiTikZ LaTeX commands to a file. Depending on the required file format, this file may be interpreted by pdflatex to create a pdf file or perhaps converted to a svg vector graphics file.

Each CircuiTikZ component is named after the identifier of the corresponding Lcapy. This allows boxes to be drawn around selected components and other operations that reference components by name.

IV. Examples

This section shows a few examples to demonstrate how Lcapy works. For more examples see the Lcapy documentation[4].

A. Resistor

The simplest example is a single component. For example, the single resistor network shown in Figure 2 is described by the netlist of Listing 1. Note the placement hint to draw the negative node of the resistor to the right of the positive node.

```
R1 1 2; right
```

Listing 1. Single resistor netlist.

Fig. 2. Single resistor schematic with labelled nodes.

```
\begin{tikzpicture}[scale=1.00,transform
    shape,/tikz/circuitikz/bipoles/length
    =1.50cm,american currents, american
    voltages]
  \coordinate (1) at (0,0);
  \coordinate (2) at (2,0);
  \draw[] (1) to [R,l_=$R_{1}$,,,o-o,n=R1]
      (2);
  \draw[anchor=south east] (1) node {1};
  \draw[anchor=south east] (2) node {2};
  \end{tikzpicture}
```

Listing 2. Generated LaTeX mark-up for the resistor shown in Figure 2.

[3]www.ctan.org/tex-archive/help/Catalogue/entries/circuitikz
[4]http://lcapy.elec.canterbury.ac.nz/lcapy/.

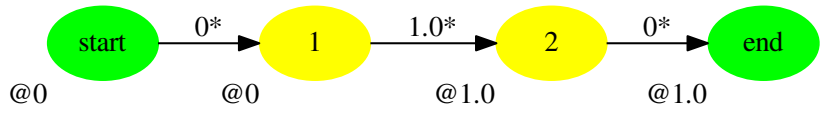

Fig. 1. Component placement graph for the x-positions of the resistor specified by Listing 1. Node names are shared electrical nodes. Edges are annotated with the minimum distance and a $\star$ if they can be stretched. Labels starting with @ are distances from the left. The yellow nodes denote the longest path between the special start and end nodes.

The generated LaTeX mark-up is shown in Listing 2. There are several sections:

1) The preamble that defines the sizes, scaling, and style.
2) Commands to define the node positions.
3) Commands to draw the components (and optionally the nodes).
4) Commands to annotate the node names.

B. Resistor network

The resistor network shown in Figure 4 is described by the netlist of Listing 3. The component placement graph for the x-direction is shown in Figure 3. In this example, the longest path is formed by R1, R2, and R3. The sizes of R4 and R5 need to be stretched equally to meet the constraints imposed by the vertical wires between nodes 1 and 5 and between nodes 4 and 7.

```
R1 1 2; right
R2 2 3; right
R3 3 4; right
W  1 5; down
W  4 7; down
R4 5 6; right
R5 6 7; right
```

Listing 3. Resistor network. Note, wires are using anonymous identifiers.

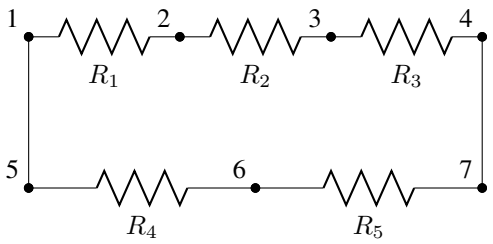

Fig. 4. Resistor network described by the netlist in Listing 3.

C. Annotations

The netlist for a more complicated example is shown in Listing 4. The resultant schematic is shown in Figure 5. This example has statements that control the style of the generated schematic. These start with a single semicolon. There are also statements starting with a pair of semicolons. These statements are CircuiTikZ macros that get appended to the CircuiTikZ macros generated by Lcapy. In this example, these statements are used to draw labelled boxes around the selected components specified by the names of their nodes. Alternatively, the component identifiers could have been used.

```
C1 1 0 100e-12; down, size=1.5, v={5\,kV}
R1 1 6 1500; right
R2 2 4 1e12; down
C2 3 5 5e-9; down
W 2 3; right
W 0 4; right
W 4 5; right
SW 6 2 no; right, l=, size=1.5
;draw_nodes=connections, label_nodes=false
    , label_ids=false
;;\node[blue,draw,dashed,inner sep=5mm,
    anchor=north, fit=(1) (6) (0), label=
    Human body model] {};
;;\node[blue,draw,dashed,inner sep=5mm,
    fit=(2) (3) (4) (5), label=CMOS input
    model]{};
```

Listing 4. An annotated schematic. Note, statements starting with ;; are CircuiTikZ commands.

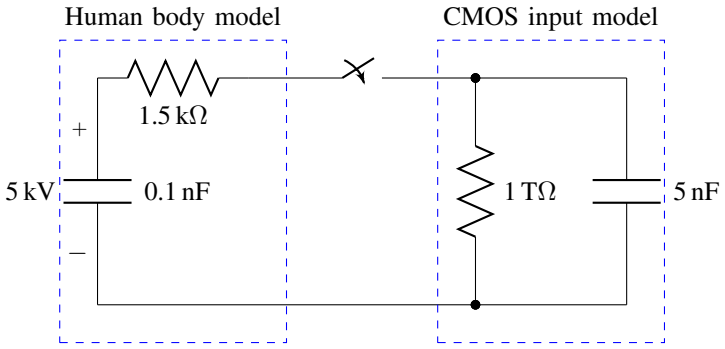

Fig. 5. Annotation of a schematic.

D. Implicit wires

Lcapy can draw implicit wires to power supply connections. For example, consider the Pierce oscillator in Figure 6.

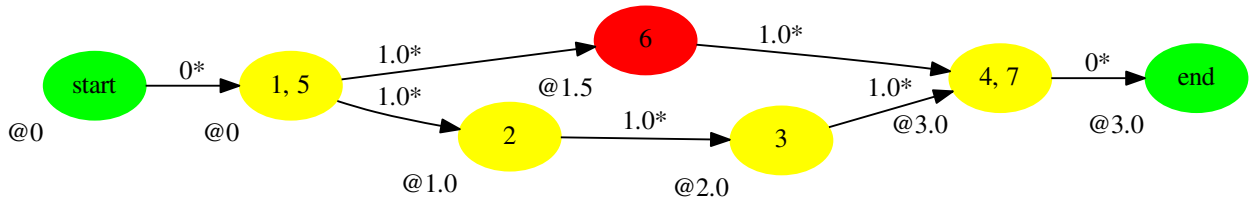

Fig. 3. Component placement graph for the x-positions of the resistor network specified by Listing 3. Node names are shared electrical nodes. Edges are annotated with the minimum distance and a $\star$ if they can be stretched. Labels starting with @ are distances from the left. The yellow nodes denote the longest path between the special start and end nodes.

This is described by the netlist shown in Listing 5. Note, this example cannot be simulated by Lcapy since it contains a non-linear logic inverter (albeit biased in its linear region by R_1).

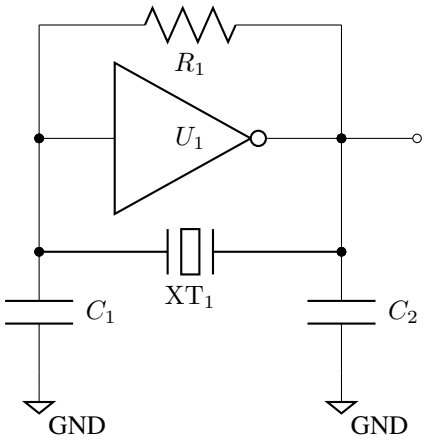

Fig. 6. Pierce oscillator.

```
U1 inverter .in .VSS .out .VDD; right
W 5 U1.in; right=0.5
W U1.out 6; right=0.5
W 6 9; right=0.5
R1 7 8; right
W 5 1; down=0.75
W 6 2; down=0.75
W 5 7; up=0.75
W 6 8; up=0.75
XT1 1 2; right
C1  1 4; down=0.8
C2  2 3; down=0.8
W 4 0; down=0.1, implicit, l=GND
W 3 0; down=0.1, implicit, l=GND
; draw_nodes=connections, label_nodes=
    false
```

Listing 5. Pierce oscillator. Note the two implicit connections to ground.

E. *Block diagrams*

Block diagrams can be generated by drawing shapes such as rectangles and circles. Figure 8 shows a signal block diagram that uses transfer functions and summing points. Listing 6 shows the corresponding netlist. Since this is a linear system, it can be simulated by Lcapy.

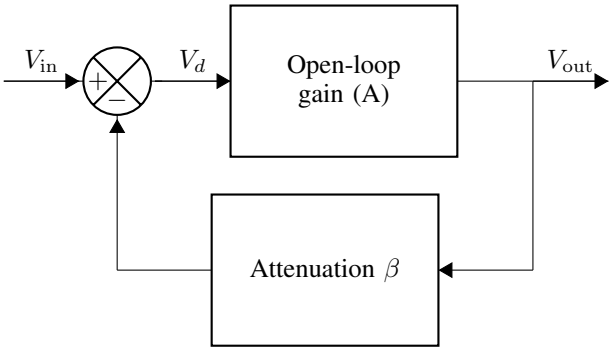

Fig. 8. Negative feedback block diagram.

```
W 1 SP1.inp; right=0.5, endarrow=tri, l=V_
    {in}
SP1 pm .inp .inm .out; right, l={}
W SP1.out TR1.in; right=0.5, endarrow=tri,
     l=V_d
TR1 .in .out A; right=1.5, l=Open-loop
    gain (A)
W TR1.out 2; right=0.5
W 2 3; down
TR2 .in .out B; left=1.5, l=Attenuation $\
    beta$
W 3 TR2.in; left=0.5, endarrow=tri
W TR2.out 4; left=0.5
W 4 SP1.inm; up, endarrow=tri
W 2 5; right=0.5, endarrow=tri, l=V_{out}
; draw_nodes=false, label_nodes=false
```

Listing 6. Negative feedback block diagram. TR denotes a transfer function; SP denotes a summing point.

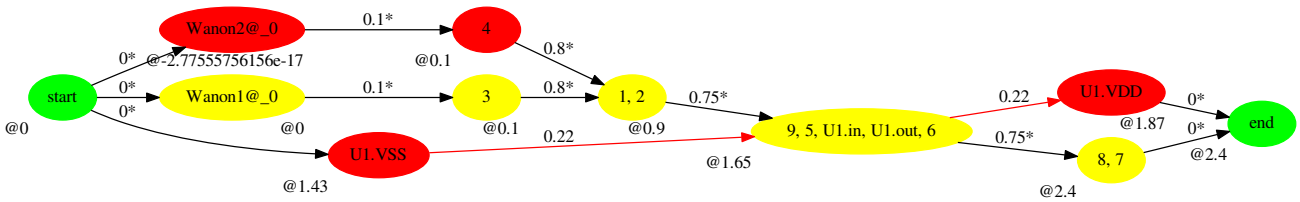

Fig. 7. Component placement graph for the y-positions of the schematic shown in Figure 6. With this example, there are multiple options from the start and to the end nodes.

F. Hierarchical diagrams

Lcapy allows large schematics to be created in a hierarchical manner by including netlists with the `.include` directive. Each of these included netlists can have their own namespaces to avoid conflicts, for example, see Listing 7 and Figure 9.

```
.include LC1.sch as s1
.include LC1.sch as s2
.include LC1.sch as s3
W s1.2 s2.1; right=0.1
W s1.3 s2.0; right=0.1
W s2.2 s3.1; right=0.1
W s2.3 s3.0; right=0.1
; label_nodes=false, draw_nodes=
    connections
```

Listing 7. Including other netlists into their own name-spaces.

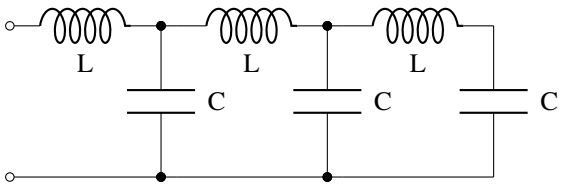

Fig. 9. Example of including other netlists.

G. Transformations

The primary motivation for Lcapy schematic drawing was to show the network created with series and parallel operators. For example, consider the Python script in Listing 8. This generates the schematic shown in Figure 10. Internally, Lcapy transforms an abstract syntax tree into a netlist by traversing the tree. The network can be transformed into an s-domain model using the `smodel` method as shown in Listing 9. This generates the schematic shown in Figure 11. Another example of a transformation is shown in Listing 10. This uses the `simplify` method and generates the schematic in Figure 12.

```
from lcapy import R, C, L
((R(1) + L(2)) | C(3)).draw()
```

Listing 8. RLC network created using series and parallel operators.

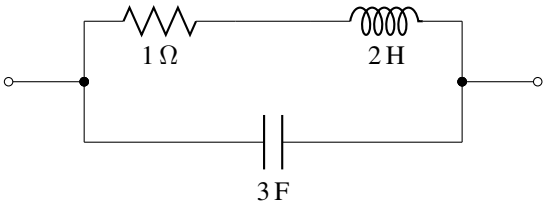

Fig. 10. RLC network created using series and parallel operators.

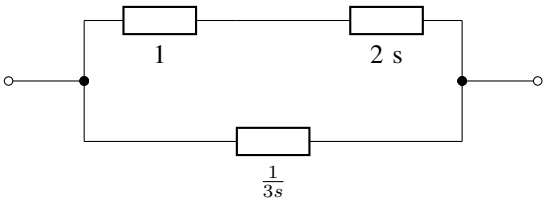

Fig. 11. s-domain model of RLC network created using series and parallel operators.

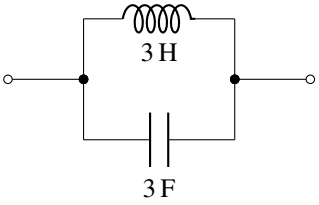

Fig. 12. Simplified RL network created using series and parallel operators.

H. Styles

Component attributes can be specified to set their colour, size, line style, etc. Attributes can also be specified for the overall schematic to control the style, see Listing 11. Three component styles are supported: American (default), British, and European, see Figure 13 for American and European examples.

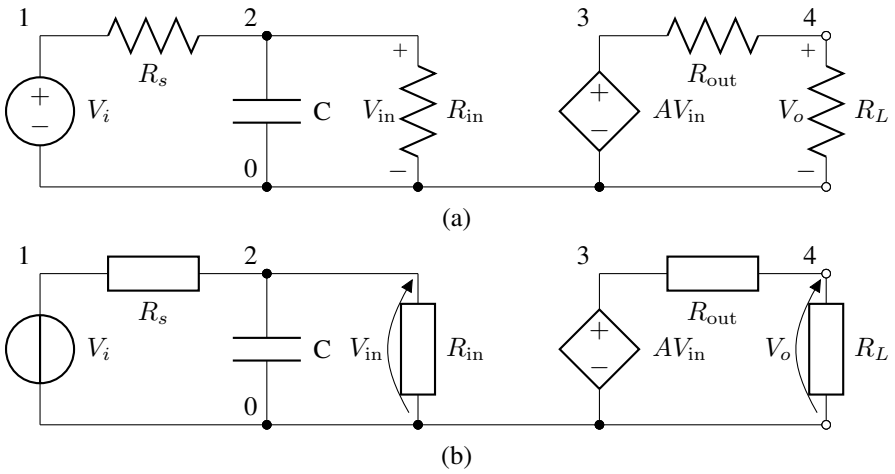

Fig. 13. American and European schematic styles.

```
from lcapy import R, C, L
((R(1) + L(2)) | C(3)).smodel().draw()
```

Listing 9. s-domain model of RLC network created using series and parallel operators.

```
from lcapy import R, C, L

((L(1) + L(2)) | C(3)).simplify().draw()
```

Listing 10. Simplification of a LC network.

V. Discussion

From a quick sketch on a piece of paper, a netlist can be quickly generated and from this a high-quality schematic can be produced. However, an error in a component attribute, say its direction, can result in a conflict in the component placement graph. Unfortunately, it is difficult to highlight the source of the error.

Further improvements would include the automatic generation of wires to connect nodes at the same potential. A graphical front-end to guide component placement would also be useful. Rules could be added to stretch automatically the spacing of components so an uncluttered, yet compact, schematic is created. Furthermore, better heuristics are required for the placement of component identifiers and values. In particular, voltage labels would be better placed on the outside of a circuit where possible. Finally, better customisation would be desirable to accommodate different tastes.

The schematics in this paper were generated by a Lcapy program called `schtex`. This reads a netlist file and then create a `.tex` file that is input into the source file for this paper.

VI. Conclusions

Lcapy is a Python package for circuit simulation, designed to help teaching linear circuit theory. It can also produce textbook quality schematics from either a netlist annotated with drawing attributes or from an abstract syntax tree.

```
Vi 1 0_1; down
Rs 1 2; right=1.5
C 2 0; down
W 0_1 0; right
W 0 0_2; right
Rin 2_2 0_2; down, v=V_{in}
W 2 2_2; right
E1 3 0_3 2 0 A; down, l=A V_{in}
Rout 3 4; right=1.5
RL 4 0_4; down, v=V_o
W 0_2 0_3; size=1.2
W 0_3 0_4; right
P1 4 0_4; down
; style=european, draw_nodes=connections
```

Listing 11. Changing the component style.

References

[1] G. Van Rossum and F. L. Drake Jr, *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam, 1995.

[2] D. Joyner, O. Čertík, A. Meurer, and B. E. Granger, "Open source computer algebra systems: SymPy," *ACM Communications in Computer Algebra*, vol. 45, no. 3/4, pp. 225–234, 2012.

[3] M. P. Hayes, "Lcapy: Linear circuit analysis with python," in *ENZCon'14 Conference Proceedings*. Hamilton, New Zealand: Electronics New Zealand Inc. (ENZCon Inc), 2014, pp. 89–94.

[4] L. W. Nagel and D. O. Pederson, *SPICE: Simulation program with integrated circuit emphasis*. Electronics Research Laboratory, College of Engineering, University of California, 1973.

[5] A. Beg and A. Beg, "Auto-generating publication-quality circuit schematics using open technologies," in *Proceedings of the 21st Western Canadian Conference on Computing Education*, 2016, p. 9.

[6] Y. Wu, "Analog-aware schematic synthesis," *Advances in Analog Circuits*, p. 247, 2011.

[7] L. D. Green and J. Andersen, "Automated generation of analog schematic diagrams," in *Circuits and Systems, 1990., IEEE International Symposium on*, 1990, pp. 3197–3200.

[8] G. M. Swinkels and L. Hafer, "Schematic generation with an expert system," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 9, no. 12, pp. 1289–1306, 1990.

[9] T. Tantau, "The TikZ and PGF packages: Manual for version 2.10," 2010.

A preliminary investigation of an electromagnetic groundwater flow measuring device

Michael P. Hayes
Department of Electrical and
Computer Engineering
University of Canterbury
Christchurch, New Zealand
Email: michael.hayes@canterbury.ac.nz

Ian Platt
Lincoln Agritech Limited
Lincoln, New Zealand
Email: ian.platt@lincolnagritech.co.nz

***Abstract*—This paper describes a preliminary simulation of an electromagnetic system for measuring groundwater flow. The simulation solves the flowmeter equation for a 1 m prototype system using a circular coil. A parametric model is used to calculate the magnetic field and Jacobi relaxation is used to determine the induced electric field due to the moving water ions. Peak potential differences of the order of a few tens of nanovolts are predicted for water speeds of 10 mm/hour and magnetic field strengths of 10 mT. The potentials are reduced compared to those for a flowmeter due to additional conductive paths outside the region where the magnetic field is strong.**

I. Introduction

Most of the World's fresh water is below the ground and moves slowly forming stream flows. This flow of groundwater is complex since it depends upon many factors, such as geology, soil composition, and topographical control (e.g., lakes, rivers, vegetation) [1]. Consequently most information about flow directions, rates, and depths come from models tuned by a few point locations consisting of monitored wells. The principal advantage of this proposed technology is in reducing the uncertainty in groundwater models to a point where biophysical models can be generated to guide land use practices [2]–[4]. This is the single biggest constraint to resolving the water quality issue facing land users and the community in New Zealand today.

This paper describes a preliminary simulation of a prototype scale model for inductive groundwater flow measurement. The overall goal is to construct a system for measuring groundwater flow. The proposed method creates a magnetic field using an electromagnetic and measures the potential difference induced by the flow of moving water ions. The induced voltage can be described by Faraday's Second Law of electromagnetic induction. This effect is routinely used for measuring the flow rates of liquids in pipes [5], [6] and open channels [7]. However, there are two main differences: the water is not constrained to a pipe and a Helmholtz coil cannot be used to create a uniform magnetic field. Thus this paper considers the effect of these differences on the measured potential difference.

The paper starts with a brief analysis of the electromagnetic theory in Section II and describes a numerical simulation in Section III. The results of the simulation are presented in Section IV. The paper concludes with a discussion and conclusions in Section V and Section VI.

II. Theory

The motion of the ions in the water as they pass through a magnetic field induces an electric potential,

$$\mathbf{E}' = \mathbf{E} + \mathbf{v} \times \mathbf{B}, \tag{1}$$

where $\mathbf{E}'$ is the electric field in the reference frame of the moving water ions, $\mathbf{E} = -\nabla V$, is the induced electric field for an induced electric potential, V, in the rest frame, $\mathbf{v}$ is the flow velocity of the water, and $\mathbf{B}$ is the applied magnetic field. Using Ohm's law,

$$\mathbf{J} = \sigma \mathbf{E}', \tag{2}$$

where $\mathbf{J}$ is the electric current density and σ is conductivity, then using the fact that the divergence of $\mathbf{J}$ is zero, yields

$$\sigma \nabla^2 V + \nabla\sigma \cdot \nabla V = \nabla\sigma \cdot \mathbf{v} \times \mathbf{B} + \sigma \mathbf{B} \cdot \nabla \times \mathbf{v} - \sigma \mathbf{v} \cdot (\nabla \times \mathbf{B}). \tag{3}$$

In practice, the induced current density, $\mathbf{J}$, is too small to have a significant effect on the applied magnetic field, so $\nabla \times \mathbf{B} = 0$. Finally, by assuming a uniform, isotropic conductivity, σ, then $\nabla\sigma = 0$ and (3) simplifies to the standard flowmeter equation [8]

$$\nabla^2 V = \nabla \cdot (\mathbf{v} \times \mathbf{B}). \tag{4}$$

This differential equation can be solved given $\mathbf{v}$ and $\mathbf{B}$ to find V subject to boundary conditions.

Assuming point electrodes positioned at $\mathbf{x}_1$ and $\mathbf{x}_2$, the measured potential difference is

$$V_m = V(\mathbf{x}_2) - V(\mathbf{x}_1). \tag{5}$$

Note, the size of the electrodes will affect the potential V. This is not explicit in (4) but alters the boundary conditions used when solving the differential equation.

From (4), the source of the electric potential is $\nabla \cdot (\mathbf{v} \times \mathbf{B})$ and we desire this to be as large as possible to maximise the signal to noise ratio of the measured potential difference. Note, $\mathbf{v} \times \mathbf{B}$ can be interpreted as the open-circuit electric field, i.e., the electric field that would occur if there were no eddy currents. However, since $\mathbf{v} \times \mathbf{B}$ is non-uniform and the medium is conductive, eddy currents flow and so where the electric field is generated the measured electric field will be smaller than $\mathbf{v} \times \mathbf{B}$.

An alternative approach to determine the measured potential difference between the electrodes of a flowmeter is to use

the concepts of virtual currents and weight vectors [8]. The measured potential difference is given by

$$V_m = \iiint_V \mathbf{v} \cdot \mathbf{W} \, \mathrm{d}V, \tag{6}$$

where the weight vector is

$$\mathbf{W} = \mathbf{B} \times \mathbf{J}_v, \tag{7}$$

for a magnetic flux density $\mathbf{B}$, virtual current $\mathbf{J}_v$, and flow volume V. The virtual current does not exist in practice. It is the current density that would exist in the fluid if a 1 A current source was connected to the electrodes and the fluid was stationary. The virtual electric potential field resulting from this current density is obtained from the Laplace equation,

$$\nabla^2 V_v = 0. \tag{8}$$

There are a number of publications that try to optimise the weight vector so that the induced voltage is proportional to the flow, even when the flow is not uniform across the pipe [9], [10]. However, these methods are unlikely to be applicable to the groundwater problem. Furthermore, the difficulty with applying (6) for a groundwater flow measuring device is that the volume of integration is not constrained by the size of the pipe.

III. Numerical simulation

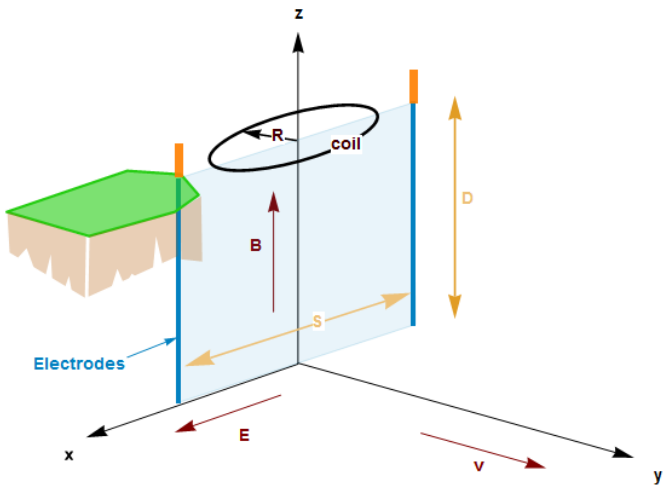

Fig. 1. System geometry. The coil sits in the $x-y$ plane at $z = 0$ and the water velocity is in the y direction.

To determine the likely potential differences that will be measured, a numerical simulation was performed in Python using the parameters given in Table I and with the geometry as shown in Figure 1. Note, the coil sits in the $x-y$ plane and the water is moving with a uniform velocity in the y direction.

A number of simplifications were assumed:

- The conductivity, σ, was assumed to be constant and isotropic for the lower half-space, $z < 0$.
- Water was flowing with a uniform velocity throughout the lower half-space, $z < 0$. In practice, the flow is zero above the depth of the water table.
- The soil matrix was ignored.
- The circular coil was ideal with an infinitesimal conductor size. This allowed the magnetic field to be calculated from a parametric model.
- The electrodes were points so that the electric field was not influenced by them.
- The load resistance between the electrodes was infinite so there was no voltage drop due to contact resistance at the electrodes.

A finite-element representation of the flowmeter equation (4) was solved using Jacobi relaxation with the method of images. This is numerically slow compared to a Fourier domain approach using Green's functions. Implicit Dirichlet boundary conditions were applied due to the finite size of the simulation. These boundary conditions are equivalent to enclosing the simulated region in a conductive box. While this erroneously forces the potential to be zero on the surface of the box the effect is minor if the box is large compared to the size of the coil.

The boundary between the moving water and the soil (and air) above requires a Neumann boundary condition; this implies that there is no potential difference across the boundary and thus no current flow across the boundary. Implementing a Neumann boundary condition when solving (4) requires the finite-element equations to be modified at the boundary. An alternative approach is to consider that the water only moves in the half-plane described by $z < 0$ and to have water in the half-plane above it. In this case, currents will flow cross the boundary between the two half-spaces but due to symmetry the potentials will be reduced by a factor of two.

Another approach requires modifying the flowmeter equation (4) to incorporate the different conductivities in the two half-spaces. It has the advantage that it can determine the potentials within the water and above it. However, this leads to a more difficult PDE to solve.

The method implemented was the method of images. Here both half-planes consist of water moving with the same velocity. However, due to symmetry no current flows across the boundary since the potentials are the same. The potentials in the negative half-space are now correct but the potentials in the positive half-space have to be ignored.

TABLE I. Simulation parameters.

Frequency	f	0 Hz (DC)
Coil radius	R	0.5 m
Coil current	I	10,000 A-turns
Water conductivity	σ_w	10 mS/m
Electrode separation	S	0.8 m
Electrode depth	D	0.25 m
Electrode length	L	0.25 m
Electrode width	w	25 mm
Electrode breadth	b	25 mm
Water velocity	$\mathbf{v}$	10 mm/hour in $\hat{\mathbf{y}}$
Sampling grid size	Δ	25 mm

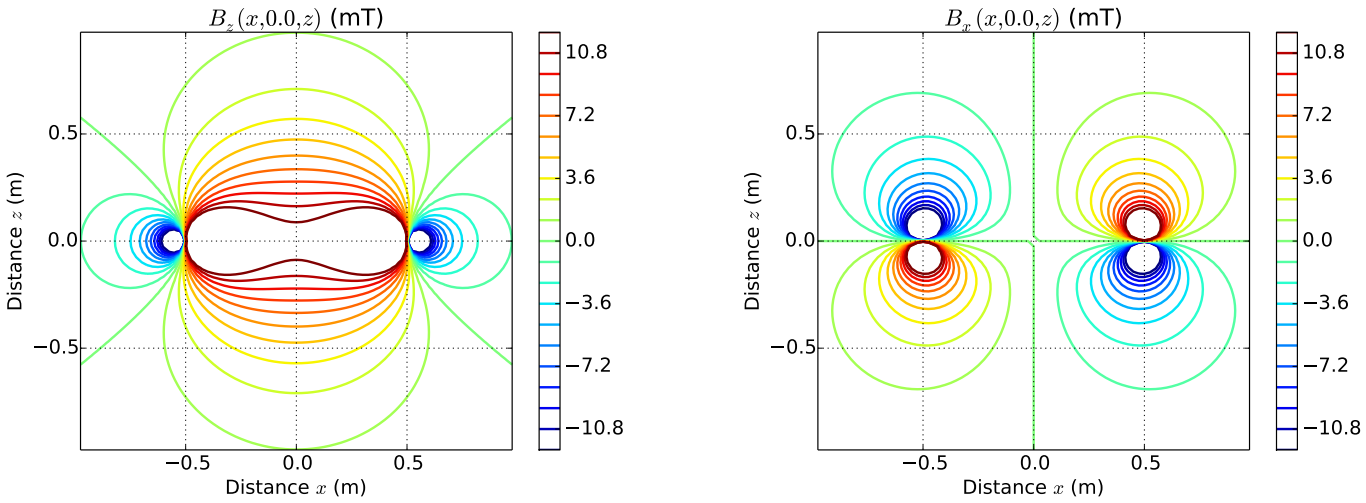

(a) (b)

Fig. 2. Magnetic field (mT) contours in x-z plane: (a) axial component, B_z, (b) lateral component, B_x.

IV. Results

The magnetic field due the circular coil is shown in Figure 2. Note, once the depth, z, is greater than the diameter of the coil, the magnetic field rapidly decreases.

The induced electric field is shown in Figure 3 with the electric potential shown in Figure 4. Note, the highest contours have been clipped to reduce the image clutter and due to the method of images, the potentials are symmetrical for both half-planes (the potentials for the half-space $z > 0$ will not occur in practice).

Slices of the contours of the electric potential are shown in Figure 5. Note, the largest potential differences are measured when the electrodes are at the coil radius (and normal to the flow direction). However, the electrodes can be placed outside the coil with a small drop in potential difference.

The effect of the implicit Dirichlet boundary conditions can be seen in Figure 5 where the potential is forced to zero at $x = \pm 2\,\text{m}$. Making the size of the simulation larger defers this effect, see Figure 5, but makes the simulation significantly slower to run.

As a comparison, a simulation was performed with the water constrained to a square channel of width 1 m and depth 1m. The induced electric potential is shown in Figure 6 with slices of the contours of the electric potential shown in Figure 7. Again note, the largest potential differences are measured when the electrodes are aligned with the coil.

V. Discussion

The potential difference for a flowmeter is often expressed as [5]

$$V_d = K B_z v S, \tag{9}$$

where K is a constant (ideally close to 1) and S is the electrode separation. For example, with $B_z = 10\,\text{mT}$, $v = 10\,\text{mm/hour}$, and $S = 1\,\text{m}$, (9) implies a measured potential difference of 27 nV assuming $K = 1$. Whereas the simulation gives a value

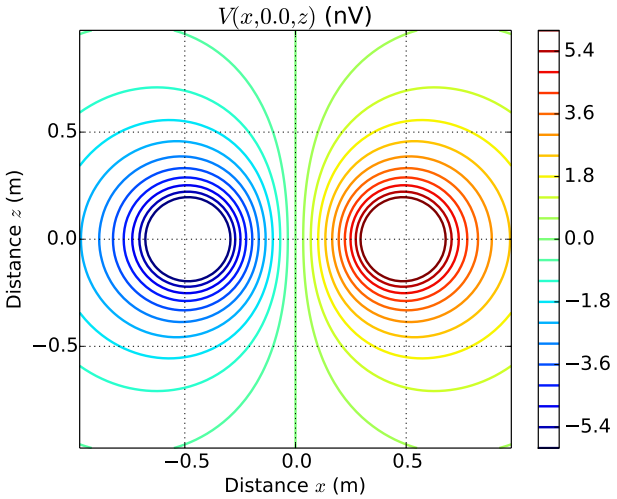

Fig. 4. Electric potential contours in x-z plane for $y = 0\,\text{m}$.

of about 12 nV, corresponding to $K = 0.44$. This constant is higher for fluid flow in a pipe due to two reasons:

1) Since both sides of the pipe are accessible, a Helmholtz coil can be used; this provides a more uniform magnetic field in the axial direction.
2) The conducting fluid is surrounded by an insulator. Thus the potentials are not reduced by currents flowing outside the region from where the electric fields are produced.
With the groundwater geometry, the resistance of the water outside the magnetic field acts as a resistive load and reduces the potentials. Remember, between the electrodes there are many paths; each can be considered a Thévenin equivalent circuit comprised of a voltage source and a series resistance. The voltage source represents the induced voltage along

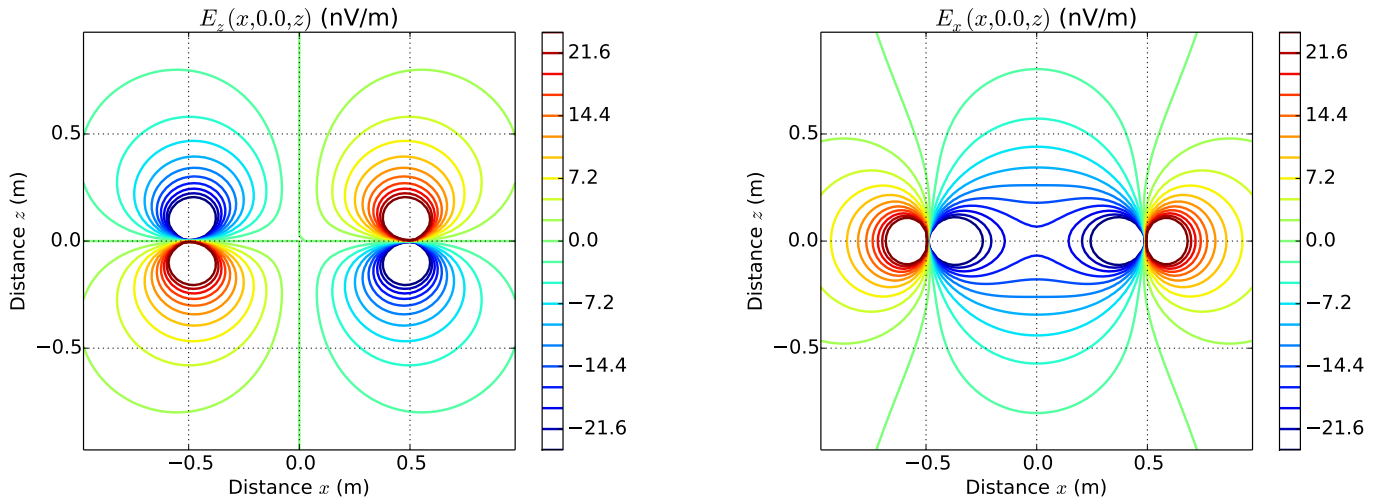

Fig. 3. Induced electric field contours in x-z plane: (a) axial component, E_z, (b) lateral component, E_x. Note, the electric field is not valid for $|z| > 0$.

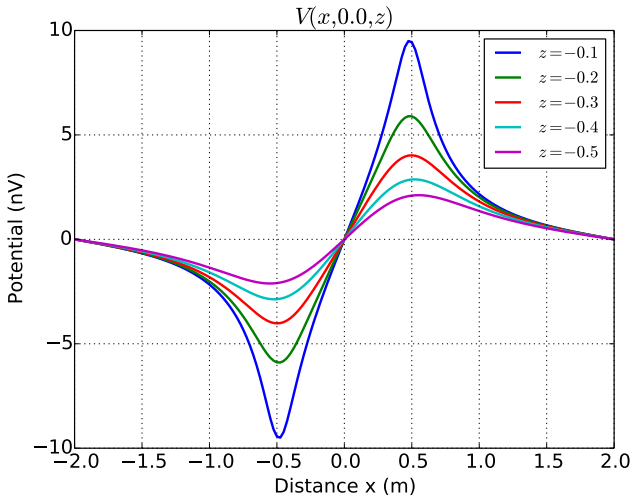

Fig. 5. Electric potentials at $y = 0$ for groundwater geometry. Note, the biggest potential difference occurs when the electrodes are below the coils.

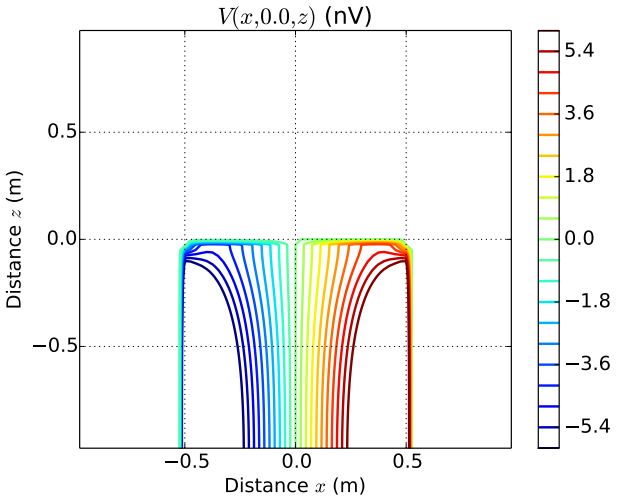

Fig. 6. Electric potentials for square insulated channel with constant axial field, $B_z = 10\,\text{mT}$. Electric potential contours in x-z plane for $y = 0\,\text{m}$. Note, the potentials are not valid for $|z| > 0$. In practice, the potentials would decay to zero.

the path and the resistance models the path resistance. The myriad of Thévenin equivalent circuits are all in parallel and are thus loaded by each other.

The simulation has assumed an average water flow speed and an average conductivity. In practice, these are spatially varying due to the soil. This inhomogeneity will introduce fluctuations in the induced electric field

VI. Conclusion

The results support the hypothesis that the potential difference for the proposed groundwater measuring system is going to be extremely small; $\approx 6\,\text{nV}$ at a depth of $0.5\,\text{m}$ for a flow speed of $10\,\text{mm/hour}$ with a magnetic field generated by an ideal circular coil with 10,000 A-turns. To achieve a sufficient signal to noise ratio so that these voltage differences can be measured requires long integration times to reduce the noise bandwidth. This assumes that there are no interfering signals within this band.

Appendix

The magnetic field at a point $\mathbf{x} = (\rho;\theta,z)$ in cylindrical coordinates produced by a current I in an ideal coil with infinitely thin wires, radius R, sitting on the x-y plane centred at the origin, has Cartesian components given by [11]

$$B_x(\mathbf{x}) = \left(\frac{R^2+z^2+\rho^2}{z^2+(\rho-R)^2}E_2(k) - E_1(k)\right)\frac{x}{\rho}K, \qquad (10)$$

$$B_y(\mathbf{x}) = \left(\frac{R^2+z^2+\rho^2}{z^2+(\rho-R)^2}E_2(k) - E_1(k)\right)\frac{y}{\rho}K, \qquad (11)$$

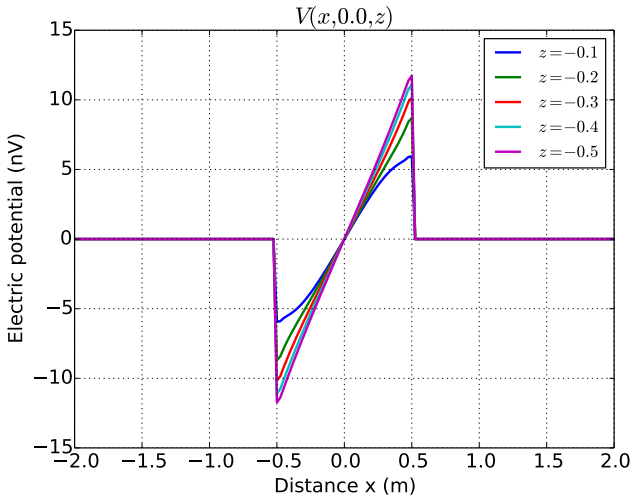

Fig. 7. Induced electric potentials at $y = 0$ for a square insulated channel with constant axial field, $B_z = 10\,\mathrm{mT}$.

$$B_z(\mathbf{x}) = \left(\frac{R^2 - z^2 - \rho^2}{z^2 + (\rho - R)^2} E_2(k) + E_1(k) \right) K, \tag{12}$$

where E_1 and E_2 denote the elliptic integrals,

$$k = 2\sqrt{\frac{\rho R}{z^2 + (R + \rho)^2}}, \tag{13}$$

and

$$K = \frac{\mu_0 I z}{2\pi\rho\sqrt{z^2 + (R + \rho)^2}}. \tag{14}$$

REFERENCES

[1] R. A. Freeze and J. A. Cherry, "Groundwater," Tech. Rep., 1979.

[2] M. Moreau and M. Bekele, "Groundwater component of the water physical stock account (wpsa)," Tech. Rep., 2015.

[3] A. P. Lovett and S. G. Cameron, "Development of a national groundwater atlas for New Zealand," Tech. Rep., 2014.

[4] P. White, M. Rosen *et al.*, "Groundwater resources in new zealand." *Groundwaters of New Zealand*, pp. 45–75, 2001.

[5] R. C. Baker, *An introductory guide to flow measurement*. John Wiley & Sons, 2002, vol. 2.

[6] J. A. Shercliff, *The theory of electromagnetic flow-measurement*. Cambridge University Press, 1987.

[7] J. Jakubowski and A. Michalski, "Application of selected linear algebra processing methods in the electromagnetic flow measurement for open channels," *Instrumentation and Measurement, IEEE Transactions on*, vol. 57, no. 8, pp. 1678–1684, 2008.

[8] M. Bevir, "The theory of induced voltage electromagnetic flowmeters," *Journal of Fluid Mechanics*, vol. 43, no. 03, pp. 577–590, 1970.

[9] J. Hemp and H. Versteeg, "Prediction of electromagnetic flowmeter characteristics," *Journal of Physics D: Applied Physics*, vol. 19, no. 8, p. 1459, 1986.

[10] S. Yin and B. Li, "A new approach for solving weight functions of electromagnetic flowmeters using resistive network modeling," *Journal of Applied Mathematics*, 2013.

[11] J. C. Simpson, J. E. Lane, C. D. Immer, and R. C. Youngquist, "Simple analytic expressions for the magnetic field of a circular current loop," 2001. [Online]. Available: http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20010038494_2001057024.pdf

Speed Control of Small Motors Through the Ćuk Converter Topology

Michael Hutchins
School of Science and Engineering
University of Waikato
New Zealand

Jonathan Scott
School of Science and Engineering
University of Waikato
New Zealand

***Abstract*—The Ćuk Converter offers simultaneous buck-boost operation, but requires careful design owing to its having a fourth-order transfer function and numerous practical design constraints. We exploit a serendipitous overlap between the converter circuit and the equivalent circuit of a dc motor to design a motor controller that can operate with supply voltage that is lower than the motor full-speed requirement. We examine the transfer function when such a topology is used to control the speed of a small motor. We conclude that the approach is relatively straightforward owing to the impact of the motor's inductance. Measurements agree with theory.**

***Index Terms*—motor speed control, Cuk converter, switchmode power conversion.**

I. Introduction

A. Ćuk Converter

When Middlebrook and his graduate student Slobodan Ćuk presented in 1983 the detail of their 1977 design for a switchmode converter, the design grabbed the imagination of researchers because it approximated the "ideal dc transformer" [1]. The topology offered conversion between input and output dc voltages, with the voltage magnitude ratio set purely by the duty cycle of drive to a single switch, one side of which could be at ground potential. The outline circuit is shown in figure 1.

For a number of reasons the design was slower to take off than its promise might have lead the reader to expect. The transfer function is fourth order; the position of poles varies not only with component values, but also the controlling duty cycle; parasitic resistance in the two inductors and the so-called "Ćuk capacitor" all tend to significantly displace the poles of the characteristic equation; and the circuit was relatively complicated compared to other topologies [2]. Only now in the 21^{st} century is the converter used with any confidence. For example, modern switching methods are just being applied [3], and a search of the Xplore Digital Library yields a total of 48 papers with "Ćuk" in the title published between 1990 and 1999, but 137 papers in the decade and a half since January 2000 [4]. Along with its variant the Single-Ended Primary-Inductance Converter (SEPIC converter) that swaps the position of output switch and inductor, and is thus able to generate an output voltage of the same polarity as the input, the Ćuk converter is the most versatile, but hardest to design, of switchmode converters.

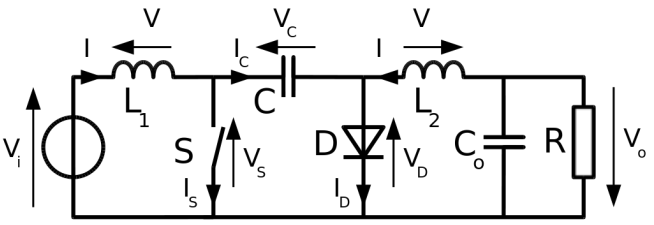

Fig. 1. Circuit of the Ćuk converter showing the conventions for currents and voltages. The middle capacitor "C" is the "Ćuk Capacitor". The switch S is usually a transistor with grounded source/emitter.

B. Brushed DC Motor

A brushed dc (BDC) motor has an equivalent circuit that consists of the series combination of a voltage source, a resistance and an inductance. The voltage source represents the energy sink that is the mechanical output of the motor, or the source of electrical energy coming from the mechanical components when the motor operates in generator mode. The I-V characteristic of this voltage source embodies the pole that arises from the mass of the mechanical components as well as the loss inherent in doing mechanical work. Previous work has shown that feedback control of the speed of small BDC motors presents greater difficulty that control of larger motors, as the mechanical pole tends not to be dominant [5].

The equivalent circuit of a BDC motor, shown in figure 2, presents 2 poles. It also bears a strong resemblance to some components of the Ćuk topology. Consider in the circuit of figure 1 that the parallel combination of C_o and R will resemble a voltage source if C_o is large, and L_2 with its inescapable parasitic resistance map exactly to L_m and R_m in figure 2. Overlaying these two circuit diagrams with the motor replacing the equivalent parts of the Ćuk circuit leads to the circuit of figure 3.

II. Transfer Function

We will now consider the transfer function of the circuit of figure 3 driving a typical, physically-small BDC motor. The mathematical derivation of the state-space equations is given in the appendix. Predictions will be made using these equations, evaluated in Matlab.

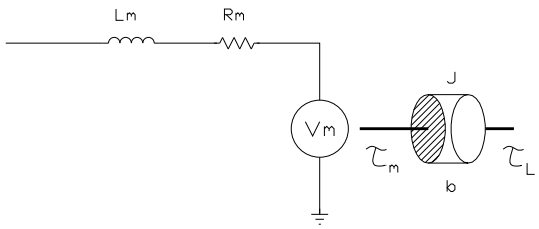

Fig. 2. The equivalent circuit of a brushed dc motor from [5].

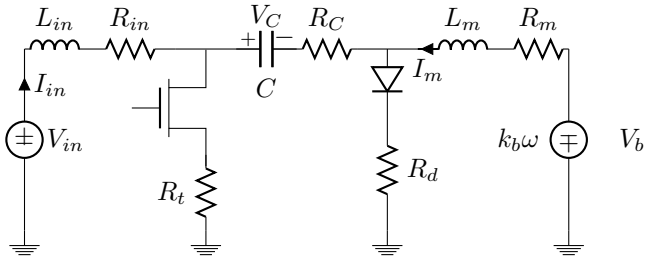

Fig. 3. The equivalent circuit of a Ćuk converter including parasitic resistances, driving a brushed dc motor that contributes L_m, R_m, and V_b. Note that V_b is proportional to the angular speed of the motor, ω.

The series inductance of a small motor is typically a few millihenries. Our example motor has $L_m = 16$mH. Motors with rotors of about 1 cubic centimetre typically have this or even smaller an inductance.

Given the ability to switch at around 100kHz, we expect values for L_{in} in the range up to a hundred μH; we will start with 10μH which is quite practical a value for small converter circuits. Practical values of the Ćuk capacitor are around a few μF; we will start with 2.2μF. We will consider at first the case where duty cycle D is 50%. Figure 4 sets the scene.

Arrows show the movement of the poles as L_m rises from a few mH to a few hundred mH. Poles 3 and 4 lie close to 40 radians/second, and figure 5 blows up this part of the plot. Pole 4 describes a small, almost-closed arc. Pole 3 nearly crosses Pole 4 as L_m moves past a value of 337mH. Figure 6 shows the movement of Pole 1 and Zero 1, which remain far out and not of much interest. It is clear that Poles 3 and 4, or rather their separation, may be of concern when applying feedback control to regulare V_b.

We next consider the variation of duty cycle, D. Motor inductance is held at 16mH, corresponding to our test motor, L_{in} is kept at 10μH and the Ćuk capacitor stays at 2.2μF. Figure 7 plots the interesting (close-in) situation. The two close poles, 3 and 4, actually separate as the duty cycle is increased, so that stability is likely to be better at higher loads. Crucially, this figure suggests that the converter-motor system will be no more difficult to control once feedback is applied than was the motor alone, that is with the motor powered with something close to an ideal voltage amplifier (or PWM equivalent).

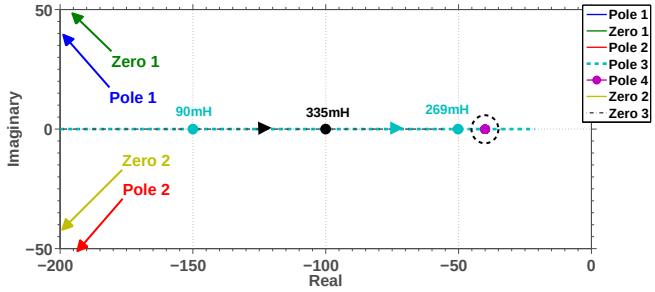

Fig. 4. The pole-zero diagram of the converter-motor combination. A complex conjugate pair of poles and a complex conjugate pair of zeroes lie ot to the left of the plot, while the dashed circle encloses a pair of poles associated with the motor components. These are expanded in figure 5.

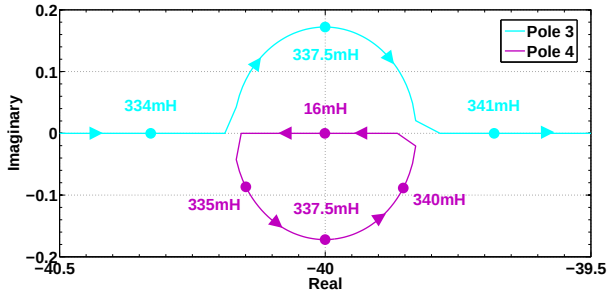

Fig. 5. The trajectory of Poles 3 and 4, the "motor inductance" and "mechanical" poles, as L_m is varied.

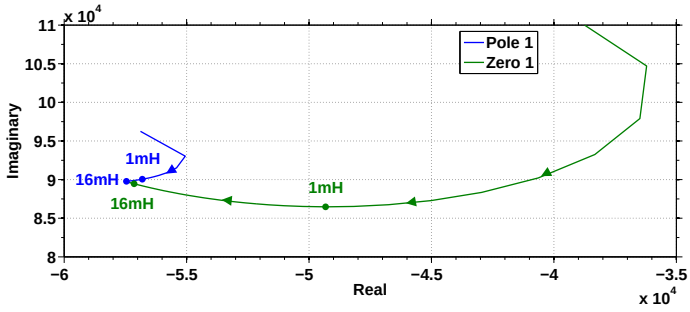

Fig. 6. The trajectory of Pole 1 and Zero 1 as L_m varies. Note the large x-axis scale.

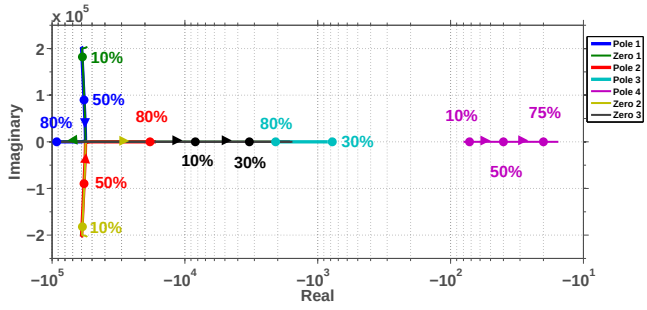

Fig. 7. Close-in pole and zero trajectories as Duty Cycle, D, is varied while other components remain at their default values. Note especially that Poles 3 & 4 *separate* as D increases, improving stability prospects as the converter increases its power transfer.

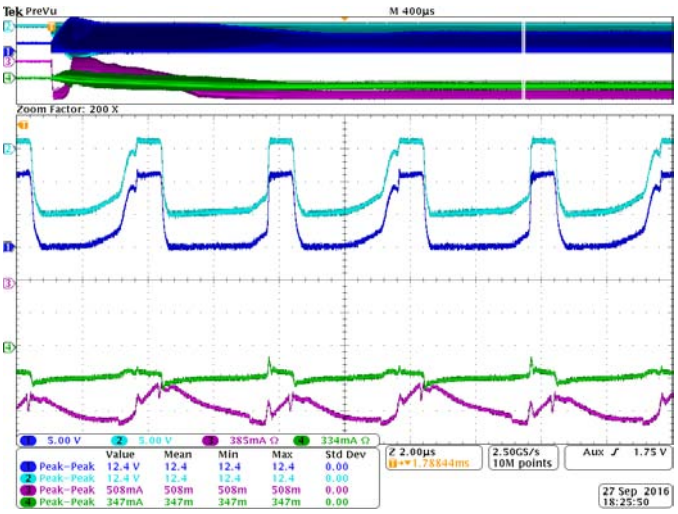

Fig. 8. Oscilloscope screen capture of the converter being subjected to a step input change in duty cycle D. The top window shows the evolution over milliseconds, the lower window the instantaneous values once the step has settled but the shaft speed has not changed. Traces show the MOSFET drain voltage, the diode voltage (with V_C being the difference of these), the input (green) and output (purple, bottom trace) current.

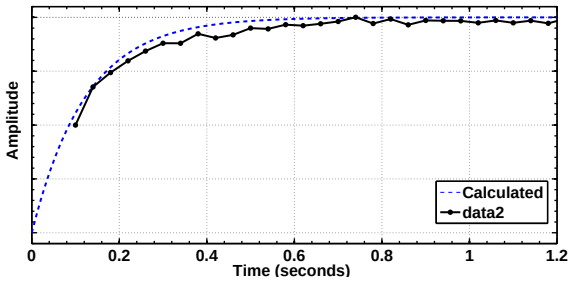

Fig. 9. Variation of shaft speed in response to a step change in D. The speed data does not extend to zero because it is measured from the period of signals from an optical encoder on the shaft and no data is available until the shaft has moved through a few degrees.

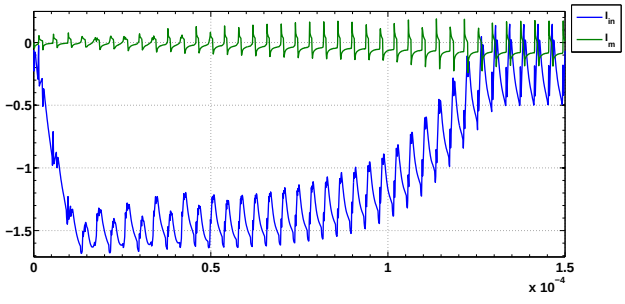

Fig. 10. Measured instantaneous input and motor current in response to a step change in D from 0 to 50%.

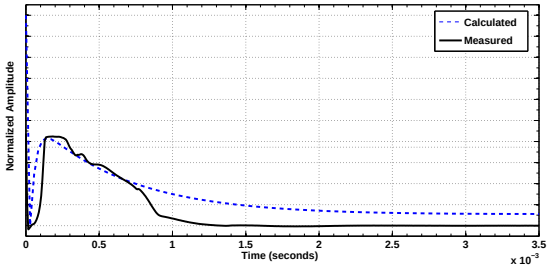

Fig. 11. Predicted and measured values of the averaged input current in response to a step input of duty cycle, D, from 0 to 50%.

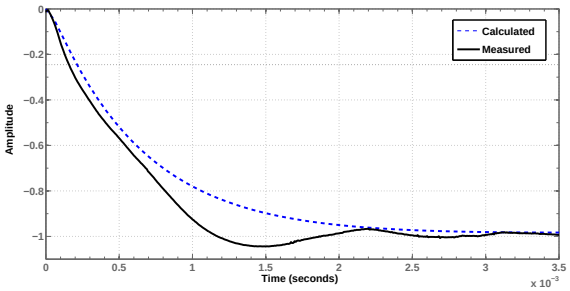

Fig. 12. Predicted and measured values of the mean output current in response to a step input of duty cycle, D, from 0 to 50%

III. Measured Results

Reassured by the above analysis of the transfer function, we construct a converter with the default values above, a 5V input supply, and a 12V BDC motor with $L_m = 16$mH series inductance and 12Ω resistance, $L_{in} = 10\mu$H, $R_{in} = 1.2\Omega$, $R_C = 0.1\Omega$, $R_t = 0.1\Omega$, $R_d = 0.1\Omega$, $J = 0.05$, $b = 0.8$, and $K_b = 0.01$.

Figure 8 is a screen capture from measurements made on the prototype. To a remarkable degree to time-domain waveforms have the instantaneous shape that is to be expected, except for finite risetimes and small "wiggles" that are attributed to measurement artefacts and extraneous parasitic impedances. The upper part of the figure shows the evolution of waveforms on a longer time scale. Figure 9 shows the variation of shaft speed on yet longer a time scale. Matlab simulation agrees with measured data where available. Figure 10 shows the input and output currents in response to the same step input. While the input inductor current may spike, the output current describes, on average, a smooth response reminiscent of a single-pole exponential change. Figures 11 and 12 show the evolution of predicted and measured mean input and output current. There is a discrepancy about 1 second that is attributed to mechanical imperfections in the apparatus and errors in our values for electrical parasitics.

IV. Conclusion

We have shown that the Ćuk topology is readily applied to drive a small motor. The overlap of circuit topologies results in poles from the motor replacing poles in the Ćuk transfer function to yield only the same number of poles (order of the system) as existed without the motor. Component values can be chosen so as to leave a dominant pole, and a system around which feedback can be applied with no more complexity than existed in the case of the motor driven by a perfect analog source.

Appendix

There are various ways to model a switchmode circuit [6]. Here we develop the state-space equations for the circuit of figure 3 [7], [8]. Including the mechanical pole of the motor and load, the system will be of fourth order. Let the state be

$$x = \begin{bmatrix} I_{in} \\ I_m \\ V_C \\ \omega \end{bmatrix} \tag{1}$$

where I_{in} is the (input) current drawn from the source through inductor L_{in}, I_m is the motor (output) current, V_C is the voltage across the Ćuk capacitor, and ω is the (desired) motor shaft output rotational speed. Next we define the input variables

$$u = \begin{bmatrix} V_{in} \\ T_L \end{bmatrix} \tag{2}$$

where V_{in} is the supply (input) voltage and T_L is the torque (load) encountered on the motor output shaft. Then

$$\dot{x} = Ax + Bu \tag{3}$$
$$y = Cx \tag{4}$$

where y is the output of the system. The state-space equations are perturbed with

$$d = D + \hat{d} \tag{5}$$
$$x = X + \hat{x} \tag{6}$$
$$y = Y + \hat{y} \tag{7}$$
$$u = U + \hat{u} \tag{8}$$

In steady state

$$\dot{x} = AX + BU = 0 \tag{9}$$
$$X = -A^{-1}BU \tag{10}$$
$$Y = CX \tag{11}$$

We have a continuous, time-varying system, as the switch has two states. The state matrix A is represented by two matrices, A_1 and A_2 representing the switch-closed and switch-open conditions. The variable D weights the two condition state matrices. Next we write

$$\dot{\hat{x}} = A\hat{x} + B\hat{u} + [(A_1 - A_2)X + (B_1 - B_2)U]\hat{d} \tag{12}$$
$$\hat{y} = C\hat{x} + (C_1 - C_2)X\hat{d} \tag{13}$$

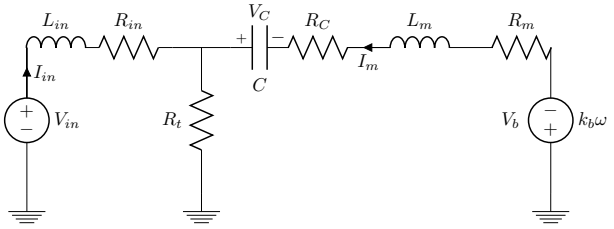

Fig. 13. Circuit with driving switch in the ON state.

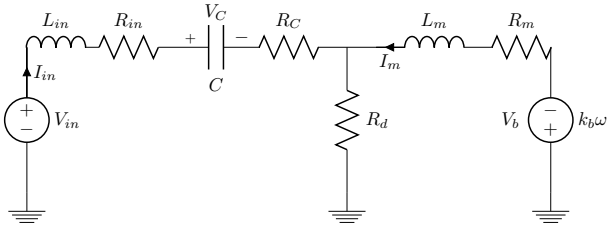

Fig. 14. Circuit with driving switch in the OFF state.

where

$$A = DA_1 + (1 - D)A_2 \tag{14}$$
$$B = DB_1 + (1 - D)B_2 \tag{15}$$
$$C = DC_1 + (1 - D)C_2 \tag{16}$$

A. On state

When the switch is closed the circuit becomes that shown in figure 13. Application of Kirchoff's laws yields:

$$\frac{dI_{in}}{dt} = \frac{1}{L_{in}}V_{in} - \frac{R_{in} + R_t}{L_{in}}I_{in} - \frac{R_t}{L_{in}}I_m \tag{17}$$
$$\frac{dI_m}{dt} = \frac{1}{L_m}V_C - \frac{k_b}{L_m}\omega - \frac{R_m + R_C + R_t}{L_m}I_m - \frac{R_t}{L_m}I_{in} \tag{18}$$
$$\frac{dV_C}{dt} = \frac{-1}{C}I_m \tag{19}$$
$$\frac{d\omega}{dt} = \frac{k_t}{J}I_m - \frac{b}{J}\omega - \frac{T_L}{J} \tag{20}$$

B. Off state

When the switch is closed the circuit becomes that shown in figure 14. Again the application of Kirchoff's laws yields:

$$\frac{dI_{in}}{dt} = \frac{1}{L_{in}}V_{in} - \frac{R_{in} + R_C + R_d}{L_{in}}I_{in} - \frac{R_d}{L_{in}}I_m - \frac{1}{L_{in}}V_C \tag{21}$$
$$\frac{dI_m}{dt} = -\frac{k_b}{L_m}\omega - \frac{R_m + R_d}{L_m}I_m - \frac{R_d}{L_m}I_{in} \tag{22}$$
$$\frac{dV_C}{dt} = \frac{1}{C}I_{in} \tag{23}$$
$$\frac{d\omega}{dt} = \frac{k_t}{J}I_m - \frac{b}{J}\omega - \frac{T_L}{J} \tag{24}$$

C. State Space Matrices

We are finally able to derive the two matrices

$$A_1 = \begin{bmatrix} -\frac{R_{in}+R_t}{L_{in}} & -\frac{R_t}{L_{in}} & 0 & 0 \\ -\frac{R_t}{L_m} & -\frac{R_m+R_C+R_t}{L_m} & \frac{1}{L_m} & -\frac{k_b}{L_m} \\ 0 & -\frac{1}{C} & 0 & 0 \\ 0 & \frac{k_t}{J} & 0 & -\frac{b}{J} \end{bmatrix} \quad (25)$$

$$A_2 = \begin{bmatrix} -\frac{R_{in}+R_C+R_d}{L_{in}} & -\frac{R_d}{L_{in}} & -\frac{1}{L_{in}} & 0 \\ -\frac{R_d}{L_m} & -\frac{R_m+R_d}{L_m} & 0 & -\frac{k_b}{L_m} \\ \frac{1}{C} & 0 & 0 & 0 \\ 0 & \frac{k_t}{J} & 0 & -\frac{b}{J} \end{bmatrix} \quad (26)$$

and

$$B = B_1 = B_2 = \begin{bmatrix} \frac{1}{L_{in}} & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & \frac{T_L}{J} \end{bmatrix} \quad (27)$$

$$C = C_1 = C_2 = \begin{bmatrix} 0 & 0 & 0 & 1 \end{bmatrix} \quad (28)$$

which will permit calculation of the performance of the circuit using a tool such as Matlab.

REFERENCES

[1] Slobodan Cuk, and R. D. Middlebrook, "Advances in Switched-Mode Power Conversion Part I," *IEEE Transactions on Industrial Electronics*, Vol. IE-30, no. 1, pp10–19, 1983.

[2] B. Bryant and M. K. Kazimierczuk, "Derivation of the Ćuk PWM DC-DC converter circuit topology", *Proceedings of the 2003 International Symposium on Circuits and Systems*, ISCAS '03, vol. 3, ppIII-292–III-295.

[3] B-R. Lin, C-L. Huang and J-F. Wan, "Analysis of a Zero Voltage Switching Cuk Converter", 33rd Annual Conference of the IEEE Industrial Electronics Society, IECON 2007, pp1972–1977.

[4] IEEE Xplore digital library, http://ieeexplore.ieee.org/, searched October 2016.

[5] John McLeish, Howell Round, and Jonathan Scott, "Speed Control with Low Armature Loss for Very Small Sensorless Brushed DC Motors", *IEEE Transactions on Industrial Electronics*, Vol 56, no. 4, April 2009, pp1223–1229.

[6] P. R. K. Chetty, "Modelling and Analysis of Ćuk Converter Using Current Injected Equivalent Circuit Approach," *IEEE Transactions on Industrial Electronics*, Vol IE-30, no. 1, February 1983, pp56–59.

[7] W. M. Polivka, P. R. K. Chetty, and R. D. Middlebrook, "State-Space Average modelling of converters with parasitics and storage-time modulation," *IEEE Power Electronics Specialists Conference*, June 1980.

[8] G.F. Franklin, J.D. Powell and A. Emami-Naeini, *Feedback Control of Dynamic Systems*, 7th Edition, Pearson, 2015.

Real Time Applications of Electrochemical Impedance Spectroscopy - A Technical Assessment

Arunkumar Jayakumar*, Tejas Trivedi†, Deepak Sharma†, Marilou Rebosura† and Latha Karthigaa Murugesan†
*School of Engineering
AUT University
Email: ajayakum@aut.ac.nz
†Department of Electrical Engineering
International College of Auckland
Email: dee056a, de001, dee109, latha.karthigaa@ica.ac.nz

***Abstract*—DC power systems such as solar cell, Fuel cells and batteries are expected to play a remarkable role for a sustainable future of the society. Electrochemical Impedance Spectroscopy (EIS) or Dielectric spectroscopy is a powerful technique for determining the characteristic of these DC electrical systems where impedance is measured over a range of frequencies to have a greater insight on their system behaviour. The application of EIS is numerous and the present paper assess the system characteristics considering solar cells and proton exchange membrane (PEM) fuel cells into consideration.**

***Keywords:* Electrochemical Impedance Spectroscopy (EIS), solar cell, Fuel cells, batteries, Nyquist, Bode plot**

I. Introduction

Power is the fundamental component of our day today activities. Soon, there will be a period where DC power systems would return back as a promising solution for Distributed Power generation systems. Fuel cells, solar cells and batteries belong to that class of DC systems and the present paper deals with the application of, Electrochemical Impedance Spectroscopy (EIS) to understand and characterize the energy storage and dissipation properties of the system under investigation. It measures the dielectric properties of a medium as a function of frequency. The frequency response of these systems are based on the interaction of an external field with the electric dipole moment of the sample, often expressed by permittivity. The common applications of EIS are in the process to solid electrolytes, electro ceramics, including polymers and glasses, dielectrics, and to devices that uses conversion of integrated energy such as Fuel cells, solar cells and batteries [1]. Electro-Chemical Impedance Spectroscopy has advantages such as its practicality in high resistance material and is a non-invasive technique with availability of time dependant and quantitative data [2]. It is also the key technique used in electrochemistry in which processes of electrodes are investigated, and for further explanation in solid-state devices, the techniques that often used is admittance spectroscopy, but in dielectrics components it is called as dielectric spectroscopy. However, it should be observed that it is not sufficient to just use EIS alone, but rather to integrate that knowledge of impedance in EIS with the other methodologies like biochemical and electrochemical or any other specific methodology that is under investigation [3]. It could lead us for better understanding of the performance incorporating complex mathematical equations that can be represented and simplified with the help of Nyquist and Bode plot [4].

II. Measurement Technique

Electrochemical Impedance Spectroscopy [5] is a sensitive instrument and hence the measurement has to be performed prudently. During an impedance measurement, a Frequency Response Analyser (FRA) [6] is used to impose a small amplitude AC signal to the electrochemical systems (Fig. 1).

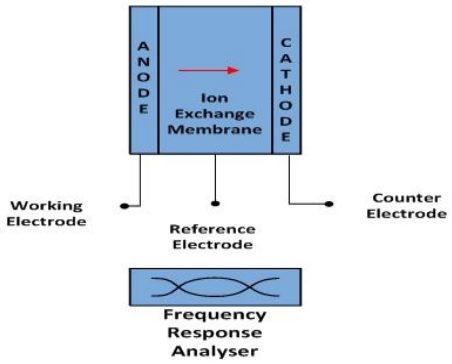

Fig. 1. Experimental Setup of Impedance measurement

The system under investigation is ideally kept at a fixed steady state by imposing stationary constraints such as the DC current, illumination intensity, etc., and the $Z(\acute{\omega})$ is measured by scanning the frequency at a multitude of values. Let us assume we have an electrical element to which we apply an electric field $E(t)$ and get the response $I(t)$, then we can disturb this system at a certain field E with a small perturbation dE and we will get at the current I a small response perturbation dI. In the first approximation, as the perturbation dE is small, the response dI will be linear. If we plot the applied sinusoidal signal on the x-axis of a graph and the sinusoidal response signal $I(t)$ on the y-axis, an oval is plotted. This elliptical figure is termed as "Lissajous figure" and is shown in Fig. 2. Analysis of Lissajous figures on oscilloscope screens was the accepted method of impedance

measurement prior to the availability of lock-in amplifiers and frequency response analysers.

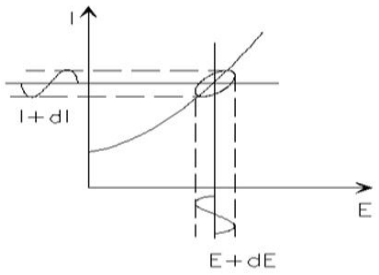

Fig. 2. Lissajous figure

The excitation signal, expressed as a function of time, has the form of:

$$E(t) = E_0\, cos(\omega t) \tag{1}$$

$E(t)$ is the potential at time t_r, E_0 is the amplitude of the signal, and ω is the radial frequency. The relationship between radial frequency ω (expressed in radians/second) and frequency f (expressed in Hertz $(1/sec)$).

$$\omega = 2\pi f \tag{2}$$

The defining relation and impedance for ideal bulk electrical elements are revealed in Table I.

TABLE I
ELECTRICAL ELEMENTS - REPRESENTATION AND FORMULA

Electrical Elements	Symbol	Formula	
Resistor		$V = IR$	$Z_R = R$
Capacitor		$I = C\,[dV/dt]$	$Z_C = 1/j\omega C$
Inductor		$V = L\,[dI/dt]$	$Z_L = j\omega L$

Equivalent Circuit modeling of EIS data is used to extract physically meaningful properties of the system under investigation (Fuel cell in the present case) by modeling the impedance data in terms of an electrical circuit composed of ideal resistors (R), capacitors (C), and inductors (L). As we are dealing with real systems that do not necessarily behave ideally with processes that occur distributed in time and space, we often use specialized circuit elements. These include the generalized Constant Phase Element (CPE) [7] and Warburg element (Z_W) [8]. The Warburg element is used to represent the diffusion or mass transport impedances of the cell. An example of a generalized equivalent circuit element for a single cell Fuel cell is shown below in Fig. 3.

III. AC IMPEDANCE SPECTRUM

This section covers the features of the spectra measured in a two electrode configuration with hydrogen passing through anode and oxygen or air passing through cathode.

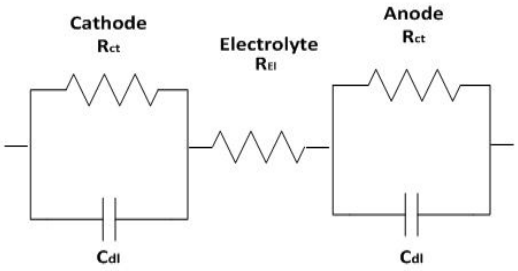

Fig. 3. Equivalent circuit of a PEM fuel cell [R_{ct}= Charge Transfer Resistance & C_{dl}= Double layer capacitance]

A. Galvanostatic and Potentiostatic measurements

To measure the impedance of a PEM fuel cell system, a Frequency Response Analyser (FRA) and a load bank are needed. The FRA is connected to a programmable DC load bank to carry out the measurements. A sinusoidal wave generated by the FRA is applied to the Fuel cell through the electric load bank. This AC signal is superimposed on the DC current. AC signals are fed back to the FRA and impedance spectra are generated [9]. A typical connection between the FRA and the Fuel cell for the impedance measurement is shown in Fig. 1. In addition, an oscilloscope may be used to monitor the noise level. Using the connection shown in Fig. 2, AC impedance can be measured either in a potentiostatic or in a galvanostatic mode. Basically, there are no significant differences between the results from the two modes. Each has its advantages and disadvantages, so depending on the application, one may be more suitable than the other. The measurements were carried out in a conductivity cell. The schematic view of conductivity cell is presented in Fig. 1.

B. Interpretation from Plot

The impedance can be illustrated as a vector of length $|Z|$ on the Nyquist plot. There is an angle between this x-axis and the vector, which is commonly entitled the "phase angle" ($f = arg\, Z$).

$$Z = \frac{E}{I} = Z_0\, \exp(i\phi) = Z_0(cos\,\phi + i\, sin\,\phi) \tag{3}$$

But where the Nyquist plot is very much effective for that system, the Nyquist Plots have one of major drawback or short come. Whenever you will look at any of the point of data which is on the plot, you will not able to tell that to record that point where frequency was used. Nyquist plot results from the electrical circuit which is shown in Fig. 4.

1) Nyquist plot: The Nyquist is normally expressed as $Z(\omega)$ which is composed of two parts real part and imaginary part. In case the imaginary part is plotted on the y-axis and the real part is plotted on the x-axis of a graph or chart, here we got a "Nyquist Plot" which is shown in Fig. 4b for the equivalent RC element shown in Fig. 4a [10]. There plot is a semicircle and reveals the characteristics of a one "time constant". Low frequency data are on the right side of the

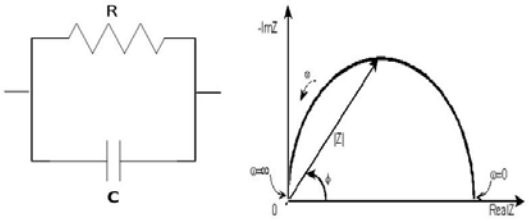

Fig. 4. a) Equivalent RC circuit b) Nyquist plot

plot and higher frequencies are on the left side. This is true for EIS data where impedance usually falls as frequency rises (this is not true of all circuits). The plot for electrochemical impedance spectroscopy often make or obtain different types of semicircles which might have different sizes (area).

2) Bode Plot: Bode Plot is the plot of the phase and magnitude of a transfer function or quantity of complex value, frequency. It has the phase which is in degrees and magnitude is on decibels are plotted versus frequency by having the use of axes of semi-logarithmic. The plot of magnitude is effectively log- log in nature, since which is describes in the form of decibels and other side frequency is based on the logarithmic axis. In the "Bode plot" the impedance is plotted with log frequency on the x-axis and both the absolute value of the impedance ($|Z| = Z_0$) and phase-shift on the y-axis [4]. The Bode plot for the RC circuit is shown in Fig. 5. Unlike the Nyquist plot, the Bode plot explicitly shows frequency information.

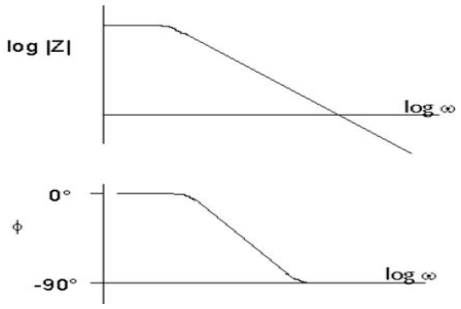

Fig. 5. Bode Plot

In the equivalent circuit analog, resistors represent conductive pathways for ion and electron transfer. As such, they represent the bulk resistance of a material to charge transport such as the resistance of the electrolyte to ion transport or the resistance of a conductor to electron transport. Resistors are also used to represent the resistance to the charge-transfer process at the electrode surface. Capacitors and inductors are associated with space-charge polarization regions, such as the electrochemical double layer, and adsorption/desorption processes at an electrode, respectively.

According to the testing system in the Real Time Application of Electrochemical Impedance Analysis bode plot is the most efficient way to make the testing plot which varies the magnitude, transfer function, sinusoidal excitation frequency, impedance phase. Real time systems are very complex or complicated, hence their analysis mostly leads to more complex derivations, sometimes extremely long where the chances of mistakes are always on the peak rating. If we will head up into the designing phase the long and complicated derivations are no longer in use. In contrary, we need to invert these problems so that manufacturer or designing engineer can work or choose the specific element values to produce a given particular behaviour which needs to be done. Including adding or changing element, make different types of new circuits, different values of elements. So we strongly need to launch the design-oriented analysis. There are some features of analytical expression which are very useful in design as well as lead to make simple equations and for the asymptotes. To make it simple ways there is an approximated method which we are describing in this report "Bode Plot". If we completely follow the rules and instruction which comes up with this method then we can reduce our work such as algebra, mathematical errors and complexity.

IV. Discussion and Interpretation

The AC voltage and current response of the electrochemical system (Fuel cell, solar cell or battery) under investigation is analysed by the FRA to determine the resistive, capacitive and inductive behaviour of the system at that specific frequency. Physicochemical processes occurring within the cell electron & ion transport, gas & solid phase reactant transport, heterogeneous reactions, etc. have different characteristic time-constants and therefore are exhibited at different AC frequencies [11]. When conducted over a broad range of frequencies, impedance spectroscopy can be used to identify and quantify the impedance associated with these various processes.

Electrochemical impedance spectroscopy (EIS), a diagnostic tool that has been widely used in studies of electrochemical systems such as battery and electrolytic cells, has been utilized by an increasing number of researchers in recent years. It is very much evident that analysing the raw measured signal, besides the classical Nyquist and Bode plots, is a very straightforward and easy to use tool which can help the user to make sure that the linearity, stability and causality conditions. Thus in the part of designing where we need to make of different types of circuit we can give the greater insight to it behaviour.

References

[1] A. Lasia, "Electrochemical impedance spectroscopy and its applications," in *Modern aspects of electrochemistry*, pp. 143–248, Springer, 2002.

[2] M. E. Orazem and B. Tribollet, *Electrochemical impedance spectroscopy*, vol. 48. John Wiley & Sons, 2011.

[3] S.-M. Park and J.-S. Yoo, "Peer reviewed: electrochemical impedance spectroscopy for better electrochemical measurements," *Analytical chemistry*, vol. 75, no. 21, pp. 455–A, 2003.

[4] Z. He and F. Mansfeld, "Exploring the use of electrochemical impedance spectroscopy (eis) in microbial fuel cell studies," *Energy & Environmental Science*, vol. 2, no. 2, pp. 215–219, 2009.

[5] C.-N. Cao and J.-Q. Zhang, "An introduction to electrochemical impedance spectroscopy," *Science, Beijing*, vol. 21, 2002.

[6] H. Allison, "Frequency response analyzer," Mar. 30 1982. US Patent 4,322,806.

[7] J.-B. Jorcin, M. E. Orazem, N. Pébère, and B. Tribollet, "Cpe analysis by local electrochemical impedance spectroscopy," *Electrochimica Acta*, vol. 51, no. 8, pp. 1473–1479, 2006.

[8] A. M. Dhirde, N. V. Dale, H. Salehfar, M. D. Mann, and T.-H. Han, "Equivalent electric circuit modeling and performance analysis of a pem fuel cell stack using impedance spectroscopy," *IEEE Transactions on Energy Conversion*, vol. 25, no. 3, pp. 778–786, 2010.

[9] X. Yuan, H. Wang, J. C. Sun, and J. Zhang, "Ac impedance technique in pem fuel cell diagnosisa review," *International Journal of Hydrogen Energy*, vol. 32, no. 17, pp. 4365–4380, 2007.

[10] Q. Wang, J.-E. Moser, and M. Grätzel, "Electrochemical impedance spectroscopic analysis of dye-sensitized solar cells," *The Journal of Physical Chemistry B*, vol. 109, no. 31, pp. 14945–14953, 2005.

[11] B. E. Conway, *Electrochemical supercapacitors: scientific fundamentals and technological applications*. Springer Science & Business Media, 2013.

A New Framework for Interactive Control of Mechatronic Instruments

Blake Johnston
Te Kōkī New Zealand School of Music
Victoria University
Wellington, New Zealand
blake.r.jo@gmail.com

Bridget Johnson
School of Music and Creative Media Production
Massey University
Wellington, New Zealand
b.d.johnson@massey.ac.nz

Ajay Kapur
Te Kōkī New Zealand School of Music
Victoria University
Wellington, New Zealand
kapuraj@ecs.vuw.ac.nz

***Abstract*—Over recent years there has been significant developments in the burgeoning field of mechatronic musical instruments. These developments have seen mechatronic instruments greatly increase in complexity, and concomitantly, their expressive capabilities. However, the development of designing control systems for these instruments has been lacking, meaning that the use of these expressive new instruments has largely being left to the builders of the instruments themselves, or other especially skilled users. This paper proposes a new framework for the control of mechatronic musical instruments that is designed to afford an intuitive interaction for all users. This design allows for more aesthetic exploration of mechatronic instruments as well as expanding the potential demographic of users. This framework explores the potential of new iPad applications that offer high level control parameters and gestures, for the real-time interaction with mechatronic instruments in both performance and installation settings without assuming prior mechanical knowledge of the user.**

This paper discusses the current state of the field and the need for this new framework. It then demonstrates and explains three case studies that have being developed by the authors as examples of this new framework.

Keywords— Control Systems, Mechatronics, Musical Robotics, Touch Interfaces, Control Framework, Interaction Design, Installation Audience Interaction Introduction (*Heading 1*)

I. Introduction

Mechatronic musical instruments have recently seen a period of rapid development in their functionality and musically expressive qualities. However, this development has often taken part purely on the instrument side, with significantly less development taking place on the control system side of the instruments. In order for mechatronic instruments to be used by a wide range of composers and performers, as well as non-expert users in an installation setting, the authors recognise a need for development in the user control systems for mechatronic instruments.

This paper begins with an assessment of the field of mechatronic instrument development, and the use of iPad applications as instruments and expressive musical tools. The development of custom iPad applications as musical interfaces is a key aspect of the new framework. Following this, the paper goes on to discuss the motivation for developing the proposed framework. After discussion of the motivating factors, the framework itself is introduced with detail provided of how the framework operates, and why each element and communication method was chosen. Finally, the paper provides three case-studies, each of which has implemented the proposed framework. These case studies include: *speaker.motion*, a mechatronic loudspeaker; *Carme*, a mechatronic string ensemble; and *mecha.space*, a control system for spatialised mechatronic percussion instruments. The paper concludes with discussion of the future direction of the framework.

II. Related Works

A. IPad and IPhone as interfaces

In the past decade the increased commercial popularity of multi-touch products has seen interfaces of this type begin to gain popularity as performance interfaces for live electronic music. In particular Apple's iPad and iPhone have proved a popular interface to develop new applications for musical expression. This popularity might be attributed to a number of factors: they have a relatively low purchase cost; their multi-touch technology is fast, accurate and reliable, particularly when compared to early large-scale multi touch tables; they require no calibration; and they are not affected by stage lighting. This set of factors make these commercial touch interfaces well-suited platforms with great potential for exploration as a musical interface.

Many artists and developers from the NIME (New Interfaces for Musical Expression), an impor community have built expressive interfaces using this platform. Ge Wang's Magic Flute [1] and Magic Fiddle [2] are new instruments

conceptually based on the design of an acoustic flute and fiddle. They were designed to fully utilize the capabilities of the tablet and smart phone technologies. Their control systems include: touch events, accelerometers, breath control (through microphone input), and GPS location. Even though these instruments are designed to replicate pre-existing acoustic instruments, they also extend their capabilities through utilizing the features of the new technology. For example, the potential for multi-user networked performance. There is also a vast array of synthesis applications written for touch screens that are used by the novice user, as well as by professional musicians, such as the Arturia iMini, Moog Minmoog and the Korg iMS-20.

A further trend in the use of portable touch-screen technologies is as a control interface (rather than a full instrument). Many music software programs are now capable of receiving control information from touch interfaces. For example, the Logic Remote app allows direct communication with Logic Pro X, a popular digital audio workstation (DAW) that allows for comprehensive editing of audio. This extends the Logic Pro X control interface for recording, mixing and performance purposes. Rather than generate audio on the app itself, the app uses the technology embedded in the iPad to send communication data back to the computer to control Logic Pro X. There are also multiple control interfaces available to interact directly with Ableton Live, another popular DAW, in both studio and performance settings, many of which are graphical representations of popular hardware interfaces.

This short list of examples shows that portable multi-touch technologies have quickly found their place as expressive interfaces both in the NIME community and in the commercial realm. The widespread adaptation of this technology makes these interfaces a desirable platform to work with in the context of the goals of the framework presented in this paper, as it has a high potential for other members of the mechatronic community to adopt and develop the proposed framework.

B. The Mechatronic Instrument Field

A main driver for research in the mechatronic musical field has been affording musical expressivity through their design. This has created a proliferation of control parameters and complexity in these systems, affording more expressivity. However, this has made the composer-mechatronic system interaction more complex and difficult. Rapid developments made in the design of these instrument systems have not been matched in the control systems, which as Jim Murphy identifies, means that '[t]o compose for these systems in their current state is to manually direct every action that the robot undertakes. To write music in this manner is quite time-consuming, requiring much actuator management rather than higher-level musical composition.' [3]

The common control solution that most designers of mechatronic systems use is to implement a custom MIDI communication framework. The MIDI communication protocol has its benefits as it is widely used (particularly in music communities), with many commercial hardware controls available, as well as many software systems that can output MIDI. The protocol is reliable, versatile, common among musicians, and simplistic, which makes it well suited to communicating control parameters to a musical mechatronic system.

However, the commercial hardware MIDI controllers available are often ill-suited to controlling the specific range of parameters that are afforded by complex mechatronic systems. The keyboard design paradigm that many MIDI keyboards are based on offers a fine degree of control over discrete pitches, with dynamics, polyphony and physical tactility all important design features. However, in many cases, this design paradigm doesn't match the specific needs of a complex mechatronic system.

This lack of suited hardware controllers often mean that composers use software to generate MIDI to control a mechatronic system, which is more customisable than its hardware counterparts. This method has its benefits, as composers can carefully control many different parameters at once by transcribing MIDI messages, and then playing their compositions back. This also makes compositions repeatable, and mechatronic builders can implement communication frameworks that can be catered for. A good example of this is Godfried-Willem Raes's research with the Logos Foundation, who have developed a large robot orchestra, and provide a comprehensive MIDI manual for each of their robotic instruments. [4]

C. Bespoke Hardware Controllers for Mechatronic Systems

There are very few examples of custom-built hardware controllers used to communicate with mechatronic instruments. Jordan Hochenbaum and Owen Vallis have designed a large multi-touch interface called Bricktable, which they used to control Ajay Kapur's MahaDeviBot, a twelve-armed percussion robot. [5] The touch interface sent positional and rotational data of tangible objects that were placed on the touch surface to MahaDeviBot, controlling rhythms and intensity of each strike. Bricktable was not specifically designed to control mechatronic systems, however, they were able to customise the user interaction to control the MahaDeviBot.

Ajay Kapur has also interacted with mechatronic systems in his piece Digital Sankirna. The work is for Kapur's own ESitar [6], a custom built hyperinstrument, that combines a sitar with sensors that detect different aspects of the performers movement and behaviour; and for the Machine Orchestra [7], a mixed ensemble of human and robotic performers developed at the California Institute of the Arts. Through using data from the sensors on the ESitar, as well as real-time analysis of his ESitar performance, streams of quick notes are distributed throughout the robotic ensemble, creating ripples of percussive strikes to accompany his sitar playing.

In a 2013 interview with the LAWeekly, Trimpin described an iOS application he has developed that allows the user to control an ensemble of his mechatronic sound sculptures [8]. The application allows for different MIDI sequences to be sent to different instruments, allowing for the control of an ensemble from one point. The application itself is specific to Trimpin's creations, and as such, not widely adaptable for other systems.

III. Motivation

As discussed in Section 2, the authors felt that portable multi-touch devices offered a great range of potential for development as a control interface for musical expression. In particular the iPad is well suited with its larger screen size, its reliability and its ease of use, especially in configuring communications over a wireless network. Furthermore, it is also a multi-functioning affordable tool that many users may already own. This was a significant factor in the decision of which design platform to implement for this new framework. It needed to be something that would allow control in multiple performance settings without requiring new hardware for each different mechatronic system. Mechatronic musical instruments themselves often require significant set up and time, and incorporate a lot of hardware to set up, so the control interface needed to be something that could be incorporated in the system without significantly adding to the complexity of the overall system.

Another factor was considering the potential for use by users coming from multiple backgrounds. It was important that there is the both the potential for nuanced, complex control that will allow an experienced performer to become highly proficient and use the platform for sophisticated musical performance; as well as being able to be implemented in an installation setting where a novice user could quickly engage with the system. This would allow them to interact with a mechatronic instrument, quickly but with enough complexity to create meaningfully musical gestural relationships.

The authors also felt that it was imperative for the development of the mechatronic musical field to develop the control elements of the robot/human interaction, to fully explore the possibilities that this research can afford.

IV. Framework Overview

The new framework developed by the authors can be broken into three major parts. The first is the control interface, which takes the form of custom iPad apps. Though each individual system has its own app developed specifically for it, there are many common features that allow the apps to fit easily within the wider frameworks. This also means that in building each new application the structure for the iOS code, can be used for each app, significantly reducing the time needed to develop each new app. The apps all send OSC data out over a wireless network, which is setup through a configuration screen that allows the user to input the IP address and port number for communication to be sent. Depending on the needs of the application, this screen can also feature other configurable settings as needed.

The OSC (Open Sound Control) [9] protocol was chosen for its ease in customisation. Each application can easily implement its own protocol based on the control data needed for the particular systems. The OSC data is then received in a customisable Max/MSP patch [10]. The Max patch essentially works as a protocol translator - receiving OSC data and translating it to the appropriate MIDI data that is needed to control the mechatronic instrument. Depending on the needs of the performance and the system, the Max patch might also generate further musical information. For example, in the mecha.space system, the Max patch also generates sequences of MIDI information that can be sent as loops to the instruments. The control interface then triggers each loop to start, and determines which instruments should play which parts of the sequence.

Max/MSP was chosen for its ease of use both in receiving OSC data and sending MIDI data, but also for its potential to be used to generate complex musical information be it in the form of audio or MIDI data. As has being identified in Section 2, MIDI is the most common control protocol implemented by mechatronic musical instruments. It was a goal of the framework to be compatible with both future and past mechatronic instruments and the use of the MIDI protocol ensures this is possible. The full framework is displayed in diagrammatic form below.

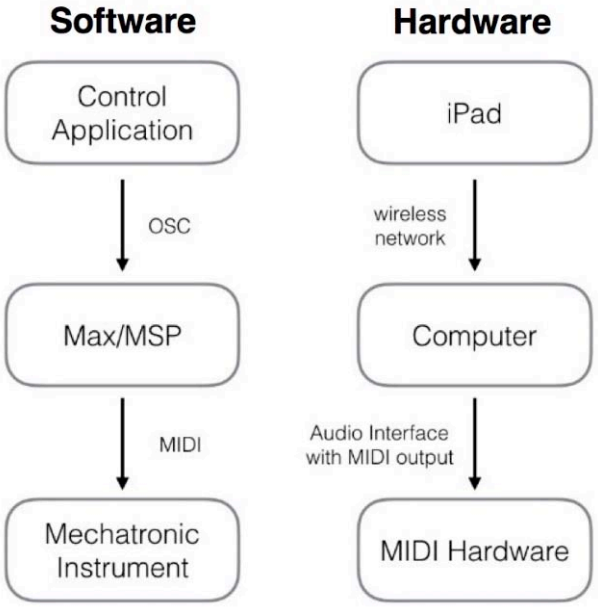

Fig. 1. Overview of New Framework

V. Case Studies

The following section will outline three case studies developed by the authors that are examples of the proposed framework. Each of these systems has being developed for both performance and installation settings. They all feature customised versions of each element of the framework built off common base foundations and communication protocols.

A. speaker.motion

speaker.motion is a mechatronic loudspeaker system that allows dynamic repositioning of the loudspeaker in real-time [11]. The loudspeaker can be fully rotated in either direction indefinitely, as well as, tilted 180 degrees. This allows its directionality to be manipulated to create complex spatial patterns and trajectories, as well as, activating the physical space in varying ways. The speaker.motion iPad application was designed so a performer or installation user could intuitively control the angle and tilt of the loudspeaker in real-time. The full speaker.motion system features four loudspeaker

units each of which can be controlled independently and simultaneously by the app.

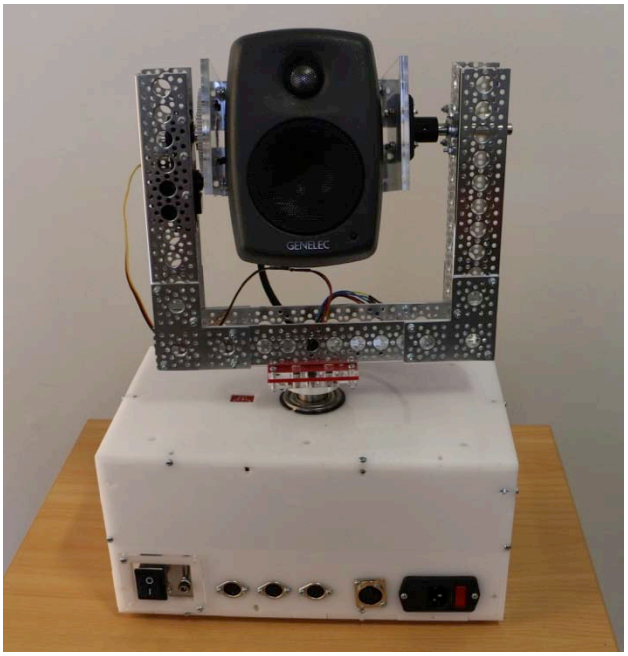

Fig. 2. *speaker.motion* mechatronic loudspeaker.

The design of the iPad app was driven by the intention of maintaining highly intuitive gestural relationships between the movement of the speaker and the physical movement the user performs. The app shows graphical representations of all four loudspeakers displayed as a series of consecutive circles. The user interacts by moving a small ball around the graphical space representing the directionality with which the speaker should aim itself. The radial movement is very intuitive as the position inside the circle that the ball is placed directly correlates to the radial direction of the loudspeaker. The tilt is controlled by the balls distance from the centre of the circle. The outer most circle will cause the speaker to tilt towards the floor, the inner circle will cause the speaker to tilt directly up, and all other positions are calculated therein.

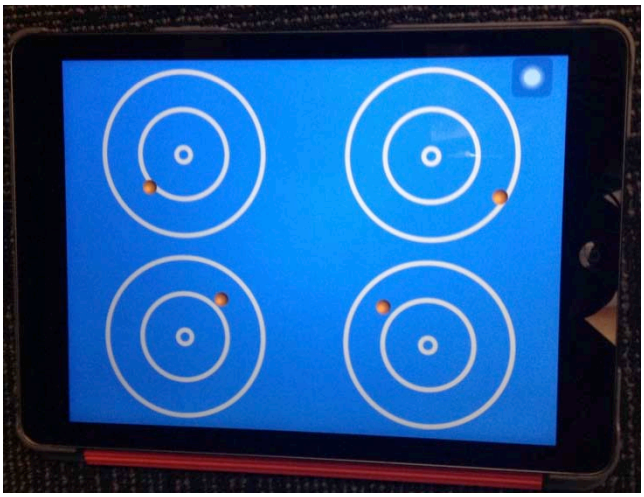

Fig. 3. The speaker.motion iPad Application

The speaker.motion iPad application and its implementation of the proposed framework means that the mechatronic loudspeakers can be interacted with in a gestural and intuitive way by performers or non-expert installation users. The direct mapping of the app to the physical movements of the loudspeaker themselves means that the instruments can be used easily by musicians and other users without any understanding of how the mechanics of the system works. The proposed framework removes the need for specialised knowledge in interacting with this mechatronic ensemble and opens speaker.motion up as an expressive tool to a much wider and more diverse range of users in both performance and installations settings.

B. Carme

Carme [12] is an application designed to control The Polus Ensemble [13], an ensemble of mechatronic, bowed string sculptures. The ensemble contains six instruments, each with a single string that is excited by a rotating bowing wheel mechanism. The design of the bowing mechanism allows for the control of the speed of rotation of the wheel, direction of rotation, and the pressure of the bow onto the string through controlling the swing arm that rotates perpendicular to the bowing wheel.

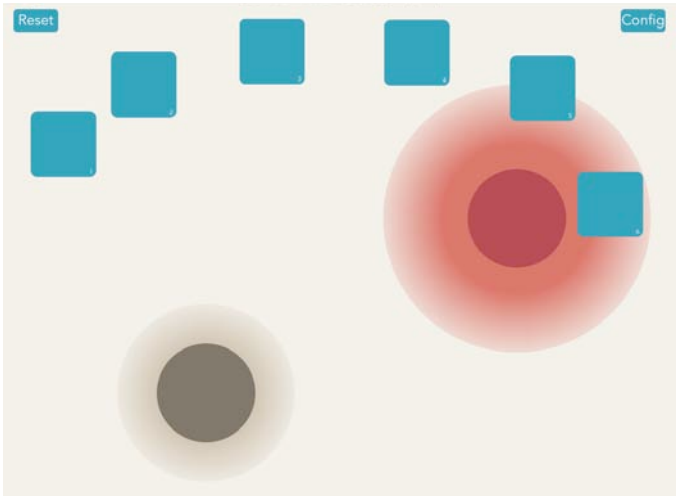

Fig. 4. The Carme Application User Interaface.

Due to each unit having four raw control parameters, high-level musical behaviours were abstracted to allow a more intuitive interaction. This means if a user wanted to create a continuous note, they would not have to set the direction of the arm and adjust the speed, as well as, control the speed of the bowing wheel. Instead these behaviours offer abstract control parameters like intensity, which are interpreted in the Max/MSP application to control these parameters accordingly.

Circles and squares represent musical behaviours and the instruments of the Polus Ensemble respectively on the iPad user interface. By dragging a musical behaviour in close proximity to an instrument, the relevant behaviour is performed by the instrument. This design can easily be customised to have any amount of instruments and behaviours on the screen, catering to the specific mechatronic system. This behaviour is only triggered when the shapes overlap. The degree of

proximity is used to control a parameter within the behaviour. An example would be the intensity of a continuous note.

Fig. 5. Polus Ensemble controlled by Carme, the iPad app.

Each shape is not fixed in space and may be grouped in any way. This allows for an interesting approach to organising sound, as sculptures can be spatially grouped together to create chords, or separated to make small sections in the ensemble. Using a spreading gesture with two fingers can alter the size of the behaviours. This allows for the influence of a behaviour to dilate and cover a large space. By making the behaviours larger, the user can have finer control over the proximity, as the resolution is effectively higher.

While to some degree, this approach of high-level musical behaviours abandons the possibility of fine-grained control over every possible control parameter; it has the benefit that it greatly simplifies the user experience, while still allowing for complex behaviours to result. This design choice was informed by the two different complexities that exist in The Polus Ensemble: controlling multiple instruments at once that are spread out through space, while treating them as one mechatronic system; and the multiple parameters of control that each sculpture presents.

Through using this framework, the complexities of the ensemble can be controlled in a nuanced way, which is both intuitive, and doesn't require a full understanding of how each of the units need to behave in order control musical behaviour.

C. mecha.space

mecha.space developed out of ideas about spatialisation in live electronic music. One of the focuses of research in acousmatic music is the dynamic spatialisation of sound across loudspeaker arrays. mecha.space seeks to explore these concepts of the spatialisation of musical ideas through an array of sound generators in the mechatronic realm.

The mecha.space application was designed for user interaction with a spatialised mechatronic percussion ensemble in both performance and installation settings. The graphical user interface allows the user to drag visual depictions of the physical instruments, represented as red squares into the space to coincide with where those instruments have been placed in the concert hall. Once satisfied that the instrument objects are placed correctly in the virtual space the user can 'lock' their positions so they aren't accidentally moved in further interactions. The user may then tap on the instrument object to send a single hit note to that particular instrument, or interact with a separate set of objects, represented by green circles to start a rhythmic sequence and dictate how this sequence is spread throughout the physical instruments. The user can dynamically change the size of the circle, ranging from very small, to taking up the full screen and the system is able to recognise which instruments fall within the scope of the object. Any instrument falling within the scope of the object will play the rhythmic sequence. The velocity, or loudness of each particular instrument is determined by its position within the object relative to the centre. Instruments on the edge of the objects parameter will receive only a low velocity and therefore play much quieter, instruments positioned right in the centre of the object will play at full loudness capabilities.

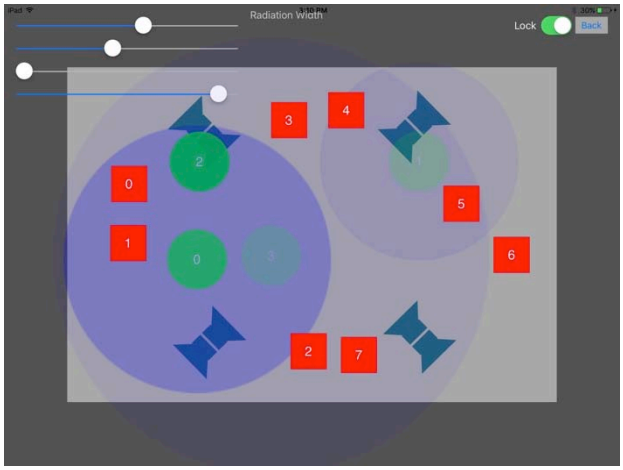

Fig. 6. The mecha.space iPad Application

In following with the proposed framework the mecha.space application has a start-up screen where the user can enter the details of the desired network and port to connect to. The data about the instruments is then sent out via OSC messages over the network. For a tap received on an instrument object a single message is sent with an ID corresponding to the instruments tapped. For the sequencing objects a message is sent for each instrument that falls within its scope, the message contains two ID tags, one for the instrument and one for the object. The message also contains a percentage value that will correspond to the velocity value, the percentage of how close to the centre of the object it is, therefore dictating the amplitude with which that instrument is played.

The sequencing takes place in Max/MSP where the MIDI data is produced. The user can easily create rhythms to sequence within Max that can be played by the mechatronic percussion ensemble. While mecha.space was designed to work with specific mechatronic percussion instruments, a major advantage of the framework is that it could very quickly and easily be adjusted for use with any mechatronic instruments that are controlled with MIDI.

VI. CONCLUSIONS AND FUTURE WORK

The strength of developing this framework lies in the ability to control multiple mechatronic instruments in a customisable, intuitive, and convenient way for a range of users from expert to novice. This expands the possibilities of these exciting new mechatronic systems to users beyond the

very small group of expert builders and makers. It also gives these users a convenient, flexible, and powerful way of interacting with their own mechatronic systems.

Future developments for this framework will be directed towards expanding the existing applications to other mechatronic systems. This will lead to development of how these touch interfaces can be interacted with in different modes to suit different systems. Also, the authors will be researching streamlining and providing a standard package of software which will allow for the quick customization of parameters and features so that many different mechatronic systems will be quickly compatible.

ACKNOWLEDGMENT

The authors would like to thank the New Zealand School of Music and the Victoria University of Wellington School of Engineering and Computer Science for their support of this research. Also Jason Long for his involvement in mecha.space, and Henry Dengate Thrush and Jim Murphy for their help with designing and building The Polus Ensemble.

REFERENCES

[1] G. Wang, "Ocarina: Designing the iPhone's Magic Flute", *Computer Music Journal*, vol. 38, no. 2, pp. 8-21, 2014.

[2] G. Wang, J. Oh, T. Lieber. “Designing for the iPad: Magic Fiddle.”, In *Proceedings of New Interfaces For Musical Expression. (NIME’11),* Oslo, Norway, 2011.

[3] J. Murphy, J. McVay, P. Mathews, D. Carnegie and A. Kapur, "Expressive Robotic Guitars: Developments in Musical Robotics for Chordophones", *Computer Music Journal*, vol. 39, no. 1, pp. 59-73, 2015.

[4] L. Maes, G. Raes and T. Rogers, "The Man and Machine Robot Orchestra at Logos", *Computer Music Journal*, vol. 35, no. 4, pp. 28-48, 2011.

[5] J. Hochenbaum and O. Vallis, "BrickTable: A Musical Tangible Multi-Touch Interface", in *Berlin Open Conference*, Berlin, 2009.

[6] A. Kapur, *et al.*. “The electronic sitar controller. *In Proceedings of New Interfaces For Musical Expression (NIME’ 04*), Singapore, 2004.

[7] A. Kapur, *et al*., “The Machine Orchestra: An Ensemble of Human Laptop Performers and Robotic Musical Instruments”, *Computer Music Journal* vol. 35, no. 4, pp. 49-63, 2011.

[8] M. Sallabedra, "A Chat With Trimpin, an Artist Who Combines Music, Apps and Robots", *L.A. Weekly*, 2013. [Online]. Available: http://www.laweekly.com/arts/a-chat-with-trimpin-an-artist-who-combines-music-apps-and-robots-4183107. [Accessed: 03- Oct- 2016].

[9] M. Wright, A. Freed, “Open sound control: A new protocol for communicating with sound synthesizers.” In *Proceedings of the 1997 International Computer Music Conference* Thessaloniki, Greece, 1997.

[10] D. Zicarelli, "Max/MSP Software." *San Francisco: Cycling ’74*, 1997.

[11] B. Johnson, M.Norris, A. Kapur, “speaker.motion: A Mechatronic Loudspeaker System for Live Spatialisation.” In *Proceedings of New Interfaces For Musical Expression (NIME’16),* Brisbane, Australia, 2016.

[12] B. Johnston, M. Norris, B. Johnson, A. Kapur, “Carme: A Multi-Touch Controller for Real Time Spatialisation of Gestural Behaviour” In *Proceedings for the Australasian Computer Music Conference (ACMC 15),* Sydney, Australia, 2015.

[13] B. Johnston, et al., “Polus: The Design and Development of a New, Mechanically Bowed String Instrument Ensemble. In *Proceedings of New Interfaces For Musical Expression (NIME’ 14),* London, United Kingdom, 2014.

Aspects of Quartz Operation in Viscous Fluid

Loren Kersey, Phill Brown, Daniel Schumayer and Timothy Molteno
Electronics group, University of Otago
Box 56, Dunedin, NEW ZEALAND
Email: kerlo392@student.otago.ac.nz

***Abstract*—A model for the oscillation of an AT-cut quartz crystal submerged in a weakly viscous fluid is developed and used to characterize the expected shifts in resonant frequency. These shifts are complex and predict a change in magnitude of around 0.034% in water. A trend for the quality factor deviation is proposed. An extension to include the oscillation amplitude is in progress. These characterizations are essential to predict the feasible limits of QCM operation in liquid and influence the selection of a frequency counter such that the shifts in frequency can be measured.**

I. Introduction

The detection of subtle physical changes at the surface of a piezoelectric crystal is emerging as a field of great potential, with established applications in sensing of fluid properties, mass accumulation and biomolecular adhesion [1]. Piezoelectric materials provide a very direct transduction mechanism through coupled mechanical and electrical behavior, i.e an applied alternating electric field induces mechanical oscillations and vise versa. The mechanical response greatly depends on the nature of any load to the crystal surface, affecting its resonance characteristics.

The versatility and stability of quartz makes it a popular piezoelectric material of choice, as it has a very high Q-factor and offers a vast collection of different cuts, some of which even exhibit temperature stability [2]. Quartz gained momentum (in the figurative sense) as an investigative device after the discovery by Sauerbrey in 1959 [3] that the resonant frequency of a quartz crystal varies linearly, to first order, with mass accumulation at the crystal surface. The Quartz Crystal Microbalance (QCM) is a device designed to exploit this relationship, capable of monitoring mass changes of the order of a nanogram [4].

It was initially thought that placing a QCM in contact with a comparatively large volume of liquid would result in significant loss in quality factor and cessation of the crystal oscillation. However countless experiments have since revealed this as physically realizable, an outcome which has significantly broadened the applications of the QCM [5]–[8]. The QCM is now widely recognized as an inexpensive biomoleular sensor, enabling precise determination of the analytes responsible for accumulation of foreign mass layers on the surface of the crystal. This application has driven a wide investigation into the operation of an electroded crystal unit placed in contact with a liquid containing analytes which bind themselves to a coating on the electrodes. This application has also influenced a drive to remove the electrodes from the surface of the crystal and excite it remotely, to maximize adhesion surfaces and enable contactless interrogation.

When an oscillating crystal is placed in contact with a liquid, it is not only the adhesion of molecules which produces shifts in resonance characteristics but the presence of the liquid itself imposes additional damping and dispersion which also alters the resonance characteristics. In fact, numerous studies have shown that a crystal oscillating in a shear mode drags a portion of fluid with it, inducing shear waves in the fluid. These shear waves couple to the shear waves in the crystal, resulting in a shift in resonance frequency which depends on the density-viscosity product of the liquid [7]–[9].

In many medical and food processing applications, a remotely queried viscosity sensor would greatly simply measurements. A remotely driven quartz crystal would provide an elegant solution, due to the dependence of resonance on fluid properties and the incorporation of the transducer and sensing element into one simple piece. Although the prospect of wirelessly exciting an unelectroded crystal with an antenna seems a bit far-fetched, numerous studies have already demonstrated successful attempts. A number of these studies have even reported wireless excitation of quartz crystals which are submerged in a liquid [10]–[15]. Despite the successful attempts, the separation distance between the crystal and antenna remains quite small, with the largest separations around 2cm [16]. Removal of the electrodes does however

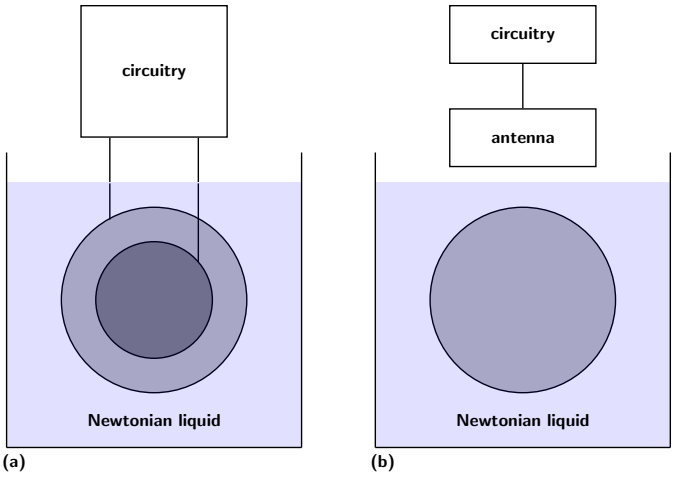

Fig. 1: (a) Schematic of an electroded crystal submerged in a Newtonian liquid. (b) Depiction of a wirelessly excited crystal.

show to improve the sensitivity of the oscillator, making the sensing area larger and more responsive [17]–[20]. In order to perform viscosity sensing with wirelessly excited quartz crystals we must first confirm that we can adequately model electroded quartz crystals operating in viscous fluid. A schematic representation of the two excitation mechanisms is shown in Figure 1. The present work will focus on the behavior of submerged electroded QCMs, the situation depicted by Figure 1 (a).

This is an in-progress work which encompasses the development of models that we intend to test experimentally on an electroded QCM, as well as a few experimental aspects of operating a QCM in a Newtonian liquid. Expected frequency shifts, changes to the Q-factor and expected oscillation amplitudes are investigated. These characteristics are dependent on the fluid properties, thus providing a sensing mechanism. Knowledge of these properties also is essential to making an informed decision about feasible limits of operation and frequency detection device requirements.

II. Model

Due to the anisotropic nature of quartz, the behavior of an individual unit strongly depends on its geometry and how it was cut from the bulk crystal. The most common crystal cut used in QCMs is the AT-cut, which is a crystal unit cut with its faces at an angle of $35.25°$ to the optical axes. This particular cut has negligible temperature coefficient between $0°$C and $50°$C, exhibiting temperature stability around room temperature [21]–[23].

Our model will be based around a Cartesian coordinate system, where the faces of an AT-cut crystal lie in the X_1X_3 plane. The thickness of the crystal unit lies in the X_2 direction, with $x_2 = 0$ passing through the center of the crystal and the faces in the planes $x_2 = a$ and $x_2 =$ -a. A schematic representation of the crystal's orientation with respect to our Cartesian coordinate system is given on Figure 2. A translation of axes is performed on all material property tensors to account for the discrepancy between the crystal's internal axes and our Cartesian coordinate system.

For a crystal submerged in a fluid we will examine a multi layer system consisting of an AT-cut quartz crystal sandwiched between two fluid layers. The thickness of the electrodes will be neglected. We will examine the motion in the crystal and fluid individually, and then combine them with stress conditions at the boundary.

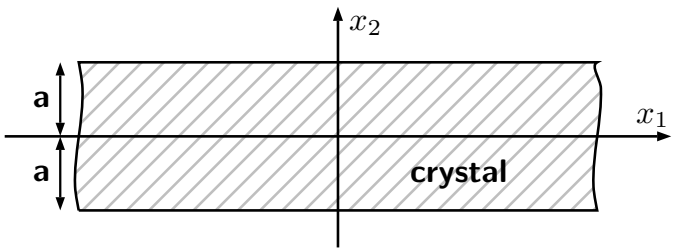

Fig. 2: Schematic representation of an AT-cut quartz crystal and our Cartesian coordinate system. The faces of the crystal lie in the X_1X_3 plane, with the thickness in the X_2 direction. The coordinate system originates in the center of the crystal.

A. Motion in crystal

In piezoelectric materials, deviations in electric field have accompanying mechanical perturbations. The piezoelectric constitutive relations connect the appropriate physical quantities:

$$T_{ij} = c_{ijkl}S_{kl} - e_{kij}E_k \tag{1a}$$

$$D_i = e_{ikl}S_{kl} + \epsilon_{ik}E_k, \tag{1b}$$

where **T** is the stress tensor, **c** is the isothermal elastic stiffness coefficient tensor, **S** is the strain tensor, **e** is the piezoelectric coefficient tensor, **E** is the electric field vector, **D** is the electric displacement vector and ϵ is the dielectric coefficient tensor. The strain is related to the displacement within the crystal by

$$S_{kl} = \tfrac{1}{2}\left(u_{k,l} + u_{l,k}\right), \tag{2}$$

where u_i is the displacement in the ith direction.
Assuming no external body forces, the equation of motion is derived from the principle that the resultant force on a body is equal to the rate of change of momentum

$$\rho\ddot{u}_i = T_{ij,j}, \tag{3}$$

and in the absence of free charges

$$D_{i,i} = 0. \tag{4}$$

The linear theory of anisotropic elasticity permits the presence of shear-horizontal waves, with the only non-zero displacement in the X_1 direction. Furthermore when the faces of the crystal are large compared to the thickness, variation in the magnitude of displacement in the plane of the crystals surface is negligible compared to variation across its thickness [24]. This reduces the equation of motion to one dimension, with $u_1 = u_1(x_2, t)$, $u_2 = u_3 = 0$. This translationally invariant motion generates an electric scalar potential, ϕ, which also only depends on x_2 and t, and the only non-zero component of the electric field vector can be written as $E_2 =$ -$\phi_{,2}$. Specifying (3) and (4) to the present case results in

$$c_{1212}u_{1,22} + e_{212}\phi_{,22} = \rho\ddot{u}_1 \tag{5a}$$

and

$$e_{212}u_{1,22} - \epsilon_{22}\phi_{,22} = 0. \tag{5b}$$

Assuming a simple sinusoidal time dependence and plane wave solutions the displacement takes the form

$$u_1 = A\mathrm{e}^{i(kx_2+\omega t)}. \tag{6}$$

Substituting this into (3) reveals the dispersion relation

$$k = \omega\sqrt{\frac{\rho}{c_{1212} + \frac{e_{212}^2}{\epsilon_{22}}}}. \tag{7}$$

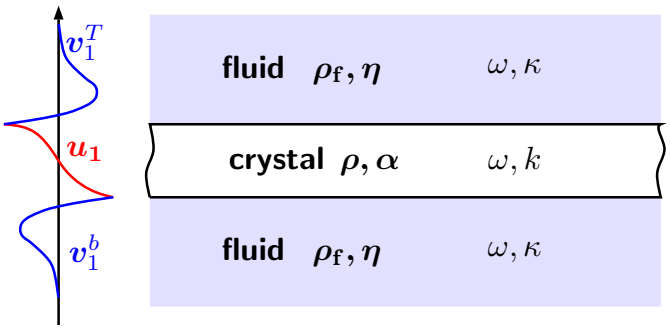

Fig. 3: Schematic representation of a quartz crystal submerged in a fluid of viscosity η. The graphs on the left show the qualitative behavior of the ideal displacement in the crystal, u_1 and the fluids, v_1^T and v_1^b.

The piezoelectrically stiffened elastic stiffness coefficient, $c_{1212} + \frac{e_{212}^2}{\epsilon_{22}}$, is often referred to as α for simplicity. The frequency, ω, and wave number k, are continuous parameters and appropriate boundary conditions provide conditions of resonance.

Omitting time dependence, equation (5b) can be integrated to find an expression for ϕ

$$\phi = \tfrac{e_{212}}{\epsilon_{22}} u_1 + Bx_2 + C, \tag{8}$$

where B and C are integration constants that are determined by the electric boundary conditions. C is immaterial and will be omitted. Assuming electrodes either side of the crystal subject it to a sinusoidally varying potential difference with amplitude V imposes the condition $\phi(a) - \phi(\text{-}a) = V$, which gives an expression for B

$$B = \tfrac{1}{2a}\left[V - A\tfrac{e_{212}}{\epsilon_{22}}\left(\mathrm{e}^{ika} - \mathrm{e}^{-ika}\right)\right]. \tag{9}$$

The expression for stress relative to the boundary conditions, T_{21}, is then given by:

$$T_{21} = \left(c_{1212} + \frac{e_{212}^2}{\epsilon_{22}}\right) iku_1(x_2) + \tfrac{e_{212}}{2a}V + \tfrac{e_{212}^2}{\epsilon_{22}2a}\left(u_1(\text{-}a) - u_1(a)\right). \tag{10}$$

B. Motion in fluid

When the crystal is submerged in a fluid, it drags a small portion of the fluid with it as it oscillates. This induces shear waves in the fluid. We will assume that the volume of fluid is such that the penetration depth of the shear wave is significantly less than the fluid depth. A schematic representation of a quartz crystal submerged in a Newtonian liquid is show in figure 3.

Due to the symmetry of the crystal we assume the displacement field in the fluid is parallel to the faces of the crystal. One can simplify the Navier-Stokes equation to arrive at

$$\rho_f \dot{v}_1 - \eta v_{1,22} = 0, \tag{11}$$

where ρ_f is the density of the fluid, v_1 is the component of the velocity parallel to the crystal surface and η is the viscosity of the fluid. The associated stress is

$$T_{21} = \eta v_{1,2}. \tag{12}$$

The fluid in contact with the upper surface requires a periodic solution diminishing with distance to the upper surface of the crystal. v_1^T takes the form

$$v_1^T = D\mathrm{e}^{i(-\kappa x_2 + \omega t)}, \tag{13}$$

where the superscript T indicates the fluid attached to the top surface. The fluid in contact with the lower surface requires a periodic solution diminishing with distance to the lower surface of the crystal. v_1^b takes the form

$$v_1^b = E\mathrm{e}^{i(\kappa x_2 + \omega t)}, \tag{14}$$

where the superscript b indicates the fluid attached to the bottom surface.

When combined with (11), both of these expressions result in the dispersion relation

$$\kappa^2 = \frac{i\rho_f \omega}{\eta}. \tag{15}$$

This dispersion relation contains the imaginary unit, i, indicating that κ and ω cannot simultaneously be real, and the presence of damping.

III. Operation of Crystal in Fluid

It is assumed that when the crystal oscillates while submerged in the fluid, the fluid layer in contact with the crystal is dragged perfectly in sync with the crystal. This assumption imposes the conditions at the boundaries

$$\dot{u}_1(a^-) = v_1^t(a^+), \tag{16a}$$

$$\dot{u}_1(\text{-}a^+) = v_1^b(a^-). \tag{16b}$$

The stress must be continuous at the boundaries, imposing the conditions

$$T_{21}(a^-) = T_{21}(a^+), \tag{17a}$$

and

$$T_{21}(\text{-}a^+) = T_{21}(\text{-}a^-). \tag{17b}$$

A. Frequency

The conditions at the boundary combine together to form the expression

$$\frac{2i\omega\frac{\eta\kappa}{\alpha k}}{1 + \omega^2\left(\frac{\eta\kappa}{\alpha k}\right)^2} = \tan 2ka, \tag{18}$$

which when combined with dispersion relations 7 and 15 reveal a frequency equation. In the absence of fluid, $\eta = 0$ and the result is the natural oscillation frequency for a crystal in free space

$$\omega_0 = \frac{n\pi}{2a}\sqrt{\frac{\alpha}{\rho}} \qquad n = 0, 1, 2, 3, \ldots \tag{19}$$

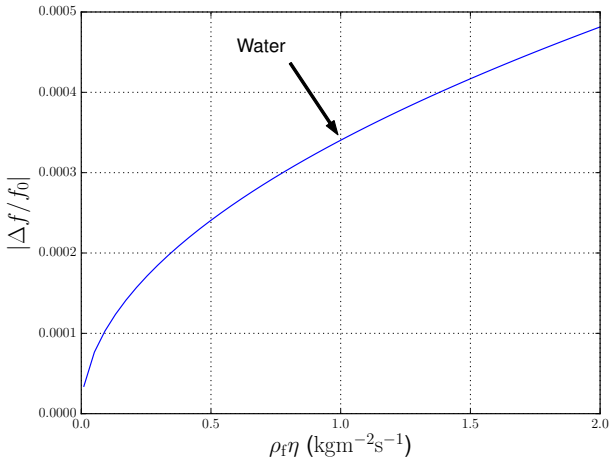

Fig. 4: Relative frequency shift due to viscous damping. The location of water on the curve is indicated.

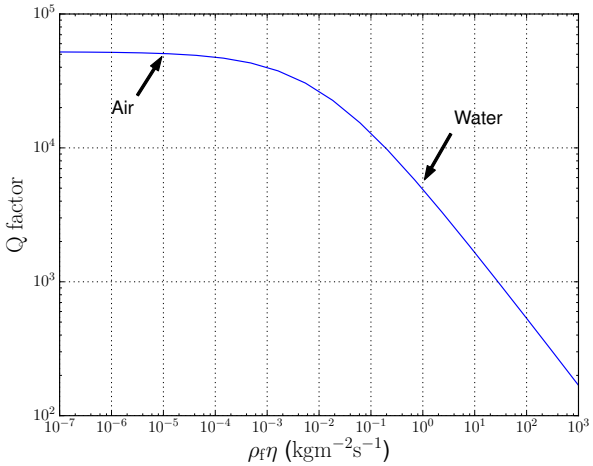

Fig. 5: Quality factor response to viscous damping. The locations on the curve corresponding to air and water are indicated.

where $n = 0, 2, 4, ...$ correspond to the symmetric modes and $n = 1, 2, 3, ...$ correspond to the antisymmetric modes. This expression agrees with the well known results of Koga [25]. It happens that only the antisymmetric modes are physically realizable.

As we only intend to examine weakly viscous fluids, any shift in resonant frequency induced by the presence of a fluid, $\Delta\omega$, is much smaller than the resonant frequency itself, ie $\frac{\Delta\omega}{\omega} \ll 1$. Linearizing both sides of (18) around $\frac{\Delta\omega}{\omega} = 0$ and keeping only first order terms leads to an expression for the frequency shift induced by submerging the crystal in a weakly viscous fluid

$$\Delta\omega = (i-1)\frac{2}{\pi}\sqrt{\frac{\rho_f\eta}{2\rho_q\alpha}}\omega_0^{3/2}\left(1 - \frac{\rho_f\eta}{\rho_q\alpha}\omega_0\right). \tag{20}$$

As expected the shift is complex, meaning the wave is dispersive. Since $\frac{\rho_f\eta}{\rho_q\alpha}\omega_0 \ll 1$, the second term in the brackets can be neglected for most practical purposes. The shift in frequency is proportional to $\sqrt{\rho_f\eta}$, which indicates that frequency measurement alone is not enough to infer the viscosity of the fluid, and more sophisticated techniques are required to pry apart the density-viscosity product.

This expression differs from those commonly stated in the literature for two main reasons. Firstly, it is a complex expression, retaining information about energy losses in the system. This helps us enormously in characterizing the quality of resonance. And secondly, it is an expression for the shift in frequency when the crystal is completely submerged in a fluid. This is reasonably unexplored, in contrast to the case where only one surface is in contact with a liquid which has been extensively studied with established results [9], [26]. Our results show that the frequency shift induced by completely submerging a crystal in a liquid is approximately twice the shift induced when the crystal is only in contact with a liquid on one surface, as our intuition might expect.

The relative frequency shift is plotted on Figure 4 against density viscosity product. The shifts are typically quite small-around 0.034% for water. This amounts to a 3.4kHz shift for a 10MHz crystal.

B. Quality factor

For a general resonant system, the quality factor is related to the complex frequency by [27]

$$2Q = \frac{\mathrm{Re}(\omega)}{\mathrm{Im}(\omega)}. \tag{21}$$

The quality factor against density viscosity product is shown on Figure 5. The trend reflects our intuition in that as the density viscosity product gets smaller the Q-factor begins to resemble the Q-factor of the crystal in free space. In the absence of all losses this value would be infinite, and although internal losses within the crystal are extremely small the electrodes attached to the crystal surface are a source of loss. The exact effect of the electrodes on the Q-factor is strongly dependent on the mounting technique and conditions. However, the free space Q-factor of an electroded quartz crystal is typically of the order of 10^5 [28].
As the density viscosity product of the fluid increases, more of the driving energy put into the crystal goes into moving the fluid rather than being stored in the crystal to sustain oscillations. This supports the sharp drop in Q-factor observed on Figure 5.

C. Oscillation amplitude

The displacement amplitude can be found by the stress boundary conditions at the upper surface, and solving for the amplitude A in terms of V, with ω given by 19 and 20. This results in the expression

$$A = -\frac{e_{212}}{2a}\left[(ik\alpha - \kappa\omega\eta)\,\mathrm{e}^{ika} + \frac{e_{212}^2}{2a\epsilon_{22}}\left(\mathrm{e}^{-ika} - \mathrm{e}^{ika}\right) + \right]^{-1} V. \tag{22}$$

In air, where the density-viscosity product is small (about $10^{-5}\mathrm{kgm}^{-2}\mathrm{s}^{-1}$), this oscillation amplitude amounts to about $5 \times 10^{-2} p\mathrm{m}$. This value is approximately 10^6 times smaller than values stated elsewhere [1], [29], [30]. However, Kanazawa [30] produces an expression agreeing with equation 22. Further work is needed in this area to obtain confidence in equation 22.

IV. Oscillator circuit

Decisions regarding the circuitry responsible for driving the crystal and receiving its response depend largely on the resonant properties of interest. For applications in viscosity sensing, many regard dissipation to be the most useful property [6], [31]. Dissipation characterizes the losses in the system, and is equal to the inverse of the Q-factor, making it proportional to $\sqrt{\rho_\mathrm{f}\eta}$ to a good approximation. The dissipation of an oscillatory system is found by scanning across a range of frequencies, and extracting the Q-factor.

It is theoretically possible to infer fluid properties from the oscillation frequency, through use of equation 20. Frequency measurements are much easier conduct than dissipation measurements, as crystal drivers such as the SN74LVC1GX04 from Texas Instruments act to drive the crystal at its resonant frequency. The circuitry involved is incredibly simple, and inexpensive. However, there are a number of practicalities to keep in mind. Firstly, the ability of the driver to drive oscillations accurately at the resonant frequency diminishes alongside the Q-factor [32]. And secondly, an accurate time base is essential in order to confidently resolve the desired shifts in frequency.

V. Conclusions and future work

In order to use wirelessly excited quartz crystals as a sensing element for fluid viscosity, a detailed understanding of the operation of a QCM completely submerged in a fluid must first be achieved. We have developed reasonable expressions for the shift in frequency due to viscous effects of a QCM submerged in a liquid. These shifts are typically quite small, around 0.03% to 0.04%. The frequency shift was used to characterize the quality of resonance, which suffers as viscous load increases.

Our model for the oscillation amplitude is still in progress, and our immediate attention will be devoted towards its development. Once this has been achieved, these characterizations need to be further extended to predict the feasible limits of QCM operation in liquid.

To infer the density-viscosity product of a fluid from frequency data, the time base of the frequency counter must have a greater stability than the the crystal itself. An appropriate counter must therefore be selected such that the desired shifts in frequency can be measured.

Once it is established that the shifts can be measured, our focus will shift to the inference of model parameters from frequency data.

Acknowledgment

This work was supported under MBIE contract UOOX-1208.

References

[1] A. Janshoff, H.-J. Galla, and C. Steinem, "Piezoelectric mass-sensing devices as biosensors-an alternative to optical biosensors?" *Angew. Chem. Int. Ed*, vol. 39, pp. 4004–4032, 2000.

[2] V. E. Bottom, "Dielectric Constants of Quartz," *Journal of Applied Physics*, vol. 43, no. 4, pp. 1493–1495, 1972.

[3] G. Sauerbrey, "Verwendung von schwingquarzen zur wägung dünner schichten und zur mikrowägung," *Zeitschrift für Physik*, vol. 155, no. 2, pp. 206–222, April 1959.

[4] S. Bruckenstein and M. Shay, "Experimental aspects of use of the quartz crystal microbalance in solution," *Electrochimica Acta*, vol. 30, no. 10, pp. 1295 – 1300, 1985.

[5] H. Muramatsu, E. Tamiya, and I. Karube, "Computation of equivalent circuit parameters of quartz crystals in contact with liquids and study of liquid properties," *Analytical Chemistry*, vol. 60, no. 19, pp. 2142–2146, 1988.

[6] D. Johannsmann, "Viscoelastic, mechanical, and dielectric measurements on complex samples with the quartz crystal microbalance," *Physical Chemistry Chemical Physics*, vol. 10, no. 31, pp. 4516–4534, 2008.

[7] S. Kurosawa, E. Tawara, N. Kamo, and Y. Kobatake, "Oscillating frequency of piezoelectric quartz crystal in solutions," *Analytica Chimica Acta*, vol. 230, no. 0, pp. 41 – 49, 1990.

[8] S. Kurosawa, H. Kitajima, Y. Ogawa, M. Muratsugu, E. Nemoto, and N. Kamo, "Resonant frequency of a piezoelectric quartz crystal in contact with solutions," *Analytica Chimica Acta*, vol. 274, no. 2, pp. 209 – 217, 1993.

[9] K. K. Kanazawa and J. G. G. II, "The oscillation frequency of a quartz resonator in contact with liquid," *Analytica Chimica Acta*, vol. 175, no. 0, pp. 99 – 105, 1985.

[10] A. Stevenson, B. Araya-Kleinsteuber, R. Sethi, H. Mehta, and C. Lowe, "Hypersonic evanescent waves generated with a planar spiral coil," *Analyst*, vol. 128, no. 9, pp. 1175–1180, 2003.

[11] A. Stevenson, B. Araya-Kleinsteuber, R. Sethi, H. Metha, and C. Lowe, "Planar coil excitation of multifrequency shear wave transducers," *Biosensors and Bioelectronics*, vol. 20, no. 7, pp. 1298 – 1304, 2005.

[12] A. Stevenson, A. Roque, B. Araya-Kleinsteuber, E. Kioupritzi, and C. Lowe, "Wireless excitation of quartz crystals immersed in an aqueous fluid," *Analyst*, vol. 131, no. 4, pp. 474–476, 2006.

[13] F. Lucklum and B. Jakoby, "Novel magneticacoustic resonator sensors for remote liquid phase measurement and mass detection," *Sensors and Actuators A: Physical*, vol. 145146, pp. 44 – 51, 2008, special Issue: Transducers/07 Eurosensors XXI, The 14th International Conference on Solid State Sensors, Actuators and Microsystems and the 21st European Conference on Solid-State TransducersThe 14th International Conference on Solid State Sensors, Actuators and Microsystems and the 21st European Conference on Solid-State Transducers.

[14] H. Ogi, "Wireless-electrodeless quartz-crystal-microbalance biosensors for studying interactions among biomolecules: A review," *Proceedings of the Japan Academy. Series B, Physical and biological sciences*, vol. 89, no. 9, p. 401, 2013.

[15] T. Shagawa, H. Torii, F. Kato, H. Ogi, and M. Hirao, "Relationship between viscosity change and specificity in protein binding reaction studied by high-frequency wireless and electrodeless mems biosensor," *Japanese Journal of Applied Physics*, vol. 54, no. 6, p. 068001, 2015.

[16] E. Tonoli, M. Ba, M. Ferrari, and V. Ferrari, "Piezoelectric resonant sensors with contactless interrogation for mass-sensitive and acoustic-load detection," *Procedia Engineering*, vol. 47, pp. 442 – 445, 2012, 26th European Conference on Solid-State Transducers, {EUROSENSOR} 2012.

[17] M. Thompson, S. M. Ballantyne, L.-E. Cheran, A. C. Stevenson, and C. R. Lowe, "Electromagnetic excitation of high frequency acoustic waves and detection in the liquid phase," *Analyst*, vol. 128, no. 8, pp. 1048–1055, 2003.

[18] F. Kato, H. Ogi, T. Yanagida, S. Nishikawa, M. Hirao, and M. Nishiyama, "Resonance acoustic microbalance with naked-embedded quartz (ramne-q) biosensor fabricated by microelectromechanical-system process," *Biosensors and Bioelectronics*, vol. 33, no. 1, pp. 139 – 145, 2012.

[19] F. Kato, H. Ogi, T. Yanagida, S. Nishikawa, M. Nishiyama, and M. Hirao, "High-frequency electrodeless quartz crystal microbalance chip with a bare quartz resonator encapsulated in a silicon microchannel," *Japanese Journal of Applied Physics*, vol. 50, no. 7S, p. 07HD03, 2011.

[20] H. Ogi, K. Motoshisa, T. Matsumoto, K. Hatanaka, and M. Hirao, "Isolated electrodeless high-frequency quartz crystal microbalance for immunosensors," *Analytical chemistry*, vol. 78, no. 19, pp. 6903–6909, 2006.
[21] B. K. Sinha and H. F. Tiersten, "First temperature derivatives of the fundamental elastic constants of quartz," *Journal of Applied Physics*, vol. 50, no. 4, pp. 2732–2739, 1979.
[22] J. A. Kosinski, J. G. Gualtieri, and A. Ballato, "Thermoelastic coefficients of alpha quartz," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 4, pp. 502–507, July 1992.
[23] R. Bechmann, A. D. Ballato, and T. J. Lukaszek, "Higher-order temperature coefficients of the elastic stiffinesses and compliances of alpha-quartz," *Proceedings of the IRE*, vol. 50, no. 8, pp. 1812–1822, Aug 1962.
[24] S. J. Martin, V. E. Granstaff, and G. C. Frye, "Characterization of a quartz crystal microbalance with simultaneous mass and liquid loading," *Analytical Chemistry*, vol. 63, no. 20, pp. 2272–2281, 1991.
[25] I. Koga, "Thickness vibrations of piezoelectric oscillating crystals," *Journal of Applied Physics*, vol. 3, no. 2, pp. 70–80, 1932.
[26] M. Thompson, A. L. Kipling, W. C. Duncan-Hewitt, L. V. Rajaković, and B. A. Čavić-Vlasak, "Thickness-shear-mode acoustic wave sensors in the liquid phase. a review," *Analyst*, vol. 116, no. 9, pp. 881–890, 1991.
[27] L. Li and C.-H. Liang, "Analysis of resonance and quality factor of antenna and scattering systems using complex frequency method combined with model-based parameter estimation," *Progress In Electromagnetics Research*, vol. 46, pp. 165–188, 2004.
[28] B. Borovsky, B. Mason, and J. Krim, "Scanning tunneling microscope measurements of the amplitude of vibration of a quartz crystal oscillator," *Journal of Applied Physics*, vol. 88, no. 7, pp. 4017–4021, 2000.
[29] B. A. Martin and H. E. Hager, "Velocity profile on quartz crystals oscillating in liquids," *Journal of Applied Physics*, vol. 65, no. 7, pp. 2630–2635, 1989.
[30] K. áKeiji Kanazawa, "Mechanical behaviour of films on the quartz microbalance," *Faraday Discussions*, vol. 107, pp. 77–90, 1997.
[31] C. Mista, M. Zalazar, A. Pealva, M. Martina, and J. M. Reta, "Open source quartz crystal microbalance with dissipation monitoring," *Journal of Physics: Conference Series*, vol. 705, no. 1, p. 012008, 2016.
[32] C. Barnes, "Development of quartz crystal oscillators for under-liquid sensing," *Sensors and Actuators A: Physical*, vol. 29, no. 1, pp. 59–69, 1991.

Wireless Power and Network Synchronisation for Agricultural and Industrial Remote Sensors using Low Voltage CMOS Harvesting and Data Demodulator IC

[1]T. Lee, [1]H. R. B. Kennedy, [1]R. A. Bodnar, [2]J. Scott, [1,2]W. Redman-White
[1]University of Southampton, UK [2]University of Waikato, Hamilton, NZ
tl7e13, hk1g10, r.bodnar@soton.ac.uk; wrw@ecs.soton.ac.uk; scottj@waikato.ac.nz

***Abstract*—This paper presents a wide area medium frequency loosely coupled magnetic energy harvesting system with power delivery and network synchronisation for remote sensors, intended for agricultural and industrial environments. Intended for situations with poor service access, power is supplied from a source via a large area loop. Receiver nodes may use ferrite cored coils for good efficiency with modest volume. Transmission of low bandwidth network synchronisation data permits very low operational duty cycle with the need for real time clocks or wake up receivers and their associated power drain. As a key enabler for the system, a full custom energy harvester and QPSK data demodulator IC has been designed and fabricated in a commercial 180nm CMOS technology. The IC occupies 0.54mm2 and can deliver 10.3µW at 3V to an external battery or capacitor. With standard MOS device thresholds the rectifier can start from cold with only 250mV peak from the antenna loop, and the battery charge output is delivered with 330mV peak input. Results are presented from laboratory evaluation and from preliminary measurements in the field with a 10m x 10m loop driven at 800kHz.**

I. INTRODUCTION

Remote sensor networks are becoming well established with the availability of very low power IC technology. For many applications, a simple battery remains the most sensible choice of power and service lifetimes of several years are common. Where access for maintenance is restricted, energy harvesting or scavenging systems have become popular allowing indefinite operation. Photovoltaic cells are a common and very successful source provided that adequate illumination is available [1]. In some industrial environments, periodic vibration from machinery can be used successfully [2].

For agricultural and industrial applications power can be obtained by a variety of methods depending on the specific circumstances. Where photovoltaic generation is not practical, wireless power becomes attractive. Rather than attempt to recover adequate power from an ill-defined ambient electromagnetic environment at some high frequency, we investigate supplying power in the near field via very weak magnetic coupling from a medium frequency source using a large-area loop. In this frequency range a ferrite-cored loop receiver antenna may be used to increase effective area and ensure acceptably small volume. The use of a lower-frequency source also reduces problems associated with localised signal attenuation due to walls etc., making deployment in less well-defined environments easier. The current in the transmitting loop is not considered a significant issue for the applications envisaged provided that current levels are not impractical and that the power required for the source does not exceed the limits of a single-phase AC mains outlet. A further benefit of supplying the power from a known source is that management and control data may be transmitted without a significant bandwidth penalty. Supplying timing information to a number of deployed sensor nodes provides the opportunity to operate at very low duty cycle ratios without the need for a real-time clock or a wake-up receiver, eliminating the disproportionate power of such functions. To achieve workable operating range in a physically small unit, the receiver power recovery circuits must be able to cold start with the lowest possible voltage and hence the weakest magnetic field in a given operation area. Rather than focus on power transfer efficiency, we instead direct our effort to ensuring that the receiver can start from an unpowered ("cold start") condition with the minimum voltage at the antenna coil. The main effort of this work is hence directed at demonstrating the feasibility of an IC that performs the functions of rectifier, data demodulator and power management.

We first describe the overall architecture for the proposed system, and then the requirements for the receiver IC. We then present some aspects of the circuit design and show measured results from laboratory and field experiments.

II. SYSTEM ARCHITECTURE

A. Wide-Area MF Loosely-Coupled Magnetic Power Transfer

Much research has been directed into harvesting or scavenging power from electromagnetic sources. In the main, these works have looked at extracting energy from the ambient conditions in the UHF and low microwave regions from cellular and wireless LAN transmissions [3]. At these frequencies, a physically small antenna can easily be of the order of a wavelength and thus achieve good efficiency. Whilst

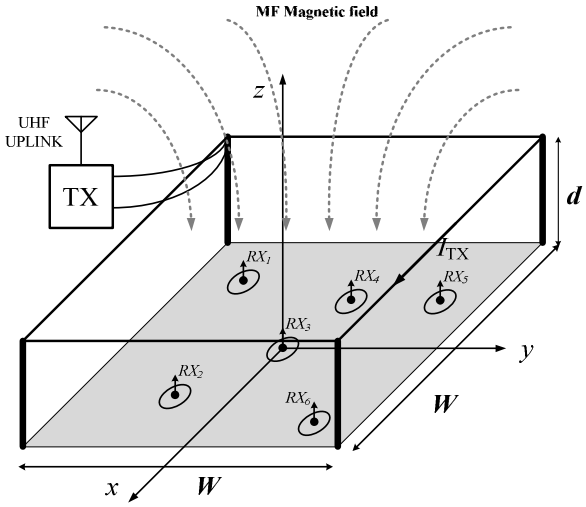

Figure 1. Wide area MF magnetic loosely coupled energy harvesting.

potentially attractive from the user's perspective, there are some drawbacks. If the source is an RF service intended for another purpose, the field strength at the intended location may vary considerably. Propagation is also restricted to essentially line of sight or with only insubstantial obstructions for the sensor antenna locations.

In this work, we seek to address the needs of applications where sensors may need to be deployed within a defined area but may be sited where direct illumination is not available, or behind or below some barrier substantially opaque at UHF and microwave frequencies. Industrial and agricultural sites may pose such challenges, but the site may well be suited to providing a dedicated RF source. Fig. 1 shows an example of a wide area MF magnetic, loosely-coupled, energy-harvesting system. A dedicated magnetic field is created by a large loop antenna that may be fixed at or near the boundary of the operational site. The power for the MF current source can be from AC mains or other large fixed supply, and is consistent with typical small industrial equipment.

The sensors, represented by several secondary coils, harvest the available energy within the transmitter loop. Air cored receive coils may be used, but the use of ferrite cores can reduce the physical size of a practical receiver to a few cm^3. The requirement for the magnetic system is to provide enough voltage at the receiver coil terminals to activate the harvesting electronics. For this system, 250mV was set as the target antenna voltage for a cold start.

Figure 2 shows a plot of the field created by an example configuration of a 10m x 10m loop (as used in practical measurements). It can be seen that the magnetic field is relatively constant within the loop, and decays rapidly outside its perimeter [4], making interference with other spectrum users unlikely.

Note that while the regulatory framework in many territories does not currently permit such proposed systems, some territories are permitting experimental work, and it is likely that the declining use of MF broadcasting will make narrow band non-radiating systems more likely elsewhere.

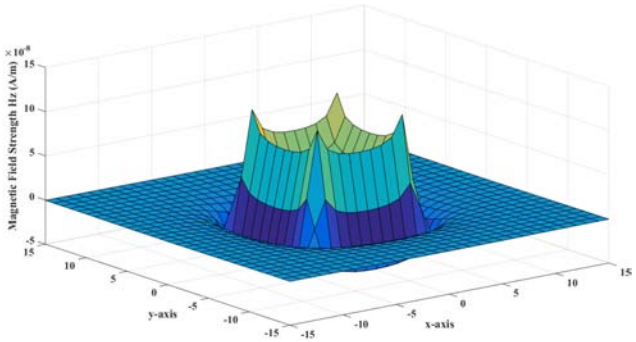

Figure 2. Calculated magnetic field pattern for a 10mx10m square transmit antenna loop.

B. Synchronisation Protocol

Almost all remote sensors are configured to operate with a low active duty cycle, as this enables very significant average power savings when data collection is only needed at a very slow sample rate. The calculation of the average power consumption of the smart sensor can be done by using equation (1). P is the average power consumption; P_n is the power consumption of the module; t_n is the active time of the module per period and t_T is the total time period.

$$P = \frac{P_1 t_1 + P_2 t_2 + P_3 t_3 + ...}{t_T} \quad (1)$$

For example; a smart sensor collects data once every hour. The sensor module requires a receiver, a sensing transducer and an uplink transmitter (typically UHF) where the power consumption and the active time of the modules are shown in table 1. By using equation (1), the average power consumption of the smart sensor is 0.24µW.

Problems arise when the duty cycle is made very low in terms of guaranteeing that a data burst sent from a node is received reliably by the controller. If the timing is not accurate, the controller itself may be in a sleep mode, or there may be a collision between data bursts from two or more remote nodes. Conventional means to control synchronization employ high accuracy real-time clock circuits, or wake-up receivers that monitor a similar frequency to the uplink. In both cases, some significant power drain is associated with the circuits that are always on, making further demands on the energy harvesting and limiting the available energy for the data acquisition.

In the proposed system, we use the MF power transfer system to carry a low-speed data channel, conveying basic system management and timing information. In this way any receiver

Table 1. Representative relative power consumption of typical components in a remote sensor.

Module	Current supply	Voltage supply	Power consumption	Active time (per hour)	Energy consumption (per hour)
Sensor node	350µA	3V	1.05mW	130.2ms	136.71µJ
Transmitter	10.5mA	3V	31.5mW	4ms	126µJ
Microcontroller	365µA (Active) 11nA (Idle)	3V	1.095mW (Active) 99nW (Idle)	200ms	219µJ (Active) 356.4µJ (Idle)

nodes within the loop can maintain precise synchronisation and avoid data collisions with the small power overhead needed for a low-rate demodulator function.

Figure 3 illustrates the flow chart of a possible synchronisation protocol for the system. Initially, the sensors harvest the magnetic energy from the primary source. When the receiver is ready to operate, an acknowledge command is transmitted to the source and the initialisation of the system is complete. In normal operating mode, the sensors are only active when they receive the *wake-up* command from the transmitter. The receiver can consume less power as it operates with a very low duty cycle without timing problems.

III. RECEIVER ARCHITECTURE

Figure 4 shows the block diagram of the complete harvester/receiver IC, as implemented in CMOS technology. Critical to the operation of the whole IC is the active AC-DC rectifier block, as its performance sets the lower operating voltage limit, and hence the range or size of transmit loop.

As the receiver loop coil is relatively small compared with the size of the transmitter loop, hence a ferrite material is used to increase the permeability of the loop coil [5]. The induced voltage received from the ferrite rod can be magnified by employing a parallel capacitor to form a high Q factor resonant circuit.

A rectifier circuit is used to convert the incoming AC voltage to DC supply. The received voltage in the intended application might be very small since the coupling coefficient is very poor. Thus, a very low start-up voltage of the rectifier is more important than the power conversion efficiency (PCE), and hence impedance matching the antenna to the rectifier in its cold state is not desirable. Implementing a rectifier in standard CMOS presents problems since the transistors will not turn on significantly until the threshold voltage is reached, typically around 0.4V – 0.5V for common processes. This threshold may be avoided by using a costly non-standard IC process with devices such as Schottky diodes [6] or near zero threshold voltage transistors [7]. To avoid the associated cost and difficulty obtaining such a special process, we use internal circuitry to cancel most of the MOS threshold apparent in the rectifier, as shown in figure 5. Using such techniques, simulations show, the circuit can start from cold with an input much less than the MOS threshold, in this case at about 250mV peak.

Once the rectifier starts, its inherent doubling function gives an output V_{RECT} at nearly twice the input. When this reaches around 600mV the power management functions are also able to operate, and a DC-DC charge pump raises the recovered power to a level suitable for slowly charging an external lithium-ion battery or other load. In a complete sensor-node application, this would typically include a low power microcontroller and a VHF/UHF transmitter as well as the sensor transducer itself.

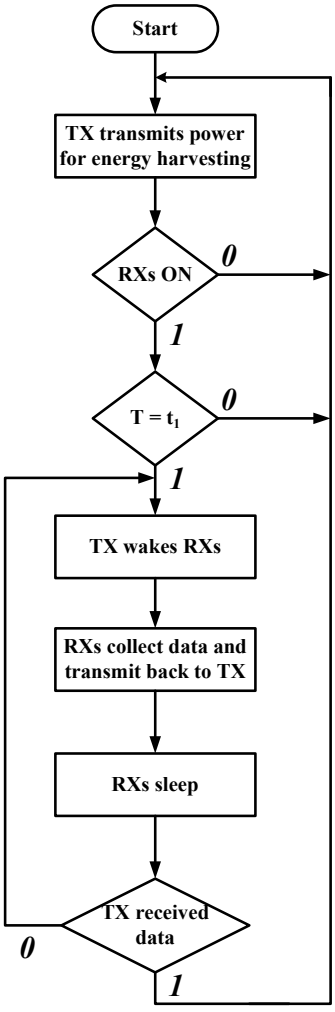

Figure 3. Example of synchronization protocol for network.

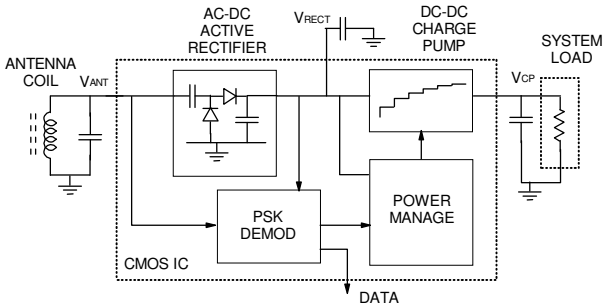

Figure 4. Block diagram of the MF harvester/receiver IC.

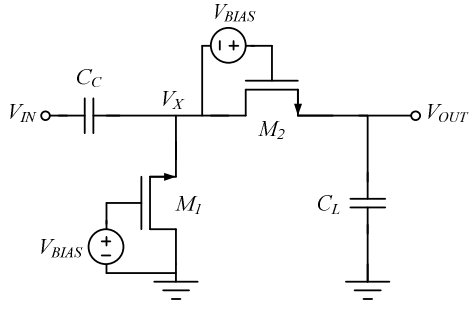

Figure 5. Active rectifier concept to reduce impact of MOS threshold voltage.

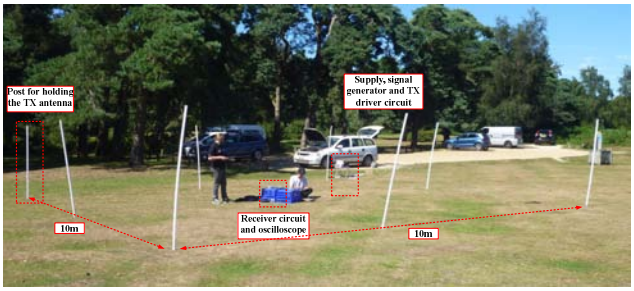

Figure. 6. Measurement setup of the wide area MF magnetic loosely coupled energy harvesting.

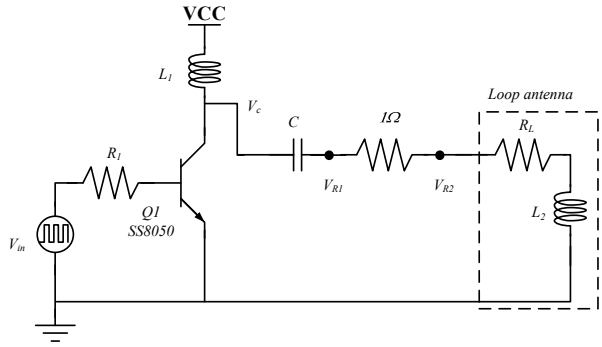

Figure. 7. Transmitter loop driver circuit.

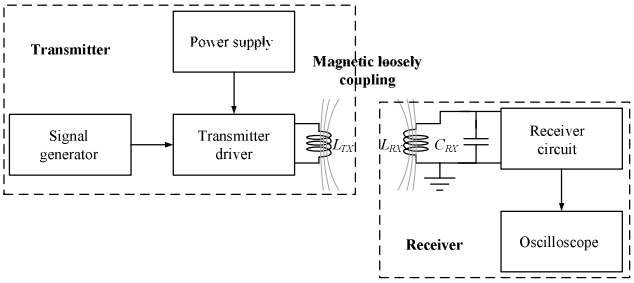

Figure. 8. Testing setup for the magnetic loosely coupled energy harvesting measurement.

IV. Measurement Results

In addition to conventional lab measurements the IC was tested with a large area loop in an electrically quiet external environment. A square loop antenna with 10x10m^2 was constructed to provide a dedicated magnetic field for the energy harvesting system operating at ~800kHz as can be seen in figure 6. The loopwas implemented using copper wire with a resistance of 0.57Ω/m is attached to the top of non-magnetic poles 2m above ground level (as might be constructed using fencing structures). The calculated inductance of the transmitting antenna is 90µH while the resonance capacitance is 413pF. Figure 7 illustrates the transmitter driver circuit implemented using the bipolar junction transistor SS8050 for the switching active device. The block diagram of the testing setup is shown in figure 8. The driver circuit is supplied from a bench power supply while the input switching signal is generated from a function generator.The receiver antenna is implemented using a 4cm x 1cm diameter ferrite rod antenna with the resonant capacitor. The loss resistance of the coil is 0.83Ω while the loop coil inductance and resonant capacitance

Table 2 Measured Harvested Signals in the Receiver IC

Position (m)		Received voltage (mV_{peak})	V_{RECT} (mV)	V_{CP} (V)	DC output power (uW)
x	*Y*				
4	-4	278	914	2.64	0.697
0	-4	246	860	-	-
0	0	209	805	-	-
2.5	2.5	227	900	-	-
-2.5	4	241	754	-	-
-2.5	-2.5	231	910	-	-
-4	-4	264	915	0.915	0.084
-4.5	-4.5	337	915	3.2	1.024

Table 3 Measured Harvested Energy of the Receiver IC at Centre of the Transmitting Coil with Different Transmitter Coil Currents

Transmitter current (mA_{RMS})	Received voltage (mV_{peak})	V_{RECT} (mV)	V_{CP} (V)	DC output power (uW)
220	209	805	-	-
230	239	0.77	0.343	0.012
570	280	928	2.709	0.734

are 53µH and 700pF respectively. The measured in circuit Q-factor of the receiving antenna was 33. The receiver was tested at different positions within the square loop antenna where the vertical distance between the coils is approximately 0m.

Table 2 shows the measured harvested power of the receiver IC in a grid where the origin (*x*=0m, *y*=0m) is at the centre of the square loop.The measured transmitting current generated from the driver circuit is at 220 mA_{RMS}. The receiver IC achieves 1.024µW maximum output power at 3.2V when the receiving coil is located close to the corner of the square loop antenna. Table 3 shows the measured harvested energy of the receiver IC at the center of the transmitting coil (*x*= 0m, *y*= 0m) when the transmitter current is varied. The results indicate that the received energy can be increased to usable levels and all the subsystems made functional by increasing the output current of the driver circuit to around 700mA_{RMS}.

V. Conclusion

A wide area magnetic loosely coupled energy harvesting system has been presented, incorporating a dedicated localised wireless power source. An embedded data channel for management and synchronization open the way for operation at very low active duty cycles with attendant average power savings. The key component for the system, the low-power harvester/receiver, has been implemented in a full custom CMOS integrated circuit. The IC has been shown to start from cold with as little as 250mV at the antenna coil terminals and with slightly larger inputs all of the internal functions are

active. The measurement of the harvesting energy using the wireless energy harvesting IC shows the maximum output power of 1.024μW at 3.2V. Measurements in an exterior setting confirm that the system can operate with practical transmit and receive power levels.

References

[1] Y. He, X. Cheng, W. Peng, G. L. Stuber, "A survey of energy harvesting communications: model and offline optimal policies," in *IEEE Communications Magazine*, vol. 53, no. 6, pp. 79-85, June 2015.

[2] Qian Zhang and Eun Sok Kim, "Vibration Energy Harvesting Based on Magnet and Coil Arrays for Watt-Level Handheld Power Source", Proc IEEE, vol. 102, no. 11, pp1747-1761, 2014.

[3] Manuel Piñuela, Paul D. Mitcheson, and Stepan Lucyszyn, "Ambient RF Energy Harvesting in Urban and Semi-Urban Environments", *IEEE Transactions on Microwave Theory and Techniques*, vol. 61, no. 7, pp2715-2726, 2013.

[4] M. Sawan, Yamu Hu and J. Coulombe, "Wireless smart implants dedicated to multichannel monitoring and microstimulation," in *IEEE Circuits and Systems Magazine*, vol. 5, no. 1, pp. 21-39, 2005.

[5] C. R. Valenta and G. D. Durgin, "Harvesting Wireless Power: Survey of Energy-Harvester Conversion Efficiency in Far-Field, Wireless Power Transfer Systems," in *IEEE Microwave Magazine*, vol. 15, no. 4, pp. 108-120, June 2014.

[6] M. Zargham and P. G. Gulak, "Maximum Achievable Efficiency in Near-Field Coupled Power-Transfer Systems," in *IEEE Transactions on Biomedical Circuits and Systems*, vol. 6, no. 3, pp. 228-245, June 2012.

[7] B. M. Badr, R. Somogyi-Gsizmazia, N. Dechev and K. R. Delaney, "Power transfer via magnetic resonant coupling for implantable mice telemetry device," *Wireless Power Transfer Conference (WPTC), 2014 IEEE*, Jeju, 2014, pp. 259-264.

[8] Y. Zhou, B. Froppier and T. Razban, "Schottky diode rectifier for power harvesting application," *RFID-Technologies and Applications (RFID-TA), 2012 IEEE International Conference on*, Nice, 2012, pp. 429-432.

[9] A. Shrivastava, N. E. Roberts, O. U. Khan, D. D. Wentzloff and B. H. Calhoun, "A 10 mV-Input Boost Converter With Inductor Peak Current Control and Zero Detection for Thermoelectric and Solar Energy Harvesting With 220 mV Cold-Start and 14.5 dBm, 915 MHz RF Kick-Start," in *IEEE Journal of Solid-State Circuits*, vol. 50, no. 8, pp. 1820-1832, Aug. 2015.

[10] S. A. S. Mohamed and Y. Manoli, "A novel fully integrated low-power CMOS BPSK demodulator for medical implantable receivers," *2014 IEEE International Symposium on Circuits and Systems (ISCAS)*, Melbourne VIC, 2014, pp. 1098-1101.

Improving the Accuracy of Musical Robots with Programmable Logic Circuits

J. Long, J.W. Murphy
New Zealand School of Music
School of Engineering and Computer Science
Victoria University
Wellington, New Zealand
jason.long@ecs.victoria.ac.nz
jim.w.murphy@gmail.com

D.A. Carnegie, A. Kapur
School of Engineering and Computer Science
Victoria University
Wellington, New Zealand
ajay@karmetik.com
dale.carnegie@vuw.ac.nz

***Abstract*—In order to provide expressive performances, musical robots must be capable of varying their playing parameters in fine-grained and consistent ways. Unfortunately, due to their often high actuator count and signal processing requirements, conventional purely microcontroller-based methods of controlling their actuators have proven insufficient at providing the accuracy required for precise and reproducible control. This paper provides an outline of the shortcomings of traditional methods of musical robot control with an example robot which outputs pulse widths with up to 100 microseconds of error, and describes a newly created control system which makes use of a hybrid microcontroller-FPGA design to achieve error rates in the order of single-digit nanoseconds. The system's components and performance is then described, and comparisons are drawn between the new system and traditional methods. This novel system leverages both the sequential logic capabilities of a microcontroller with the massively parallel capabilities of an FPGA to provide a fine-grained, reliable, low latency and high resolution control for musical robots.**

***Keywords*—Musical Robotics, Solenoid, Servo, Stepper, Music, FPGA, Microcontroller**

I. Introduction

The field of musical robotics is concerned with the activation of real-world acoustic sound objects by way of computer control and electronic actuators. This often takes the form of robotic musical instruments such as mechatronic guitars[1], self-playing drum kits[2], and automatic wind instruments[3], but also encompasses areas such as kinetic sculpture and sound installations[4], and anthropomorphic robots playing musical instruments[5].

There are many reasons that composers and other artists make use of musical robots. Some prefer the acoustic and spacial characteristics of real-world sound objects over loudspeakers[6], while others may enjoy the visual cause and effect relationship that robots can create in a musical performance[7]. Musical robots offer the ability to instantly audition musical material in an acoustic environment during the process of composition and are especially capable of playing very fast, complex music and performances of long duration. They also offer the ability to realise algorithmic, generative or interactive music with real instruments.

However, there remain aspects in which human performers excel over musical robots, and much research is being done to make improvements in these areas. Experienced musicians are very flexible, being equipped with a range of standard and extended playing techniques with which to play their instrument. While it is straight forward to simply instruct a performer to play an instrument in a different or novel way, for musical robots, any extra technique will require extra design, construction and programming to add. These techniques allow human musicians to deliver varied, expressive performances with a very large dynamic range and fine grained and precise control over timing.

Creating expressive performances with musical robots generally involves dynamically varying the amount of power applied to actuators, when the power is applied, and the length of time the power is applied for. The more accurately these variations can be programmed, the more precisely composers can write convincing and expressive performances for these instruments. When these parameters are less accurate, the result may be perceived as a variation in the robot's performance that makes the machine sound 'more natural' or 'more human', or conversely, when the parameters are unfavourably inaccurate, it can create negative issues with timing and dynamics. It is for this reason that it is necessary to improve the accuracy of these signals, so that composers have the choice of whether or not to inject expressive characteristics into the performance of their pieces, and in which way.

This paper begins by providing some background of the area in section 2, introducing the various methods of actuating musical robots and the electronic methods that have been used to drive them. A newly created robotic instrument that is an attempt to improve the accuracy of musical robots is then introduced in section 2A, and the shortcomings of the traditional methods of actuator control that this instrument revealed are outlined. The reasons for these shortcomings and potential solutions are discussed, before the newly developed hybrid microcontroller-FPGA musical robot control hardware is presented in section 3. The various sections of the hardware and its features are described, and its performance is evaluated and compared with traditional musical robot control methods in section 4.

II. Background

While the breakthrough of musical robotics into the realm of mainstream popular music could be considered a 21st

century phenomenon, the field dates back to the 1970s, with Trimpin, Godfried Willem Raes, and Ken Caulkins' early work. However, much of this work was itself based on the rich history of musical automata which dates back to the 9th century. [8] and [9] provide more information on both the recent history of musical robots and the earlier history of musical automata.

While there are many different varieties of musical robot, this paper focuses on the control of the most common group, percussion. There are also a number of methods of actuating percussion robots, from pneumatic devices, to servos and electromagnets, each with their own strengths and weaknesses, outlined in [10]. That study showed that while pneumatics offer great power and servo-based methods offer a lot of flexibility, solenoid based methods are a capable all-round solution, which explains why many of the most prolific musical roboticists such as Trimpin and Godfried Willem Raes utilise them in their projects.

Regarding solenoids, there are also several methods used to control them electronically. Earlier methods often utilised Darlington transistors to switch current to and from the electromagnets[12], but this has since been superseded by the more efficient MOSFET-based circuit illustrated in [11]. The most common method of generating the required pulses that trigger those MOSFETs is utilising the GPIO pins of microcontrollers, however dedicated 16-bit timer ICs such as the Intel 8254[13], or other methods such as analogue control[14] may also be used.

A. *The Problem*

In late 2015 to early 2016, a project was started to improve the accuracy, reliability and usability of musical robots by embedding hardware DSP inside the robots, providing them with musical information retrieval functionality. The first example of a musical robot with this functionality is the Closed Loop Robotic Glockenspiel, described in [15] and shown in Figure 1. The instrument consists of a 21-key glockenspiel with an array of linear tubular solenoids and coil pickups mounted beneath the keys. It utilises two Arm Cortex-M3-based microcontroller boards which process input from external MIDI devices, and control the solenoids. In addition, the conditioned analogue audio signals from the coil pickups are received by the microcontrollers, and they conduct onset and latency detection, continuously and automatically calibrate the instrument in real-time, and perform latency compensation.

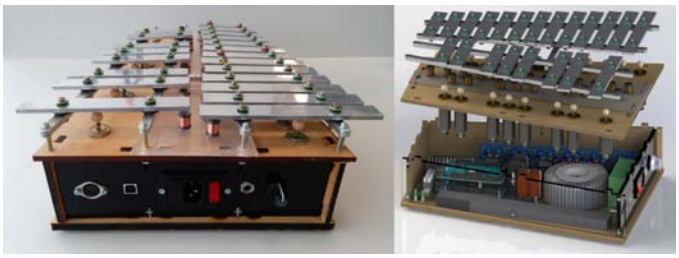

Fig. 1. The Closed-Loop Robotic Glockenspiel.

While evaluating the performance of the auto-calibration algorithms, it was observed that instead of settling on the correct values, the calibrated velocity values were slowly oscillating up and down, decreasing the effectiveness of the calibration. Closer inspection revealed that the cause of the issue could be traced back to variations in the widths of the control pulses output by the microcontrollers. This was caused by the fact that the microcontrollers were undertaking multiple tasks related to input processing, audio-rate analysis, output checking and other tasks, resulting in pulses that were intended to be of identical width varying in actual width by up to 100 microseconds.

The dynamic range of robotic percussion mechanisms is restricted on the low end by the softest hit that will reliably strike the key, and at the high end by the longest pulse that finishes before the solenoid strikes the key, so as to let the solenoid bounce off the key naturally. The Closed-Loop Robotic Glockenspiel achieves a low mechanical latency by mounting the striker solenoids very close to the glockenspiel keys, in turn limiting the dynamic range at the high end. To improve upon this, the voltage applied to the solenoids is increased to approximately 60 V, and the length of control pulses decreased correspondingly. As a result, the Closed-Loop Robotic Glockenspiel's entire dynamic range is controlled by a narrow range of pulses between approximately 1500 and 3000 microseconds in width.

In addition, the velocity response of solenoid actuators is highly non linear, as discussed in detail in [12] and [16]. As a result, there are certain regions of the dynamic range which have increased significance, being particularly sensitive to minute variations in the width of control pulses. Unreliable pulse widths at the low end of the dynamic range are also particularly problematic as they can result in a potentially very musically important programmed note not sounding at all. For these reasons, the accuracy, resolution, consistency and reliability of control pulses is vital to creating musical robots with precisely controllable expressive ability.

There are several potential approaches that could yield improved accuracy in this area. Utilising the timer peripherals of microcontrollers, configuring them to act as one-shot timers and routing their outputs directly to the GPIO pins is one approach that provides greatly increased accuracy. However, many musical robots, especially those with a large number of actuators or musical keys like the Closed-Loop Robotic Glockenspiel, would require individual one-shot timers for every single actuator under control. A survey of available microcontrollers on the market did not produce an offering that provides the high number of high resolution timer peripherals that a controller of such an instrument would require.

Alternatively, multiple microcontrollers could be included, each which would control a few actuators. While possible, this would increase cost, physical space required, and require individually programming each microcontroller separately for each change in firmware - a time consuming and error-prone work flow. Rather than individual microcontrollers, another possibility is utilising multiple programmable interval timer ICs such as the Intel 8254 mentioned by Raes[13]. However these chips also consume a lot of physical space, are costly and inflexible in functionality.

Upon investigation into these possibilities it was decided that the optimal solution to achieve sub-microsecond accuracy of pulse-length for large numbers of channels simultaneously,

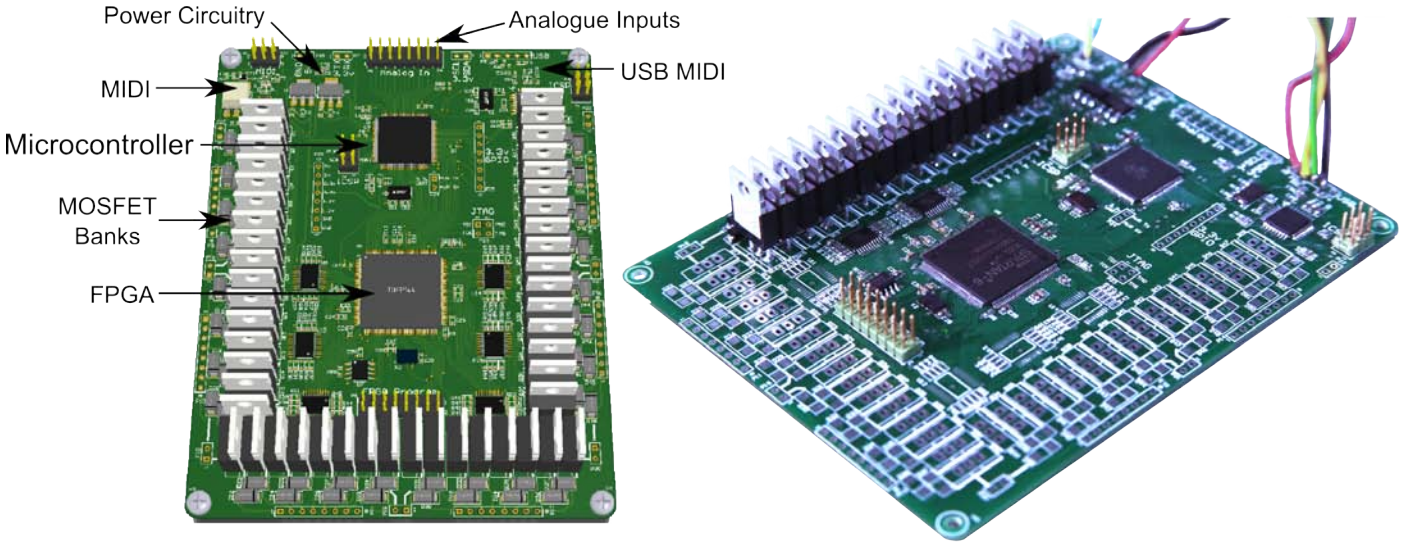

Fig. 2. The new musical robot controller board. Fully populated rendering with main sections labeled (left), and photograph with two banks populated (right).

independently of the work load of the rest of the system, would be to utilize a field-programmable gate array. This FPGA would then be configured to house as many high resolution hardware timers and other circuitry as required. The result would be a single small board that can control large numbers of a variety of actuators with a high level of precision. The following section describes the hybrid microcontroller-FPGA musical robot control board developed to achieve these objectives.

III. The Hybrid Microcontroller-FPGA Board

This section outlines the design concepts of the robot control circuit board shown in Figure 2, complete with descriptions of the board's internal and external communication protocols and capabilities, how the board is programmed and powered, and how it can be used in a large variety of different musical robots.

A. Design Concepts

There were several design goals for the robot controller board. First and foremost, it should provide at least 32 channels of sub-microsecond accurate one-shot control pulses to an equal number of MOSFETs and provide outputs to that many actuators. Secondly, the board should be flexible as to which actuators, MOSFETs and external hardware can be connected to it, and thirdly it should be straightforward to program, compatible with streamlined development environments such as Arduino and Mojo. The following subsections illustrate how these design goals were met, and explain the architecture and capabilities of the circuit board.

The 130 x 100 mm board attempts to support as many use-cases as possible by making use of a modular design. This means that several different communication interfaces such as standard MIDI, USB MIDI, serial communication, SPI, I^2C and GPIO are included, and a channel count of 48 is offered, but any functionality that is not required may not be populated on the board. An example population is shown on the right side of Figure 2 where 2 banks of 8 actuators are required, and both USB and standard MIDI communication is utilised. This leaves the remaining 4 banks of buffers, MOSFETs and flyback diodes unpopulated.

B. Communication with the Outside World

As shown on the left side of Figure 2, traditional MIDI In, Out and Thru connectivity is provided for, adhering to the 2014 update to the MIDI 1.0 specification. MIDI In is primarily utilised for receiving musical control messages from a sequencer or live performance interface, MIDI Out can relay information about the robot and performance back to the sequencer, and the buffered MIDI Thru functionality can relay incoming commands to additional boards should more than 48 channels be required. USB MIDI functionality is also provided on the top right side of the board by an Atmel Atmega8u2 microcontroller equipped with built in USB connectivity, loaded with a modified version of the Hiduino firmware[17].

16 general purpose analogue inputs line the top center of the board, which may be used for reading sensors utilised in closed-loop control systems or other purposes. Additional bi-directional UART, SPI and I^2C ports are broken out for external communication to and from the main microcontroller. An 8-pin GPIO port is also made accessible to accommodate potential future requirements.

C. Internal Communication

In response to external information received via the MIDI, GPIO, analogue inputs and other sources, the main microcontroller processes all of the control logic of the musical robot, including arranging the control of the actuators, processing any closed-loop input signals and conducting any auto-calibration, latency compensation and other functions. The Atmel Atmega1280 microcontroller was chosen due to its 4 UARTS, high number of GPIO ports, low cost and direct compatibility with Atmel Studio, AVR-GCC and the multi-platform and uncomplicated Arduino IDE.

Fig. 3. Diagram of the robot controller circuit board's signal flow.

It is widely documented that latency is an important characteristic in the design and performance of digital musical instruments, and minimizing latency is the subject of much research and development in the area[18]. In order to minimize the latency of the system, the main microcontroller communicates with the FPGA in a parallel configuration. The communication system, as shown in Figure 3, makes use of three 8-bit GPIO ports and an extra single pin to clock the data into the FPGA. This allows the microcontroller to transmit high resolution pulsewidth information along with corresponding actuator numbers in as little as 8 clock cycles, or 1000 ns at 8 MHz. Additionally, a bidirectional UART connection between the microcontroller and the FPGA may also be utilized for less time-critical information transfers.

D. The FPGA

In order to create high resolution, precisely timed signals for a large number of actuators simultaneously, a Xilinx Spartan 6 FPGA is used. The 6SLX9 device was selected due to its appropriate number of I/O ports, number of logic cells that exceeds the requirements of this project, and compatibility with both the Xilinx ISE tools and the streamlined Mojo IDE, with its simplified Lucid hardware description language.

The FPGA is highly flexible in that it can be reconfigured to provide control signals to drive a number of different musical robot actuators from high resolution PWM signals to drive servomotors, to DC and stepper motors control signals. It is also possible to create complex signals such as dynamically changing PWM curves and pulse-wave audio tones for electromagnetic actuation techniques[20]. However, since solenoids are the most frequently used musical robot actuator, a configuration for solenoid control will be described henceforth.

While PWM is occasionally use in conjunction with solenoids, they will more commonly be controlled by either one-shot or toggle style pulses. The castanet and egg-shaker robots shown in [19] are examples of one-shot triggered and toggle controlled instruments respectively. A custom configuration was designed in Verilog which contains an address decoder and an array of 48 one-shot timers, clocked by the FPGA's undivided 50MHz master clock, allowing the device to deliver pulses with a granularity of 20ns. Upon receiving a positive edge on the Latch line, the output line is brought high, the 8 bit address bus is decoded, and the 16 bit pulse-width bus is loaded into the corresponding counter comparison register. In this case the pulse-width values are specified in microseconds, so the counter value is multiplied by a factor of 50 to match the 50MHz master clock. The master clock then decrements the counter each cycle until it reaches zero, at which time the output line is brought down to ground.

As shown in Figure 3, the output banks of the FPGA are sent to up to 6 octal buffer chips which adhere to HCMOS logic standards, allowing 3.3 V logic to reliably register as 'high' in a 5 V powered buffer. These buffers serve the dual purposes of protecting the internal circuitry from potentially damaging voltage spikes from the outside, and lifting the voltage level from 3.3 V to 5 V in order to activate a larger variety of MOSFETs, analogue servomotors and other devices that require the higher voltage levels.

E. The Programmer

The musical robot control board is equipped with JTAG pins in order to uploaded configurations to and debug the on-board FPGA, but in order to maintain compatibility with the Mojo IDE and with it the Lucid hardware description language, and to provide a low-cost and shallow learning curve option, a separate programmer was also created. The small 35 x 50 mm board shown in Figure 4 is built around an Atmel Atmega32u4 microcontroller with included USB connectivity and is based on part of the Mojo V3 circuit board schematics.

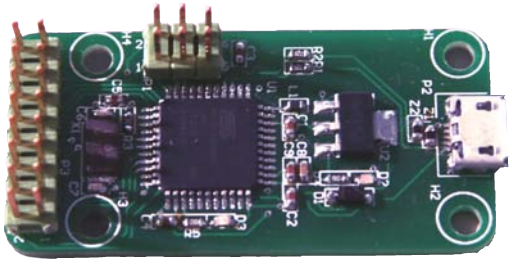

Fig. 4. The detachable USB programmer for the FPGA.

By connecting a PC via a standard micro-USB cable, the board is enumerated and it is possible to select it as a programmer in the Mojo IDE or standalone Mojo Loader software. By connecting the dual-row header shown on the left side of Figure 4 to the main solenoid control board via a ribbon cable, a bit stream can be established which loads a custom configuration to the FPGA. Alternatively, the programmer was designed in such a way so as to allow a user to solder the programmer board directly to the main control board in case of permanent installation. The programmer then sits above the FPGA on the board, and does not increase the footprint of the main board or interfere with other componentry.

F. Power

In order to ease the integration of the control board into conventional electronics environments, a power conditioning section is included to derive all of the necessary supply

voltages from a standard 5 V line. Since the FPGA requires both a primary power source of 3.3 V and an internal power level of 1.2 V, 3.3 V and 1.2 V regulators are included on board. In order to adhere to those voltage levels, all other logic on board is also powered by 3.3 V, with the included buffer ICs offering the ability to be powered by 5 V for compatibility with 5 V components.

In addition to the power that the electronics of the system requires, musical robots often require additional higher voltage supplies to their actuators. This can be just 5 or 6 V for servomotors, 12 or 24 V for stepper motors or solenoids, and solenoids will often be powered at 48 or 60 V for increased performance. In order to maintain flexibility and compatibility with a large number of musical robot systems, each MOSFET bank has independent higher voltage lines. This allows each bank to accommodate actuators with different supply voltage requirements, up to a maximum of 6. However, if fewer or only a single higher voltage level is required, temporary or permanent jumpers can be placed between the supply lines of the banks, connecting them.

IV. Evaluation

In order to evaluate whether or not the created system achieved the desired performance, a trial was conducted to determine the accuracy and consistency of the generated pulses. The output pulses were measured with an Agilent DSO-X 2024A digital oscilloscope at a sample rate of 1 Gigasample per second, and 100 measurements were taken to provide a margin of error for the trial.

Figure 5 shows that of the 100 measurements of the pulses taken, 39 of them were measured to be of the correct width to the nanosecond, with the rest of the measurements showing predominantly 1 and no more than 3 nanoseconds of deviation. This accuracy is two orders of magnitude higher than the required specification, and an improvement of 5 orders of magnitude over the Closed-Loop Robotic Glockenspiel.

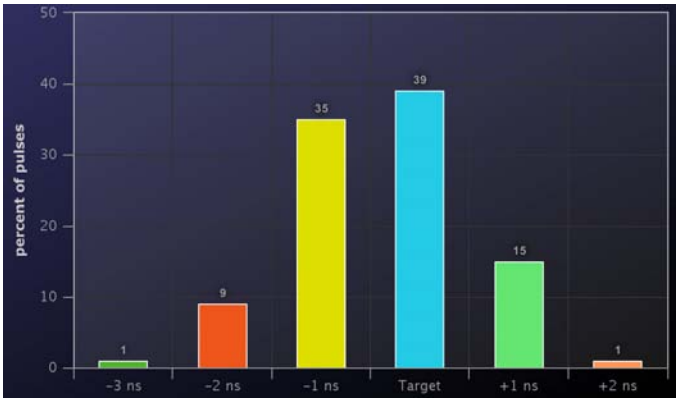

Fig. 5. Graph showing the accuracy of measured output pulses.

V. Conclusion

This paper has described significant hurdles that are encountered when attempting to achieve high levels of accuracy, consistency and reliability with musical robot actuation, and presented a novel method of overcoming these difficulties by utilizing the massively parallel capabilities of FPGAs. The design and capabilities of a new set of hardware created around this principle were then presented in order to demonstrate an implementation of the described concept. This implementation is capable of consistently generating control signals for musical robot actuators with a timing granularity of as low as 20 nanoseconds, a significant improvement over microcontroller only designs, which in the author's experience with the Closed-Loop Robotic Glockenspiel could generate timing errors as high as 100 microseconds.

The secondary requirement that the hardware implementation should be flexible with regards to actuator and MOSFET number and choice was fulfilled by utilising a modular design, organising the actuators into independent banks, each with individual voltage supplies, and buffer circuitry. Additionally, the requirement that it must be possible to program and configure the board with both very high level tools such as the Arduino and Mojo IDEs and low level tools such as the Xilinx ISE and AVR-GCC, was fulfilled by selecting compatible FPGA and microcontroller parts and complying with the publicly available hardware schematics of these systems.

It is hoped that the description of this concept and the accompanying system will aid other musical roboticists in creating new automatic musical instruments that can conduct actuations accurately and consistently. This ability will make musical robots capable of performances that more precisely reflect the composer's original intentions while also facilitating the improvement of auto-calibration systems that increase the robustness and usability of the instruments.

References

[1] J. Murphy, J. McVay, P. Mathews, D.A. Carnegie, and A. Kapur, "Expressive Robotic Guitars: Developments in Musical Robotics for Chordophones" *Computer Music Journal*, 2015.

[2] F. Biedrzycki, J. Knast, M. Nowak, and J. Paszkowski, "AutoPRK-Automatic Drum Player" *Proceedings of the International Conference on Computer Information Systems and Industrial Management*, 2015.

[3] S. Takashima and T. Miyawaki, "Control of an Automatic Performance Robot of Saxophone: Performance Control Using Standard MIDI Files" *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems-Workshop: Musical Performance Robots and Its Applications*, 2006.

[4] A. Focker, *Trimpin: Contraptions for Art and Sound.* Marquand Books, Seattle, Washington, 2011.

[5] J. Solis, T. Ninomiya, K. Petersen, and A. Takanishi, "Anthropomorphic Musical Performance Robots at Waseda University: Increasing Understanding of the Nature of Human Musical Interaction." *Proceedings of the International Conference on New Interfaces for Musical Expression*, 2009.

[6] E. Singer, J. Feddersen, C. Redmon, and B. Bowen, "LEMUR's Musical Robots" *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME, June 2004.

[7] S. Leitman, "Trimpin: An Interview" *Computer Music Journal.*, 35(4):12-27, 2011.

[8] A. Kapur, "A History of Robotic Musical Instruments." *Proceedings of the International Computer Music Conference* 2005.

[9] J.W. Murphy, A. Kapur, and D.A Carnegie, "Musical Robotics in a Loudspeaker World: Developments in Alternative Approaches to Localization and Spatialization." *Leonard Music Journal* 22:41-48, 2012.

[10] J. Long, J.W. Murphy, A. Kapur, and D.A. Carnegie, "A Comparative Evaluation of Percussion Mechanisms for Musical Robotics Applications." *Proceedings of the 6th International Conference on Automation, Robotics and Applications.* February 2015.

[11] J. Long, J.W. Murphy, A. Kapur, and D.A. Carnegie, "A Methodology for Evaluating Robotic Striking Mechanisms for Musical Contexts." *Proceedings of the International Conference on New Interfaces for Musical Expression* 2015.

[12] A. Kapur, Trimpin, A. Suleman, and G. Tzanetakis, "A comparison of solenoid-based strategies for robotic drumming." *Proceedings of the International Computer Music Conference* 2007.

[13] G.W. Raes, "Expression control in automated musical instruments " *Logos Foundation* http://logosfoundation.org/g_texts/expression-control.html. Accessed 28 September 2016.

[14] J. Long, A. Kapur, D.A. Carnegie, "An Analogue Interface for Musical Robots." *Proceedings of the International Conference on New Interfaces for Musical Expression* 2016.

[15] J. Long, A. Kapur, D.A. Carnegie, "The Closed-Loop Robotic Glockenspiel: Improving Musical Robots Using Embedded Musical Information Retrieval" *Proceedings of the International Conference on New Interfaces for Musical Expression* 2016.

[16] J. Murphy, A. Kapur, D.A. Carnegie, "Interacting with Solenoid Drummers: A Quantitative Approach to Composing and Performing with Open-Loop Solenoid-Based Robotic Percussion Systems" *Proceedings of the Australasian Computer Music Conference* 2012.

[17] D. Diakopoulos and A. Kapur, "HIDUINO: A firmware for building driverless USB-MIDI devices using the Arduino microcontroller." *Proceedings of the International Conference on New Interfaces for Musical Expression.*, NIME, May 2011.

[18] A. McPherson, R. Jack and G. Moro, "Action-Sound Latency: Are Our Tools Fast Enough?" *Proceedings of the International Conference on New Interfaces for Musical Expression* 2016.

[19] J. Long, "Augmenting Virtual Worlds with Musical Robotics" *Proceedings of the International Symposium on Electronic Art* 2015.

[20] A. McPherson, "Techniques and Circuits for Electromagnetic Instrument Actuation" *Proceedings of the International Conference on New Interfaces for Musical Expression* 2012.

Architecture for the Transient Array Radio Telescope version 3 (TART3)

J.P.W. McPherson and T.C.A. Molteno
Electronics Research Group
Department of Physics
University of Otago
mcpherson@protonmail.com

***Abstract*—We propose a new version of the Transient Array Radio Telescope (TART) based on the currently functional TART2. The TART3 will utilise a more powerful central single board computer (BeagleBone Green) that will replace two separate computational components (Xilinx Spartan 6 LX9 FPGA and Raspberry Pi 3). Due to the increase in power and replacement of certain components, we hope to produce a more efficient and cost effective aperture synthesis radio telescope with a smaller physical footprint.**

I. Introduction

A. Basic functionality and processes of a radio telescope

High energy solar particles collide with and pass through the atmosphere every day causing hadronic interactions. These events are referred to as solar scattering and often the larger particles cascade; deteriorating further into more fundamental particles in a process that is known as an extensive air shower [3]. Each of these charged particles emit radio waves as they spiral by basic electrodynamic law, a fact that has been experimentally ratified as early as 1965 [7]. These radio pulses are mainly produced due to geomagnetic emission. Geomagnetic emission is caused when charged particles are accelerated within a magnetic field [13]. In the case of atmospheric scattering, the magnetic field is that of the earth and the acceleration is caused by the change in velocity (both magnitude and direction) due to the collisions with atmospheric air particles. These interactions occur continuously, like a chain reaction until the kinetic energy of the solar particles is sufficiently reduced. During this event, a net drift of charged particles occurs, with electrons and positrons moving in the opposite directions. This is due to the influence of the Lorentz force, denoted by

$$\vec{F} = q\vec{v} \times \vec{B}$$

Where q is the particle charge, $\vec{v}$ is the vector for velocity, and $\vec{B}$ is the vector for magnetic field. This means that there exists a transverse current due to the movement of charged particles with respect to the direction of the shower axis. These currents themselves will obviously vary with time, and it is the time variation of these currents that produces the electromagnetic signature which we can observe as radio waves. These radio waves, if received by an antenna, will induce a voltage across said antenna, thus producing an electronic signal. By monitoring the change in these signals, the particles can be tracked as they further cascade, giving data on what particular hadronic particles deteriorate into in such a circumstance. Furthermore, the secondary particles themselves can further split depending on the energies involved. This can lead to a greater understanding of subatomic structure, the process of creation and the nature of elementary particles [13].

Arrays of antennas are used to give a more accurate reading of the radio waves, and to give a more accurate measurement of distance. By having multiple antennas receive the same signal, the time delay between each antenna pulse will give a path difference. This path difference, along with the incident angles of the radio signal taken from the azimuth can be used to formulate a trajectory from each antenna as to the source of the electromagnetic pulse. The intersect of these trajectories will then give an accurate reading as to the location of this source [13] in a full spherical coordinate system. It must also be noted that due to the rotation of the earth, each antenna will have a unique velocity in three dimensional space relative to the source, so each antenna will undergo a disparate Dopler shift. This needs to be accounted for (along with other confounding factors) in the analysis of the data. For obvious reasons, it is then important to have more antennas in an array, as well as an efficient mapping for the placement of the antennas within an array. As such, the development of the TART3 aims to accommodate a minimum of 24 antennas. Utilising this many antennas within a transient array has many applications, especially when the array is economical to produce. For more details on placement of antennas within an array, Huege [3] section five is recommended.

B. The TART2 radio telescope

The Transient Array Radio Telescope version 2 (TART2) is a small scale antenna based radio telescope array designed and created by the Electronics group at the University of Otago. The TART2 is a 24-element aperture synthesis array radio telescope [9]. Aperture synthesis telescopes simultaneously sample a radio signal over an array of separate antennas. Each signal recorded by the individual antennas are then digitised at a 1-bit resolution and correlated [2]. As the antennae are running synchronously, these digital results can then generate an image of the celestial sphere. The TART2 in particular has an angular resolution of approximately 3 arc-minutes [9]. As

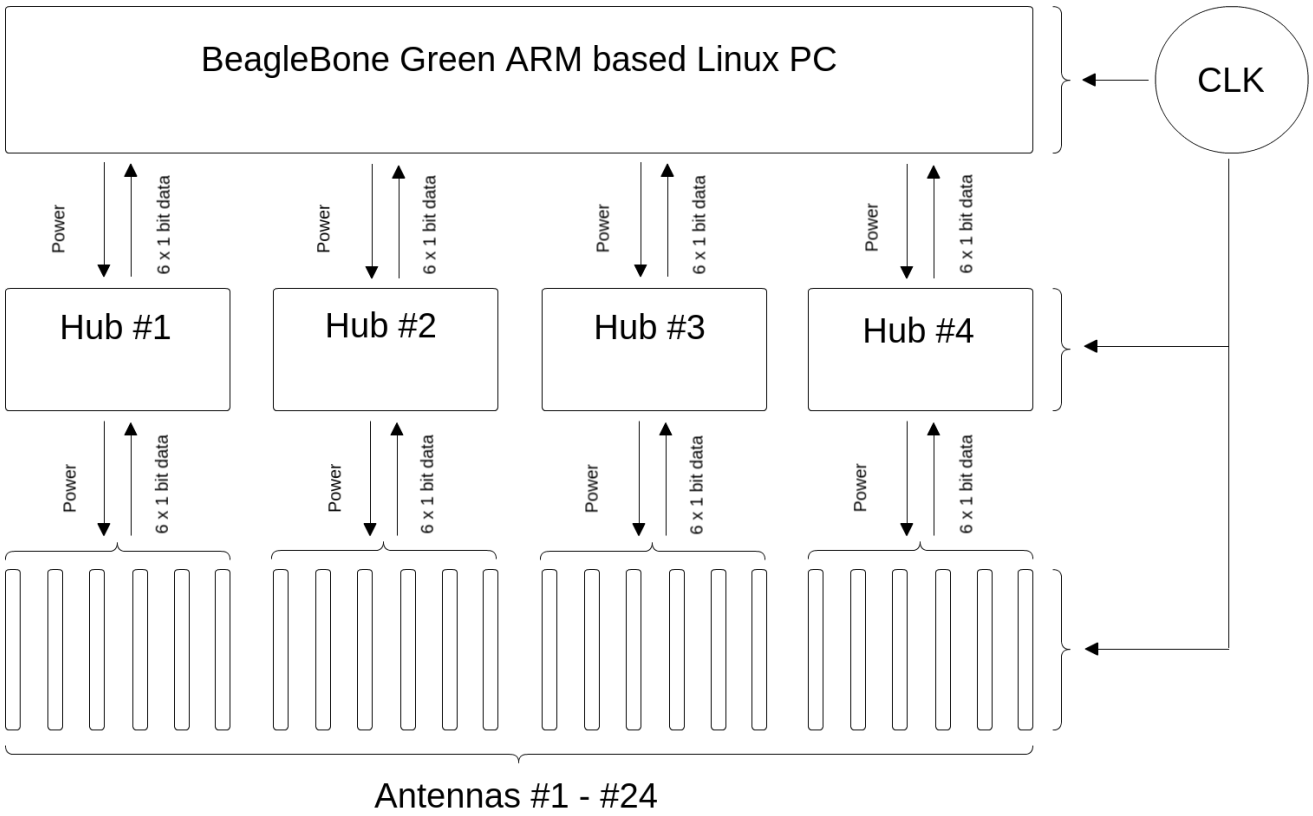

Fig. 1. A basic system diagram of the proposed TART3.

a comparison, the angular size of our sun is roughly 32 arc-minutes.

The TART family operates within the L1 GPS band. This is centered at 1.575GHz. This is a useful band to use as satellite navigation signals generate insufficient noise to interfere with the received data, due to the fact that they transmit below the noise floor [1]. Therefore, they do not alter the processed signal in any meaningful way. GPS satellites are also useful for imaging [13] and antenna calibration [9] as they are known locations in the celestial sphere, and can therefore be used as a reference point.

The 24 antenna elements of the TART2 consist of four separate hubs that run six antennae each. Each antenna hub connects to the baseboard, which consists of a Raspberry Pi 3, Spartan 6 LX9 FPGA and oscillator. The FPGA is a modified Papilio Pro version, which includes 64Mb SDRAM and 48 I/O lines, along with a JTAG programmer [4]. The oscillator sets the clock for the Spartan 6, which connects directly to the antenna hubs. Each antenna supplies a single bit-stream of data back to the Spartan 6 at a rate of approximately 16.368MHz [9]. These data streams are then passed back through the FPGA and storred within the Spartan 6's SDRAM. Due to the limited size of the the Spartan 6's SDRAM (64Mb), the TART2 can only record for intervals of approximately one second before the SDRAM is full. Once the SDRAM fulls up, data collection halts and the information stored within the SDRAM is released to the Raspberry Pi for processing. Therefore, data is recorded in small snap-shots approximately once per minute.

The TART system currently utilises Taoglas model A.01.C.301111 GPS antennas. The effective coverage for each of these is close to hemispherical, which results in the TART having the ability to retrieve signals from the entire sky [12].

The TART2 utilises a clock conditioning system in order to nullify any clock jitter present at the radio receivers based within the antenna hubs [9] [8]. This allows the distributed clock signal to be scrubbed, buffered and redistributed to these receivers via a low bandwidth phase locked loop [10].

All TART designs are open-source. This includes both hardware and software, released under the GNU GPLv3 license. These designs are currently available through the TART website (https://github.com/tmolteno/TART).

II. Design of the TART3

A. Capabilities

We wish to develop a more powerful and efficient radio telescope. While the TART2 is sufficiently capable for our needs, there is always room for improvement. The main areas of interest are decreasing the form factor, increasing the data processing speed, increasing the number of antennae, and reducing the cost of the TART system. Most of this is achievable by a re-design of the control baseboard, then adapting the antenna hubs and connections as needed in order to make the system as efficient as possible.

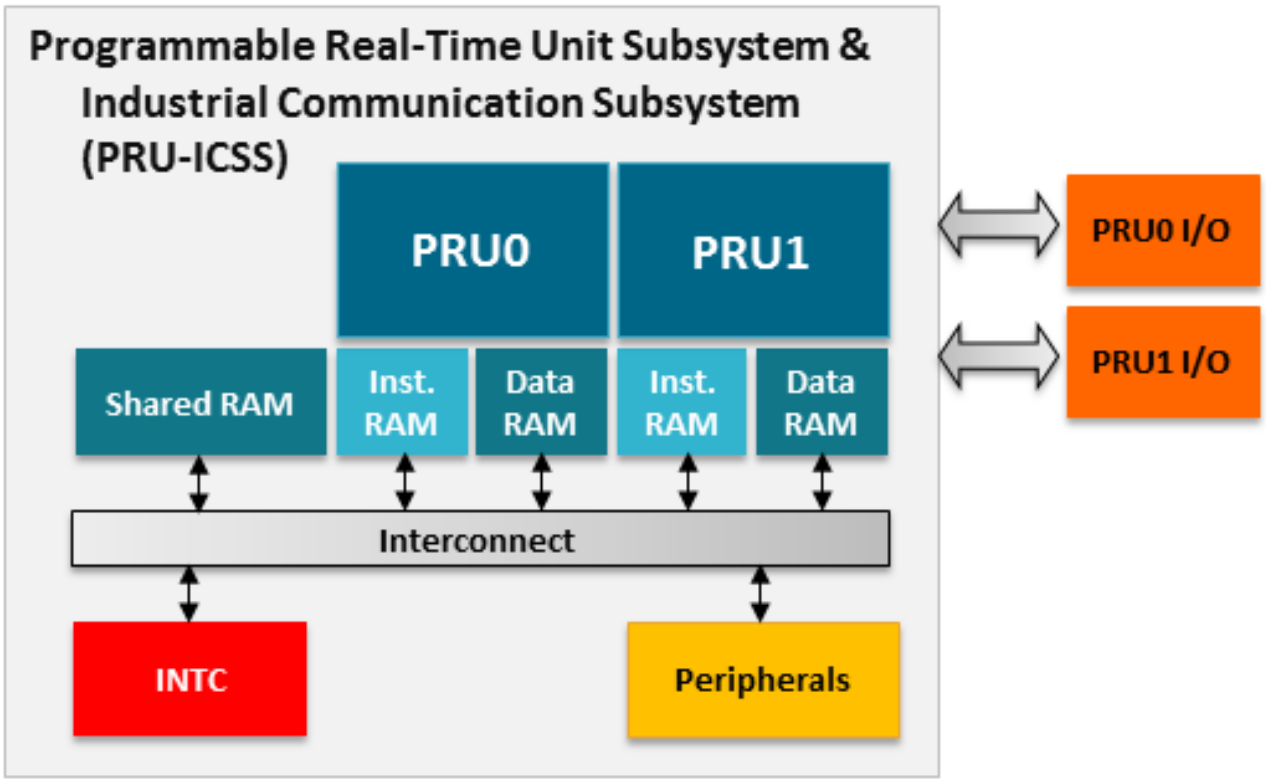

Fig. 2. System diagram of the PRU-ICSS. Taken from processors.wiki.ti.com/index.php/PRU-ICSS

Parameter	Value
Processor	AM335x 1GHz ARM Cortex-A8
RAM	512MB DDR3
On-board Flash Storage	4GB eMMC
CPU Supports	NEON floating-point & 3D GA
Micro USB	Supports powering & communications
USB Host	1
Grove Connectors	2 (One I2C and One UART)
GPIO	2 x 46 pin headers
Ethernet	1
Operating Temperature	0 ~ 75°C

TABLE I
SPECIFICATION TABLE FOR THE BEAGLEBONE GREEN

B. Hardware

The core of the TART3 is the BeagleBone Green single board computer, which will be running a Debian Linux operating system. The BeagleBone Green (BBG) features an AM335x 1GHz ARM Cortex-A8 processor, with 512MB of DDR3 RAM. The AM335x is a 32-bit processor. The BeagleBone Green will replace both the Spartan 6 and the Raspberry Pi 3. A basic system diagram of the TART3 can be seen in Figure 1, while the full specifications of the BeagleBone Green can be found in Table I.

It is unknown at this point if the external clock supplied by the oscillator will still be required, although it is very likely. As a minimum, the implementation of the BeagelBone Green will help to reduce the cost and form factor of the TART.

The antenna hubs will most likely be compromised of the same hardware components, although this is open to change as there may be more powerful of efficient models now available. In saying this, the layout of the hubs will need to be altered to best accommodate the new central base station.

C. Software

The main strength of the BeagleBone Green lies with its AM335x Sitara series processor. In particular its two Programmable Real-Time Unit Subsystem (PRUs) [6]. The PRU used within the AM335x is the PRU-ICSS, which contains the two 32-bit RISC PRU cores. ICSS stands for Industrial Communication Sub-System [5]. The architecture of the AM335x can be seen in Figure 3 while the PRU-ICSS system can be seen in Figure 2. These PRUs are what allows the BeagleBone to replace the Spartan 6. They are fast (200MHz) with a single cycle I/O access, along with full access to the internal memory. The PRU cores also contain their own DRAM with single-error detection. In addition, an enhanced GPIO module generates shift I/O support and parallel latch on external signal. This allows implementation of real-time responses [5]. Being freely programmable, this is where the core of the work on this iteration of the TART will take place.

D. Goals and Time-Line

The TART3 is due to begin development in earnest in early October 2016. The initial stages will consist of the writing of

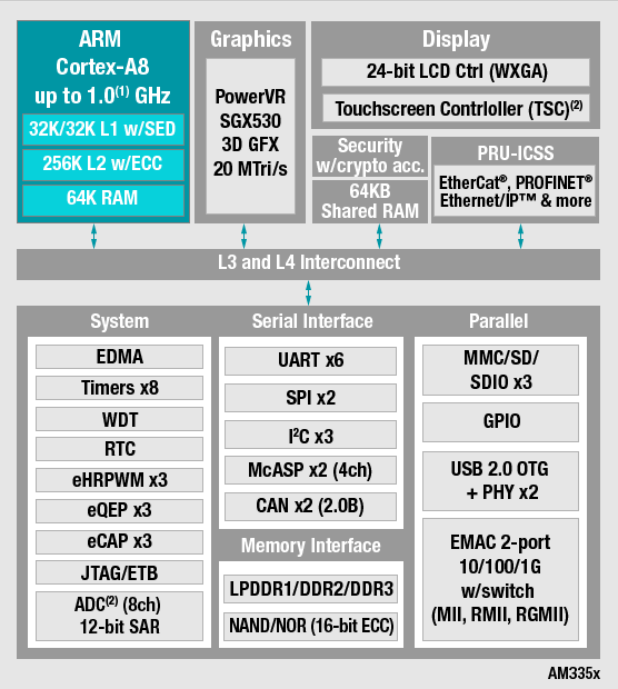

Fig. 3. Architecture of the AM335x Processor. Taken from www.ti.com/lsds/ti/processors/sitara/arm_cortex-a8/am335x/overview.page

fundamental code in order to allow the BeagleBone Green to replace both the FPGA and Raspberry Pi. This means that it will have to be functionally equivalent to the combined system of the Spartan 6 and Raspberry Pi. Development will initially focus primarily on the PRU-ICSS cores. In order to accommodate this, the rest of the TART3 system will clearly be subject to change and alteration to best make use of the BeagleBone. As such the system diagram previously mentioned is an outline only. The BBG will be coded in C++. In order to test functionality, the BeagleBone will be passed manufactured (or pre-analysed) data with a known noise factor. The effectiveness of the system will be measured by comparison of the true (known) distribution and that determined by the new algorithms.

Due to the 32 bit processor within the BeagleBone Green, there is the possibility that the TART3 will be able to process one bit streams from up to 32 antennae synchronously. This will give a huge boost to the resolution and accuracy of the TART system if the system can be improved in this way. After the initial coding this will be one of the main points of investigation.

Once the BeagleBone is operational, modifying the rest of the base-station will be the next major task. This will be done in conjunction with any changes that need to be made to the antenna hubs. This is due to that fact that these pieces of hardware directly interact, and so developing them in isolation could cause problems.

Past this, the TART3 will enter a full system testing phase and the design and components should be finalised.

III. Conclusion

The TART project is completely open-source. As they are developed, all design assets will be made available under the GPLv3 [11] open-source license.

In summation, we plan to create a new version of the Transient Array Ratio Telescope that has a smaller form factor, is cheaper to produce, is more efficient and has the capability to run more antennas. The TART3 is a logical progression from the currently functional TART2, and as such the finished product will have many similarities. The main area of innovation from the TART2 to the TART3 is the inclusion of the BeagleBone Green with its AM335x processor and in-built PRU-ICSS cores.

Acknowledgement

The TART project is supported with grants received from the University of Otago, along with generous donations from Helicopter Line. The authors would like to thank Maximilian Scheel for clarifying the workings of the TART2, and the volunteers who assisted in the restoration of the telescope hut site.

References

[1] Kai Borre, Dennis M Akos, Nicolaj Bertelsen, Peter Rinder, and Søren Holdt Jensen. *A software-defined GPS and Galileo receiver: a single-frequency approach*. Springer, 2007.

[2] R Hanbury Brown and Richard Q Twiss. A new type of interferometer for use in radio astronomy. *Philosophical Magazine*, 45(366):663–682, 1954.

[3] Tim Huege. Radio detection of cosmic ray air showers in the digital era. *Physics Reports*, 620:1–52, 2016.

[4] Gadget Factory Inc. *Papilio Pro*, Accessed September 2016. Available at http://papilio.cc/index.php?n=Papilio.PapilioPro.

[5] Texas Instruments. *PRU-ICSS*, Accessed 29-09-2016. Available at processors.wiki.ti.com/index.php/PRU-ICSS.

[6] Texas Instruments. *Sitara AM335x Processors*, Accessed 29-09-2016. Available at www.ti.com/lsds/ti/processors/sitara/arm_cortex-a8/am335x/overview.page.

[7] J.V. Jelley, J.H. Fruin, N.A. Porter, T.C. Weekes, F.G. Smith, and R.A. Porter. Radio pulses from extensive cosmic-ray air showers. *Nature*, 205:327–328, 1965.

[8] Charles Shaw and T.C.A. Molteno. A GPS Signal Synthesiser for Systems Integration Testing of the Transient Array Radio Telescope. In *The 19th Electronics New Zealand Conference (ENZCon)*, pages 79–84. University of Otago, December 2012.

[9] Charles Shaw, Maximilian Scheel, and T.C.A. Molteno. Transient Array Radio Telescope: Design. In *The 21st Electronics New Zealand Conference (ENZCon)*, pages 28–33. University of Waikato, November 2014.

[10] Silicon Laboratories Inc. *Si5317 Pin Controlled 1-711 MHz Jitter Cleaning Clock*, April 2011. Version 1.1.

[11] Brett Smith. A quick guide to gplv3. 2007, Accessed 25-09-16. Available at http://www.gnu.org/licenses/quick-guide-gplv3.

[12] Taoglas antenna solutions. *A.01.C.301111: GPS Hercules Heavy Duty Screw Mount Antenna*, 2013.

[13] A.R. Thompson, J.M. Moran, and G.W. Swenson. *Interferometry and Synthesis in Radio Astronomy*. Wiley, 2008.

Real-time control using sequential inference and Bayesian decision theory

A. D. Martin and T. C. A. Molteno
Department of Physics
University of Otago
Dunedin, 9016, New Zealand
Email: amartin@elec.ac.nz

***Abstract*—This paper formulates a framework for real-time control using a Bayesian control-system. A Bayesian controller uses a Bayesian Sensor, which can characterise its own uncertainty. The sensor outputs a posterior distribution, P, and the controller, given this distribution, can calculate the cost of different actions, a. It chooses the action, a_P that minimizes the expected cost; this is known as the Bayes act. This framework is demonstrated by an example of detecting mains frequency fluctuation, and shutting down equipment to mitigate the resulting damage. The Bayesian controller is contrasted with a conventional control system, which uses Fourier methods. The Bayes act is calculated for a range of equipment and shutdown cost ratios and the algorithm performance is verified. The Bayesian controller detects frequency changes much faster than the conventional controller, and shuts the equipment down at the optimum time to minimize loss.**

I. Introduction

Bayesian decision theory is a probabilistic approach to decision making [1]. It is the optimal way of making decisions given uncertainty, and becomes particularly important when making decisions in a noisy environment. Its utility is illustrated by the human nervous system during sensorimotor control, which has evolved to act consistently with Bayesian decision theory [2]. The primary use of Bayesian decision theory has been in the field of economics [3]. In engineering, it has been mainly limited to decisions on equipment and plant design, see, e.g., Ref. [4] although, the robotics community has recently begun to apply its techniques, e.g., in machine learning [5]. However, Bayesian estimation is widely used in engineering [6], [7], [8], [9], [10], [11], and the linear-quadratic-Gaussian controller [12] can be cast in the form of Bayesian decision theory.

In response to increased computational power and the proliferation of embedded sensors, we propose a framework for control system design which uses Bayesian decision-making on sensor data in real-time. This method will allow control systems to operate optimally in noisy environments, or when the system is otherwise uncertain. Figure 1 shows the difference between a conventional control system, and what we name a Bayesian control system. The conventional control system compares sensor measurements, or some derived quantity, y_m with a setpoint s, which represents the ideal signal. The difference, dy is used by the controller to send the control signal, u to the system. In contrast, the Bayesian controller uses a Bayesian Sensor, which can characterise its own uncertainty [11]. Specifically, the sensor outputs a posterior distribution, P, and the controller, given this distribution, can calculate the cost of different actions, a. It chooses the action, a_P that minimizes the expected cost. The Bayesian sensor presents several advanages: its rigorous uncertainty quantification allows the controller to be better informed. Furthermore, the controller will make more sophisticated decisions based on a wide range of business factors.

In this paper, we demonstrate a practical example of a Bayesian control system to protect equipment from frequency fluctuations in a noisy mains signal. To illustrate, we simulate a noisy voltage signal, and use a fast Bayesian inference algorithm to estimate the frequency at each sample point, along with a corresponding uncertainty estimate. We show that the Bayesian method detects the frequency change much more quickly than conventional Fourier methods. The Bayesian controller takes the cost of the equipment and the cost of equipment shut-down as inputs, and uses Bayesian decision theory to shut the equipment down at the optimal time to minimise loss. We present results for different noise models, sampling rates and cost ratios.

II. Example System

To demonstrate the principles of Bayesian decision theory as straightforwardly as possible, we simulated an idealized system. Of course, a real system will be more complicated, and will require modified methods; these methods should, nevertheless, follow the same principles that we demonstrate here. In our ideal system, we assume equipment that is rated to operate at a voltage frequency of 50 Hz, with a tolerance δ. We assume that outside this tolerance, the equipment will break practically instantaneously. We tested our algorithms for frequency-change detection on simulated a mains voltage time-series, when the frequency was allowed to vary. We simulated a mains voltage time-series with an initial constant frequency of 50 Hz. After a time, we ramped the simulated frequency at a constant rate to a new value. The signal remained at this new frequency for the remainder of the simulation. We sampled the simulated signal at different sampling frequencies (5000 Hz and 500 Hz), to test our methods when the signal was more or less well resolved by the sample points.

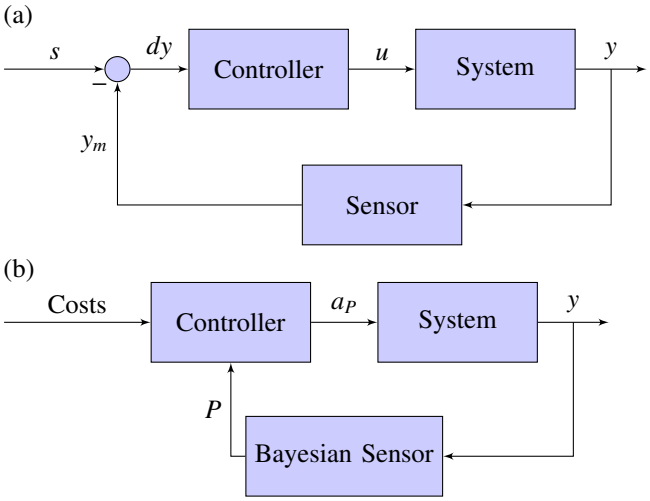

Fig. 1. (a) Conventional control system which compares sensor measurements y_m with a setpoint, s, to inform a controller to send signal control signal u to the system. Here, $dy = s - y_m$. (b) Bayesian control system, which requires a Bayesian sensor characterises its measurement uncertainty. In general, the Bayesian sensor issues a series of posterior distributions, P to the controller. The controller uses this distribution along with the costs encoded in a loss function to apply the Bayes act, a_P, to the system.

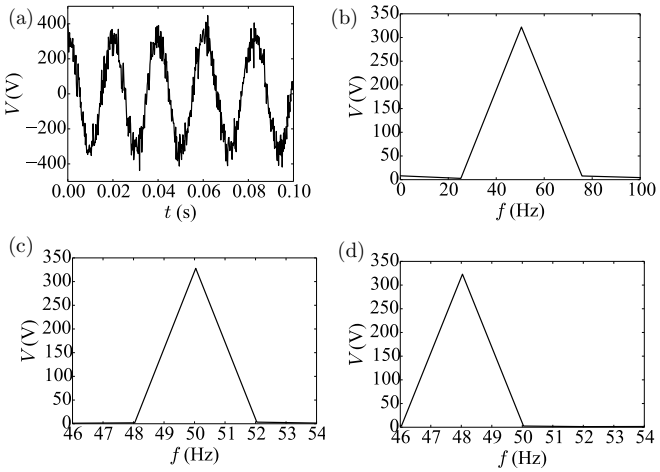

Fig. 2. (a) Simulated mains voltage time-series, with frequency ramped from 50 Hz to 48 Hz starting after 0.05 s. The noise at each timestep is Gaussian with covariance 3000 Hz^2. (b) Section of the fast Fourier transform of the first 0.04 s of the mains voltage timeseries shown in figure (a), i.e., before the ramp. (c) Section of the The Fourier transform of the first 0.5 s of a longer simulated mains voltage timeseries than in figure (a), taken before the frequency ramp. (d) The Fourier transform of a later 0.5 s of the timeseries than treated in figure (c), taken after the frequency ramp.

The mains voltage can be corrupted by distortion, which can be categorised as non-harmonic distortion, flat-top harmonic distortion and notching harmonic distortion [13]. Non-harmonic distortion is essentially random noise. To test the robustness of our methods to noise, we simulated additive Gaussian observation noise at each sample point. Harmonic distortion is a distortion with the same period as the waveform, usually caused by distorting loads. This can cause a flattening of the peaks of the sinusoid, or notching effects [13]. These effects are easy to model by a finite superposition of sinusoids. However, we will leave consideration of harmonic distortion to future studies.

Figure 2(a) shows an example simulation where the initial frequency of 50 Hz was held constant for 0.05 s, followed by a constant frequency ramp, reaching to 48 Hz at time t = 0.09 s after the start of the simulation. The observation noise at each timestep was drawn from a Gaussian distribution with variance 3000 Hz^2. This is much more noisy than we anticipate for a real mains voltage series, but is useful to demonstrate the value of our methods in noisy systems. The system was sampled at 5000 Hz. Figure 3 shows similar simulated time-series, but sampled at 500 Hz. At this sampling rate, there are only ∼6 samples per period, which presents a bigger challenge for frequency inference, especially for the noisier time-series.

III. Sensor Design

Here we compare sensor design for a conventional and Bayesian sensor for our example system. The conventional sensor uses a Fourier method to return frequency measurements, while the Bayesian sensor uses sequential Bayesian inference to return posterior distributions of freqeuncy.

A. Fourier Analysis

The Fourier resolution of a time-series of length T is $1/T$. Therefore to detect a frequency change δ by Fourier methods we need to sample for a time $T = 1/\delta$. This could be a problem if one needs to detect the change more speedily. Figure 2(a) represents a time-series for which the frequency change $\delta = 2$ Hz. To resolve this change, one requires a sample of 0.5 s. This timescale, corresponding to around 24 periods of the mains signal, might be too long to prevent equipment damage. We exemplify this fact by computing the fast Fourier transform (FFT) of noisy simulation data. We show in Figure 2(b) the FFT of the first ∼ 0.04 s of the time-series in figure 2(a) (i.e., most of the time before the ramp). Clearly, the resolution of ∼ 25 Hz is insufficient to resolve any frequency change of the order of 2 Hz. Figures 2(c)-(d) show the FFTs of 0.5 s regions of a similar but longer time-series, before and after the ramp. Here, the FFT obtains sufficient resolution to resolve the frequency change.

B. Sequential Bayesian Inference

Model-based Bayesian sequential inference is a method for updating knowledge a system's state from sequential measurements on the system. If we can assume an observation model for the shape of the signal, model-based inference will resolve frequency changes much faster than Fourier methods. For mains frequency inference, we can safely model the signal as a sinusoid, and we should be able to infer the signal well within one period. Our task is made easier still by Bayesian inference as we can make use of a prior estimate of the frequency and a model for how the frequency changes.

Sequential Bayesian inference represents the state as a probability distribution or probability density function. The

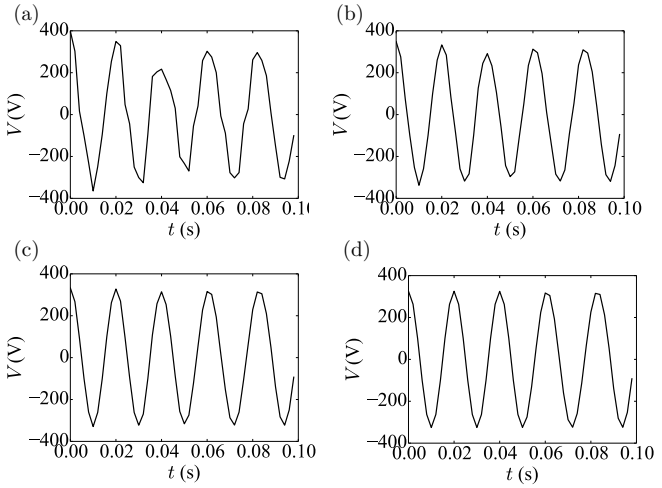

Fig. 3. Voltage time series sampled at 500 Hz, with observation noise from a Gaussian distribution of variance (a) 3000 Hz^2, (a) 300 Hz^2, (a) 30 Hz^2, (d) ~0 Hz^2.

state distribution evolves according to a 'forward map', which encodes the model system dynamics. The algorithm predicts measurements by applying an observation model to the updated distribution. It compares these predictions to the actual measurements, and updates the distribution sequentially using Bayes' rule.

We assume that the frequency and phase of the signal are initially well known. This knowledge is encoded in a narrow prior distribution; we assume a Gaussian prior distribution of frequency with mean $\mu = 50$ Hz and variance $\sigma^2 = 0.0025$ Hz^2. For the inference, we use a state model in which the frequency changes due to a random (Brownian) process. This is different from the simulated process that we used to generate the data, in which the frequency is ramped gradually from a constant value. Our use of a Brownian model instead of a ramp reflects the fact that in a real situation we will be ignorant of the prospective frequency behaviour. It allows us freedom to detect a range of frequency processes, provided the frequency change between samples is not larger than the characteristic frequency scale of the Brownian model. Our inference assumes a process noise at each time-step from a Gaussian distribution of zero mean and variance 0.0025 Hz^2 per time-step. This corresponds to a characteristic rate of 250 Hz/s; for comparison, in the simulated data, the frequency is either stationary or ramped at at rates between 50 and 250 Hz/s.

Our sensor's inference algorithm must be able to handle the sinusoidal (i.e., nonlinear) observation model. We use the Unscented Kalman Filter (UKF) [6], [7], [8], [9], [10], [11], which is suitable for nonlinear systems. The UKF approximates the state distributions by a number of points, a much faster inference than more for general Bayesian inference algorithms, e.g., particle filters and Markov Chain Monte Carlo algorithms. This enables real-time inference at high sample rates.

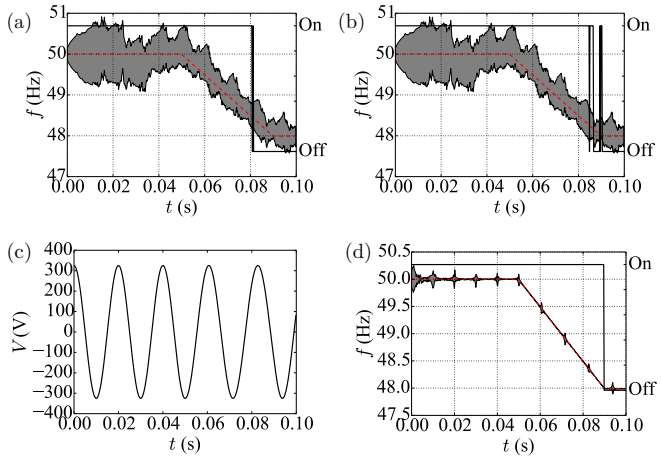

Fig. 4. (a) True frequency from the timeseries 2(a) (red dashed line), with 95% credible region from the UKF output (grey region). The black solid line represent the Bayes act when the ratio, k, between the cost of the equipment and the cost of shut-down is 1000, and the equipment's frequency tolerance $\delta = 2$ Hz. (b) as (a) but with the ratio, $k = 10$. (c) represents a clean timeseries with negligible observation noise. The frequency dynamics are the same as in figure 2(a). (d) as (a), but for the clean timeseries shown in (c). In all cases, the frequency is ramped down at 50 Hz/s.

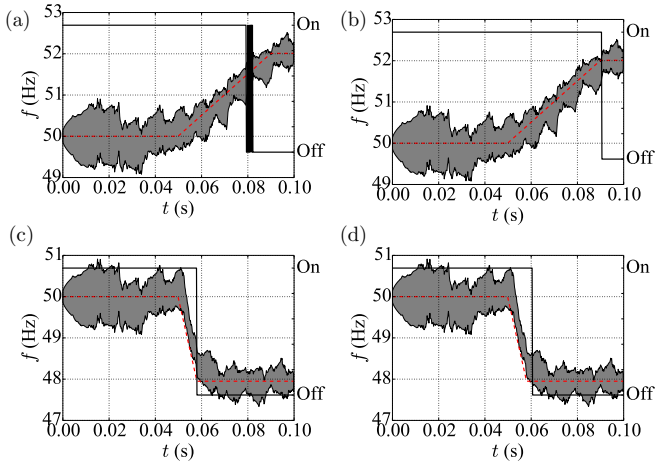

Fig. 5. (a), (b) as Fig. 1 (a) and (b), but with the frequency ramped up, not down. (c), (d) as Fig. 1 (a) and (b), but with the frequency ramped down at rate 250 Hz/s.

The UKF output is a time-series of means $\{\mu_i\}_i$ and variances $\left\{\sigma_i^2\right\}_i$. We interpret each pair $\left\{\mu_i, \sigma_i^2\right\}$ to represent a Gaussian posterior distribution and calculate 95% credible regions for the frequency. Figures 4(a)-(b) show the 95% credible region for the posteriors generated from the time-series shown in figure 2(a). These credible regions show the width and locations of the posterior distributions as a function of time. The credible regions follow changes in the frequency very quickly. When the observation noise is negligible (figure 4(c)), the width of the posterior narrows around the true frequency (figure 4(d)). We show in figures 5(a) and (b), that the frequency inference is similarly effective for up-ramps as for down-ramps. We also

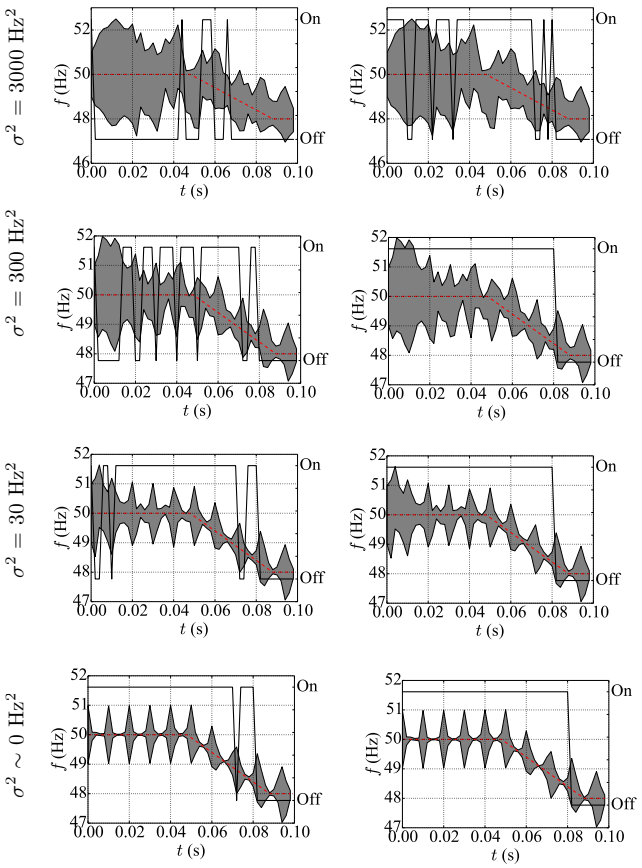

Fig. 6. Inferences from the voltage time-series from Fig. 3 (sampled at 500 Hz). The left column has Bayes act for $k = 1000$, the right for $k = 10$.

tested fast ramps, at a rate 250 Hz/s (see figures 5(c) and (d)). At this faster rate, the true value falls outside the 95% credible region during the ramp, although it falls inside the region afterwards. The loss of accuracy reflects the limit of our Brownian model, in which many consecutive frequency increments of 250 Hz/s are unlikely. Despite a slight loss of accuracy during the ramp, the Bayesian inference detects the rapid frequency change much faster than the Fourier method. Figure 6 shows frequency inference on the simulated time-series of figure 3, where the sample rate is only 500 Hz. It is clear from comparing, e.g., figure 2(a) with figure 3(a) that the small sample rates introduce more ambiguity into the signal, especially when the signal is corrupted by observation noise. This ambiguity is reflected in the larger width of the 95% credible regions for low sample rates (compare the top row of figure 4 with that of figure 6). Reducing the observation noise noise to something more realistic (figures 3(b)-(d)) leads to narrower 95% credible regions, which allows better decision-making as we will see in section IV.

IV. Bayesian control system framework

Figure 1(a) contrasts different control system frameworks. A conventional control system might use the Fourier method, described in section III-A. In this case, the measurements, y_m, in 1(a) are the largest Fourier components over some buffered data. These are compared with a setpoint, s, of 50 Hz, and the controller issues the shutoff signal when the difference dy crosses the threshold, δ, for equipment damage.

In contrast, the Bayesian controller uses a Bayesian Sensor, which can characterise its own uncertainty [11]. Specifically, the sensor outputs a posterior distribution, P of the measured quantity. In our example, frequency posteriors are generated from the UKF. The controller knows the costs of all possible actions available to it, and uses P to make a decision, a, which minimises the expected loss. In our highly simplified example, we assume the cost of shut-down is a constant, s, which is related to the cost of equipment repair, $e = ks$, by some factor k. If the equipment fails, the process will require repair, in addition to shut down, thus incurring total cost $e + s$. The loss function therefore takes the form:

$$L(f, a = \text{On}) = \begin{cases} 0 & \text{if } f \in [f_0 - \delta, f_0 + \delta] \\ e + s & \text{otherwise} \end{cases} \qquad (1)$$
$$L(f, a = \text{Off}) = s,$$

where $f_0 = 50$ Hz. For a probability density function P, the Bayes act against P, a_P, is the act that minimizes the expected loss [1]. In our demonstration, P is the (Gaussian) posterior for frequency generated by the UKF, with mean μ and variance σ^2. Then the Bayes act takes the form

$$a_P = \begin{cases} \text{On}, & \text{if } (k+1)\, F_P < 1 \\ \text{Off}, & \text{otherwise}, \end{cases} \qquad (2)$$

where

$$F_P = 1 - \frac{1}{2}\left[\text{Erf}\left(\frac{\delta + f_0 - \mu}{\sqrt{2}\sigma}\right) + \text{Erf}\left(\frac{\delta - f_0 + \mu}{\sqrt{2}\sigma}\right)\right] \qquad (3)$$

is the probability that $f \notin [f_0 - \delta, f_0 + \delta]$ and Erf is the error function.

As well as the 95% critical regions, Figures 4(a) and (b) also show the Bayes act for different values of the cost ratio k, where the tolerance $\delta = 2$ Hz, for the time-series with the negative frequency ramp shown in figure 2(a). When $k = 1000$, the "off" act is triggered around 0.01 s before the frequency reaches the critical value of 48 Hz; at this time the 95% critical region does not yet overlap the 48 Hz. For a smaller cost ratio $k = 10$, the decision is less cautious, and the "off" act is triggered later, around 0.005s before the frequency reaches 48 Hz, although in this case the 95% critical region now overlaps this value. When there is no observation noise on the signal, and the UKF produces a very narrow posterior, then the "off" act is triggered almost instantaneously as the frequency reaches 48 Hz. The Bayes act displays similar behaviour for the positive frequency ramps in figures 5(a)-(b). When the frequency is ramped faster, near the limit of that the inference algorithm can handle, the Bayes act switches to "off" almost instantaneously as the frequency reaches 48 Hz, and is slightly late when the cost ratio as small ($k = 10$). However, this lag is tiny, and is a factor of 200 smaller than the lag for the conventional control system.

For low sampling rates, as the posteriors are much broader than for high sampling rates (all else being equal), the Bayesian controller is more cautious (see figure 6). For the noisiest data, when the cost ratio $k = 1000$, the Bayes "off" act is triggered soon after the start of the simulation. This caution is necessary, due to the ambiguity in the signal. As the noise amplitude is reduced, or as k is reduced, the controller is less cautious, making the Bayes "off" act closer to time the frequency reaches 48 Hz. These results show that for noisier systems, a larger sampling rate will improve decision-making, particularly when the stakes are high, and controller needs to resolve small frequency changes.

V. Conclusions and further work

In conclusion, we have introduced the concept of a Bayesian control system, which provides a superior method of automated control and decision-making. We have demonstrated the effectiveness of Bayesian sensors that use sequential Bayesian inference, using an industrial example of detection fluctuations in mains frequency. The Bayesian sensor far out-performs conventional Fourier methods. Moreover, it allows decision-making based on a full set of operating considerations and prior information. This allows the control system automatically to take decisions that minimise business costs, in this case detecting a damaging fluctuation in the power supply and shutting down the system. We note that a conventional controller could make use of a loss function using a heuristic estimate of the system's state (e.g., the frequency in our example) as its argument, and calculating the act that would minimize that loss. However, the calculated loss would not be a rigorous quantification of the expected loss. In our example, for all values of k, this would lead to a trivial result where the control system switches off the equipment exactly when the estimate of f reaches the critical value.

In future work, we propose a more realistic model, which will allow for phase uncertainty, and harmonic distortion. In order to verify the performance in a real system, we intend to test our algorithms on real mains data, which we expect to be less noisy than for the highest amplitude noise considered in this paper. Moreover, in a real setting, the loss function will need be much more sophisticated than in equation (1). Additionally, the work can straightforwardly be extended to many other systems, e.g., thermostats and strain gauges. We propose that there is huge scope for Bayesian control systems to proliferate as businesses strive to maximise performance and minimize losses.

Acknowledgment

This work was funded by grant UOOX1208 from the Ministry of Business, Innovation & Employment. We thank Matt Parry for useful discussions.

References

[1] P. D. Grünwald and A. P. Dawid, "Game theory, maximum entropy, minimum discrepancy and robust bayesian decision theory," *The Annals of Statistics*, vol. 32, no. 4, pp. 1367–1433, 2004. [Online]. Available: http://www.jstor.org/stable/3448538

[2] K. P. Körding and D. M. Wolpert, "Bayesian decision theory in sensorimotor control," *Trends in Cognitive Sciences*, vol. 10, no. 7, pp. 319–326, 2016. [Online]. Available: http://dx.doi.org/10.1016/j.tics.2006.05.003

[3] J. C. Harsanyi, "Bayesian decision theory and utilitarian ethics," *The American Economic Review*, vol. 68, no. 2, pp. 223–228, 1978. [Online]. Available: http://www.jstor.org/stable/1816692

[4] D. R. Davis, C. C. Kisiel, and L. Duckstein, "Bayesian decision theory applied to design in hydrology," *Water Resources Research*, vol. 8, no. 1, pp. 33–41, 1972. [Online]. Available: http://dx.doi.org/10.1029/WR008i001p00033

[5] K. Werling, A. T. Chaganty, P. S. Liang, and C. D. Manning, "On-the-job learning with bayesian decision theory," in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Curran Associates, Inc., 2015, pp. 3447–3455. [Online]. Available: http://papers.nips.cc/paper/5860-on-the-job-learning-with-bayesian-decision-theory.pdf

[6] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Fluids Engineering*, vol. 82, no. 1, pp. 35–45, 03 1960. [Online]. Available: http://dx.doi.org/10.1115/1.3662552

[7] S. Julier and J. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, Mar 2004.

[8] E. Wan and R. van der Merwe, "The unscented kalman filter for nonlinear estimation," in *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, 2000, pp. 153–158.

[9] R. van der Merwe and E. A. Wan, "The square-root unscented kalman filter for state and parameter-estimation," in *in International Conference on Acoustics, Speech, and Signal Processing*, 2001, pp. 3461–3464.

[10] Y. Wu, D. Hu, M. Wu, and X. Hu, "Unscented kalman filtering for additive noise case: augmented vs. non-augmented," in *American Control Conference, 2005. Proceedings of the 2005*, June 2005, pp. 4051–4055 vol. 6.

[11] A. D. Martin and T. C. A. Molteno, "Automated weighing by sequential inference in dynamic environments," in *Automation, Robotics and Applications (ICARA), 2015 6th International Conference on*, Feb 2015, pp. 274–278.

[12] M. Athans, "The role and use of the stochastic linear-quadratic-gaussian problem in control system design," *IEEE Transactions on Automatic Control*, vol. 16, no. 6, pp. 529–552, Dec 1971.

[13] V. Gosbell, S. Petera, and V. Smith, "Harmonic distortion in the electrical supply system," *Integral Energy Power Quality Centre, University of Wollongong, Tech. Rep*, 2000.

Measurement of Implant Electrode Leads using Time-Domain Reflectometry to Predict the Resonant Length for MRI Heating

Steven McCabe
Department of Engineering
The University of Waikato
Hamilton, New Zealand
Email: som1@students.waikato.ac.nz

Jonathan Scott
Department of Engineering
The University of Waikato
Hamilton, New Zealand
Email: jonathanscott@ieee.org

***Abstract*—Magnetic Resonance Imaging (MRI) machines can generate hazardous RF heating of patients with implanted neurostimulation leads. Consequently, most patients with these implants are contraindicated from having MRI scans. The level of RF heating has a strong dependence on lead length and is most severe when the length is close to a specific resonant length. Recent studies have shown that simple modifications to the lead construction and insulating material can alter the resonant length and significantly ameliorate this heating hazard, achieving MRI safety. We propose a technique using time domain reflectometry (TDR) to find the resonant length of an arbitrary lead such to minimise the amount of MRI machine time needed to find the length of highest heating. The results are compared with temperature measurements made in a 3-Tesla MRI machine and with a CW dipole radiator in the lab.**

Keywords*—Pulse measurements, reflectometry, biomedical electrodes, medical diagnostic imaging, electrical stimulation, electromagnetic modeling, specific absorption rate, heating, safety.

I. Introduction

Implant leads such as those found in pacemakers, Spinal Cord Stimulators (SCS), and Deep Brain Stimulators (DBS), can be hazardous to a patient undergoing a Magnetic Resonance Imaging (MRI) scan. Fig. 1 shows a typical implant lead for SCS. The RF field generated by a 3 T MRI machine can deliver peak pulses exceeding 30 kW in power at 128 MHz. [1] This can induce significant heating of patient tissue at the distal electrodes [2], [3], well beyond the 1–2 °C safety limits recommended by the International Commission on Non-Ionizing Radiation Protection (ICNIRP) in [4]. This has lead engineers to develop implant leads that are insusceptible to RF heating from MRI. [5]–[8]

An implant lead is most hazardous at a specific resonant length, l_{res}, a parameter which can be exceedingly difficult to calculate and resource intensive to simulate. Typically it is found through measurement in an MRI machine or in the lab with dipole radiators, where heating tests are made on several implant leads varying by length [9]. We present an alternative approach that employs Time Domain Reflectometry (TDR) to allow simple and rapid extraction of l_{res} from a single measurement.

Fig. 1. Distal end of a 62 cm SCS "Octrode lead". Eight platinum electrodes, each 3 mm long and 1.3 mm in diameter, are separated by 4 mm insulating spacers. Electrode numbering starts from '1' (left-most electrode) to '8' (right-most electrode).

II. Distal Heating

An implanted lead can behave as a resonant dipole to the incoming RF field from an MRI machine. [10] Stored energy is transferred along the length of the dipole and can permeate out into the surrounding tissue, especially around the bare electrodes. This gives rise to joule heating and in some circumstances, can reach hazardous levels. Works by [3] and [11] have shown that peak heating occurs when the lead length is about $0.41\lambda_{Pn}$, where λ_{Pn} is the wavelength along the lead, largely determined by the tissue composition and lead design. Calculation of λ_{Pn} is possible but is limited to simple coaxial-like lead structures. [10] Simulation can provide predictions when the complexity is higher but demands considerable resources and impractical computation run times.

Fig. 2 shows the simulated heating induced by an implanted wire within an MRI birdcage. The birdcage was calibrated to deliver a whole-body Specific Absorption Rate (SAR) of 1 W/kg. The three-dimensional simulation took 5 hours in COMSOL Multiphysics 4.4 running on a 3.5 GHz quad-core Intel CPU and consumed more than 60 GB of memory. Additional simulations were also needed in order to confirm the worst-case length for highest heating. For further simulation details refer to [11].

Experimental measurements are usually performed within a torso-shaped phantom inside of an MRI machine. Fig. 3 shows our phantom on the bed of a 3T MRI machine. The phantom is comprised of saline gel with electrical and thermal properties similar to that of human tissue. [12], [13] An implant lead under test is positioned within the gel and a fiber-optic based thermometer is aligned to the distal electrode where the heating is expected to occur.

As λ_{Pn} varies significantly with insulation thickness, permittivity, and geometry, its value is often unknown and is

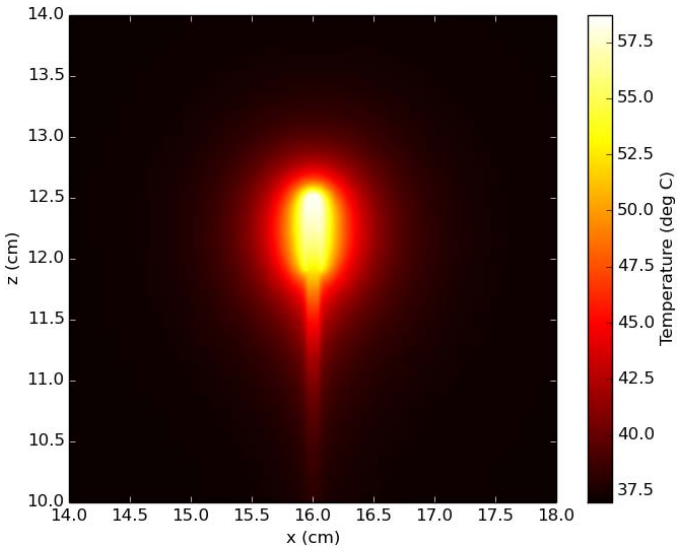

Fig. 2. Simulated heating generated at the 6 mm bared end of an insulated wire after 5 minutes of excitation from an MRI birdcage antenna. The 800 µm diameter wire was $0.41\lambda_{Pn} = 25$ cm in length, coated with plastic insulation 350 µm thick. Blood perfusion was included in the phantom model.

Fig. 3. Operators preparing a wire sample for testing within a gelled saline phantom in a 3T MRI machine.

usually found by individually measuring the distal heating of several wires differing only in length (≥ 10 cm), and looking for the length at which maximal heating occurs. A typical test set as shown in Fig. 4 contains 10 or more samples for a given lead type. The set-up and scan time to measure a single wire sample can easily exceed 30 minutes.

III. TDR Technique

In addition to antenna effects, an implanted lead behaves as an unbalanced transmission line to RF currents. [10] The conductive tissue along the surface of the lead jacket forms the return path for currents, like a shield to a coaxial cable, with the tissue surrounding the distal electrode forming the dissipative load.

TDR is a well established technique for measuring the propagation of signals along transmission lines [14]. Similarly, TDR can be applied to implanted electrode leads. The phase

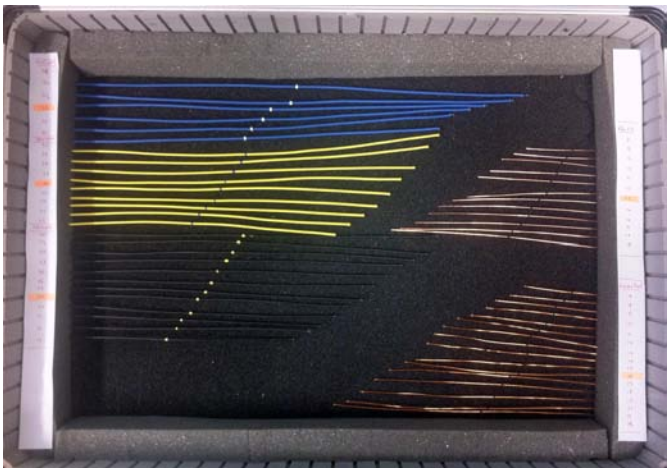

Fig. 4. Briefcase containing a range of wire samples to be tested.

velocity v_p for an implant lead of arbitrary length can be determined by measuring the time taken t_D for a voltage pulse to propagate to the end electrode and reflect back again:

$$v_p = \frac{2l_i}{t_D} \tag{1}$$

where l_i is the length of the insulated portion of the implant lead conductor. As the phase velocity is independent of length, the resonant length, l_{res}, of an implant lead at the MRI RF frequency f_{MRI} is therefore:

$$l_{res} = 0.41\lambda_{Pn} = 0.41\frac{v_p}{f_{MRI}} \tag{2}$$

The RF frequency of MRI is proportional to the strength of the static magnetic field. In a 3 T machine, the RF frequency is 128 MHz. When immersed in a gelled saline phantom from [12], a typical implant lead will have a λ_{Pn} in the order of tens of centimeters [9].

A phantom comprising 28 L of gelled saline in the shape of a torso-and-head was built from clear acrylic after [12]. The ratio of NaCl and polyacrylic acid (PAA) to distilled water was 1.32 g/L and 10 g/L, respectively, with an overall conductivity of 0.47 S/m.

To facilitate TDR measurements of implant leads, the test fixture in Fig. 5 was constructed. A thin aluminium disc rests on the surface of the gelled phantom, providing an electrical path from the shield of the coax cable to the gel. A screw terminal secures the implant lead under test and provides electrical connection to the inner conductor of the coax cable. The implant lead is immersed within the gel, uncoiled, and with the distal electrode unobstructed. Close up views of the fixture are shown in Fig. 6.

Measurements of various wire samples representative of implant leads, were captured with an Agilent 54754A TDR with 40 ps system rise time. The reflection produced by a 29 cm wire sample with 6 mm distal electrode is shown in Fig. 7. The initial sharp change corresponds to the impedance mismatch between the 50 Ω SMA connector and the characteristic impedance of the wire sample. The reflection from the end of the wire sample can be identified by the second discontinuity, where the electrode comes into contact with the dissipative

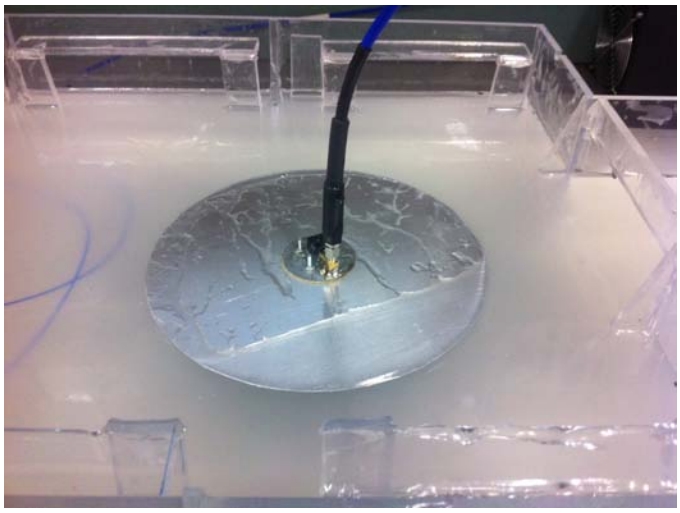

Fig. 5. An aluminium disc with a 20 cm diameter provides electrical connection from the coax cable shield to the gelled saline. Immersed within the gel, a wire sample is connected to the coax cable inner conductor via. a small hole in the center of the disc and secured in place with a screw terminal.

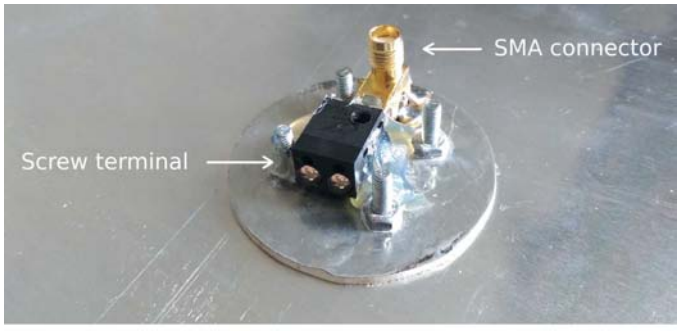

(a)

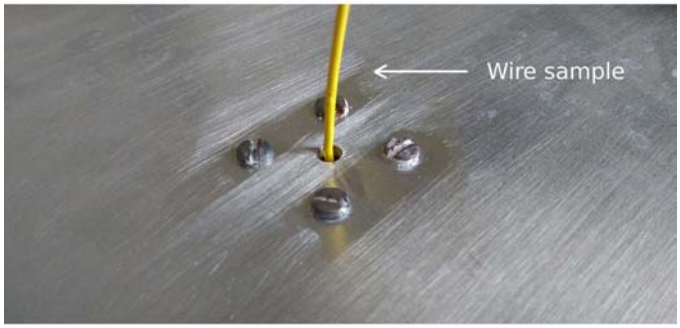

(b)

Fig. 6. Close up view of the disc (a) topside (b) underside.

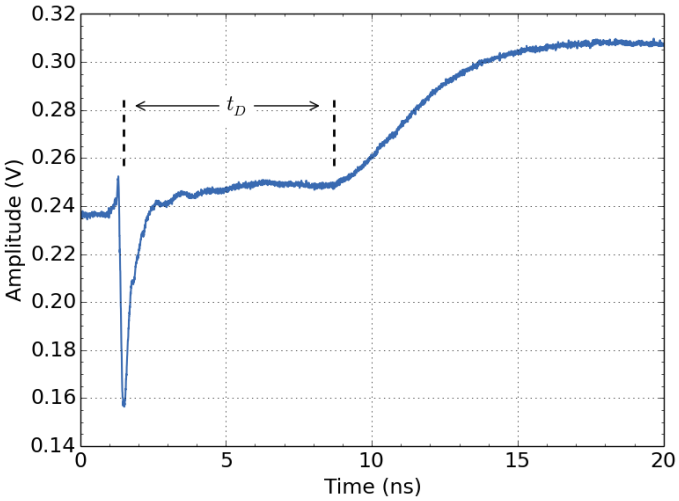

Fig. 7. Reflected TDR signal from a wire sample 29 cm in length, with insulation 350 µm thick. The time delay for a TDR pulse to propagate to and from the distal electrode is 7.18 ns.

gel. Inserting a time delay of 7.18 ns into (2) yields a resonant length of 25.9 cm. Verification of this result is provided in section V.

IV. Verification Technique

Recent work in [9] demonstrated a lab technique for predicting the level of RF heating induced by an implant lead from MRI, without requiring high energy RF pulses from an MRI machine. A dipole antenna driven by a low-power CW is used to provide excitation of the implant under test, with a fiber optic probe to monitor the distal temperature.

The same experimental set-up from [9] was used and is shown in Fig. 8. A dipole antenna was made from rigid 2.1 mm diameter copper wire with 350 µm of insulation covering the entire 32 cm length. A close up view is shown in Fig. 9. A wire sample under test is spaced 6 cm from the dipole antenna and is centered about its midpoint. Temperature of the distal electrode is monitored with a GaAs-based fiber-optic temperature probe with 0.1 °C resolution.

Calibration of the experiment is achieved by scaling the result 1.5 times, such that the reference wire sample generates equivalent heating when exposed to the RF field from the 3 T MRI machine in Fig. 3. The reference measurement along with simulated predictions is shown in Fig. 10.

V. Measured Results

A range of wire samples varying in insulation thickness from 21–700 µm were tested. Each sample consisted of an 800 µm diameter copper core, with plastic insulation covering all but 6 mm from one end. Epoxy resin insulated the opposing end. The reflected TDR waveforms were captured for each wire sample using the test fixture in Fig. 5. The time delay t_D as measured from each waveform are listed in Table I along with the associated resonant length l_{res}, calculated using equations 1 and 2.

Heating tests using the dipole radiator for excitation were performed on each wire sample including several additional

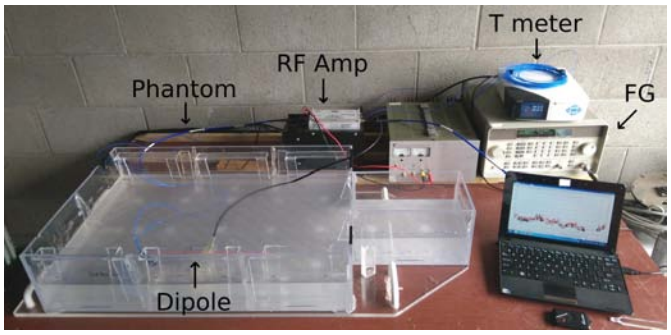

Fig. 8. Set-up for testing the RF heating of implant leads in the lab. A dipole antenna is immersed in the phantom alongside the wire sample with a function generator and 30 W RF power amplifier supplying excitation. A fiber optic thermometer monitors the temperature of the distal electrode.

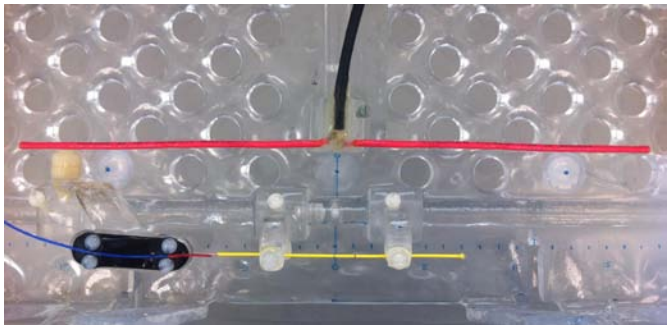

Fig. 9. Close-up view of the temperature probe aligned to the bared end of the wire sample (yellow) before immersion within the gelled saline. The dipole antenna (red) is spaced 6 cm away from the wire sample.

lengths. The change in temperature ΔT was recorded for each sample after 5-minutes of applied RF stimulus. The results shown in Fig. 11 are consistent with [3] and [11]. The lengths for peak heating as extracted from the same figure are summarised in Table I. The TDR-measured values are within 3% of expected values.

TDR measurements were also performed on the electrodes of the SCS lead shown in Fig. 1. The predicted resonant length for electrodes 1 and 8 are twice that of the 700 µm and 350 µm wire samples, respectively. This is owing to the much smaller 100 µm diameter filars that comprise the lead. Heating tests on the SCS lead were not performed.

VI. Conclusion

We explain a measurement technique that predicts the length an implanted lead will experience peak heating during MRI, without an MRI machine. We show this is possible through a TDR measurement of a single lead of arbitrary length.

Acknowledgement

The authors would like to thank Stephen Butler of Midland MRI for his kind assistance. They also acknowledge the Waikato Medical Research Foundation (WMRF) for funding, and WaikatoLink and KiwiNet for their valuable support.

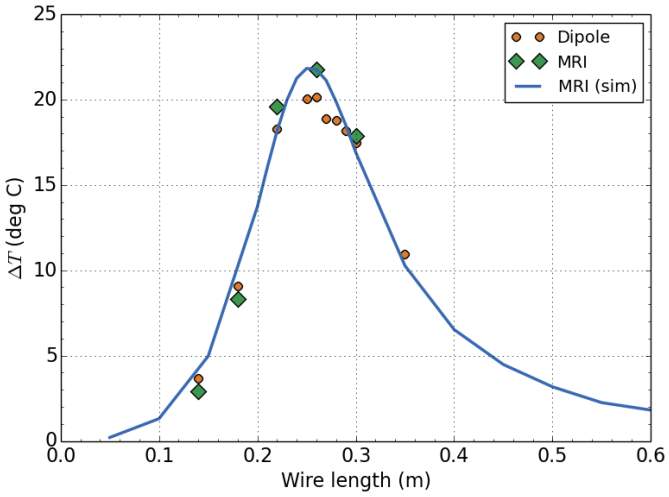

Fig. 10. Distal heating generated near the distal electrode of a reference wire sample coated with 350 µm of insulation, after 5 minutes of CW excitation from a dipole antenna. Measurements of the same wire sample in a 3T MRI machine are also shown along with simulated predictions.

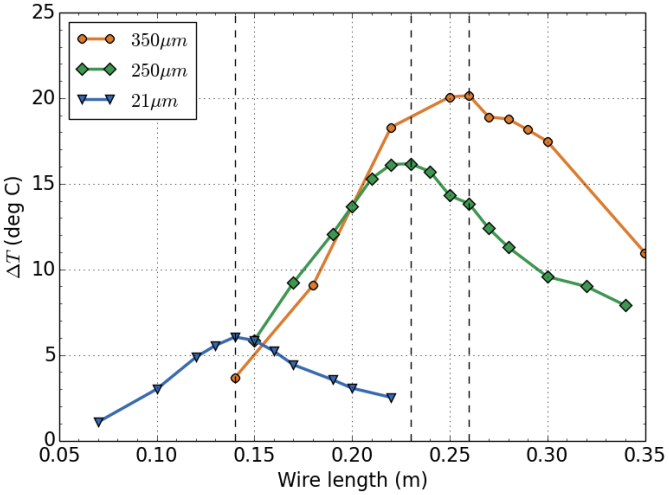

Fig. 11. Distal heating generated for wire samples varying by insulation thickness. Extraction of $0.41\lambda_{Pn}$ and ΔT_{pk} for a given lead type is achieved by measuring several lead lengths.

References

[1] R. Caverly, G. Breed, W. Cantrell, M. Eron, J. Garcia, N. Kondrath, D. Myer, M. Ruiz, and J. Walker, "Advancements at the lower end: Advances in HF, VHF, and UHF systems and technology", *IEEE Microwave Magazine*, vol.16, no.1, pp.28-49, Feb. 2015.

[2] J. A. Nyenhuis, P. Sung-Min, R. Kamondetdacha, A. Amjad, F. G. Shellock and A. R. Rezai "MRI and implanted medical devices: basic interactions with an emphasis on heating", IEEE Transactions on Device and Materials Reliability, vol.5, no.3, pp.467-480, Sept. 2005.

[3] S. McCabe and J. Scott, "Electromagnetic techniques to minimize the risk of hazardous local heating around medical implant electrodes during MRI scanning," *European Microwave Conference 2015*, Paris, 7-10 Sept. 2015.

[4] "ICNIRP statement on: Medical magnetic resonance (MR) procedures: protection of patients", International Commission on Non-Ionizing Radiation Protection (ICNIRP), *Health Phys.*, vol.87 no.2, pp.197-216, Aug. 2004.

[5] S. McCabe and J. Scott, "New MRI-safe implant electrode design," *International Microwave Symposium 2016*, San Francisco, 22–27 May. 2016.

TABLE I. RESONANT LENGTHS OF VARIOUS IMPLANT LEADS: EXPECTED AND TDR-MEASURED VALUES.

Insulation Thickness	Conductor Length (cm)	t_D (ns)	l_{res} from TDR (cm)	l_{res} from ΔT_{pk} (cm)
21 µm	21.2	9.74	13.9	14
250 µm	29.8	8.55	22.3	23
350 µm	29.0	7.18	25.9	26
700 µm	34.3	7.42	29.6	30
Octrode (electrode 1)	61.8	6.68	59.3	-
Octrode (electrode 8)	54.8	7.12	49.3	-

[6] S. O. McCabe and J. B. Scott, "Implant conductor assembly with improved radio frequency properties", New Zealand provisional patent *714212*, Filed on Nov 16, 2015.

[7] J. B. Scott and S. O. McCabe, "Implant conductor with improved radio frequency properties", New Zealand provisional patent *708633*, Filed on May 29, 2015.

[8] *MRI and SCS, SureScan system* Medtronic, USA, Implant Manual, 2013. http://www.mrisurescan.com/au/for-radiologists/spinal-cord-stimulators/mri-and-scs/index.htm#tab2 retrieved Aug 2014.

[9] S. McCabe and J. Scott, "Technique to assess the compatibility of medical implants to the RF field in MRI," *Asia-Pacific Microwave Conference 2015*. Nanjing, 6-9 Dec. 2015.

[10] R. W. P. King and G. S. Smith, "Antennas in matter: fundamentals, theory, and applications", *The MIT Press*, Cambridge, MA, 1981.

[11] S. McCabe and J. Scott, "Cause and amelioration of MRI-induced heating through medical implant lead wires," *Proceedings of the 21st Electronics New Zealand Conference*, Hamilton, 20-21 Nov. 2014, pp.34–40.

[12] *F2182-11a Standard test method for measurement of radio frequency induced heating on or near passive implants during magnetic resonance imaging,* ASTM International, West Conshohocken, PA, USA, 2010.

[13] A. T. Mobashsher and A. M. Abbosh, "Artificial human phantoms: Human proxy in testing microwave apparatuses that have electromagnetic interaction with the human body", IEEE Microwave Magazine, vol.16, no.6, pp.42-62, Jul. 2015.

[14] S. O. McCabe, "Impulse TDR and its application to characterisation of antennas", MPhil thesis, The University of Waikato, 2011.

X-ray free-electron lasers: A revolution in biomolecular imaging

Rick P. Millane, David H. Wojtas and Romain D. Arnal
Computational Imaging Group
Department of Electrical and Computer Engineering
University of Canterbury
Christchurch, New Zealand.
Corresponding email: rick.millane@canterbury.ac.nz

***Abstract*—X-ray crystallography is a technique for imaging molecules at atomic resolution that is based on the diffraction of x-rays by a crystalline specimen. X-ray free-electron lasers are a new x-ray source that can replace conventional synchrotron sources used for x-ray crystallography experiments. They produce extremely bright and brief x-ray pulses, and are revolutionising high resolution imaging of biological macromolecules as they circumvent difficulties with specimen preparation and radiation damage. In these experiments, diffraction data are collected before the specimen is destroyed by the intense x-ray pulse. This requires new experimental protocols and new data processing techniques. Key aspects of biomolecular imaging using x-ray free-electron lasers are reviewed and future opportunities discussed.**

Keywords*— *X-ray free-electron laser; x-ray crystallography; protein crystallography; diffractive imaging; molecular imaging; phase problem; structural biology

I. Introduction

X-ray crystallography is a technique for imaging molecules at atomic resolution. It involves irradiating a crystalline specimen of a molecule with a beam of x-rays, and deducing from the measured pattern of diffracted x-rays the spatial arrangement of the atoms in the molecule (Fig. 1) [1]. The penetrating power of x-rays and their short wavelength (which is comparable to the minimum distance between atoms) allow the three-dimensional structure (i.e. the coordinates of the atoms) of the molecule to be determined. The crystalline state of the specimen serves to amplify the signal in order for useful diffraction measurements to be made. X-ray crystallography is an example of "diffraction imaging," in which an "image" of a scatterer is obtained by measuring its diffraction. There are many other examples of diffraction imaging, including microwave, acoustic, ultrasonic, electron, and neutron imaging, such as in radio astronomy, electron microscopy, sonar, synthetic aperture radar, and medical ultrasonic imaging. X-ray crystallography is over one hundred years old and has been an extraordinarily successful method for elucidating the structures, properties and functionalities of both inorganic and organic (biological) molecules. This has had wide-ranging implications in materials science, chemistry, biology, medicine, and drug design.

The physics of diffraction imaging, as long as the diffraction by the specimen is "weak," is rather easily formulated using Fourier theory. Using an appropriate coordinate system in the measurement space, the measured complex diffracted field is the 3D spatial Fourier transform of the scattering density of the object. Therefore, as long as a sufficiently complete set of measurements is made, the object can be reconstructed by calculating the inverse Fourier transform of the data. In the case of x-ray crystallography, it is necessary to use a crystalline specimen because the diffraction from a single molecule would be far too weak to measure. A crystalline specimen consists of trillions of molecules arranged in a regular manner, such that the diffraction from the individual molecules adds coherently, dramatically amplifying the signal to a measurable level. However, in the case of x-ray crystallography, there are two factors that complicate this simple procedure. First, only the amplitude, but not the phase, of the diffracted x-rays can be measured. Recovery of an object from a measurement of only the amplitude of its Fourier transform is called a "phase problem" [2]. Surprisingly, loss of phase measurements does not preclude reconstruction of the object, so that this, by itself, is not a serious difficulty, as long as the Fourier transform is measured effectively continuously in Fourier (or "reciprocal") space [2]. Secondly, as a result of having a crystalline scatterer, the diffraction occurs only in specific scattering directions, the so-called Bragg reflections, and the result is that the Fourier transform is measured only at specific points in Fourier space. This means that the Fourier amplitude is not measured continuously in Fourier space, or is undersampled, and the above uniqueness properties of the phase problem do not apply [2]. Therefore, measurement of the diffracted amplitude from a crystalline specimen is insufficient to uniquely reconstruct the object. In other words, loss of the phase and a crystalline specimen together conspire to produce an indeterminate solution, whereas one alone would allow a unique solution.

Various imaginative methods for solving the phase problem in x-ray crystallography have been developed over the last 70 years. For small molecules (i.e. containing a small number of atoms) and with high resolution x-ray data, special conditions exist that, despite the nonuniqueness mentioned above, allow

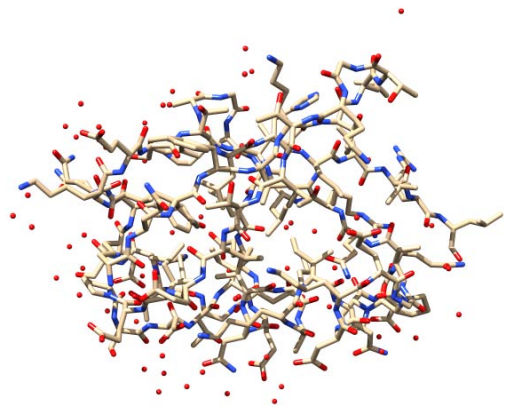

Fig. 1: The atomic structure of the protein ubiquitin [5]. Atoms are shown as balls and bonds as sticks. The oxygen atoms depicted in red are protein-bound water molecules.

the phase problem to be solved from the diffracted amplitudes alone. These are aptly called "direct methods," but they are effective only for small molecules [3]. For many molecules of current interest, particularly biological macromolecules, direct methods are ineffective. In macromolecular crystallography, experimental methods that rely on collecting additional diffraction data from either modified crystals or using different x-ray wavelengths have traditionally been used to solve the phase problem [1]. As the number of molecular structures determined has increased, it has also been possible to use these solved structures in a boot-strap procedure, called "molecular replacement," to solve the phase problem for unknown molecules that are structurally similar to a solved structure. In fact, molecular replacement is now the method of choice for solving the phase problem in macromolecular crystallography and represents over $60\,\%$ of the structure solutions in recent years [4]. However, since this method depends on knowledge of the structure of a similar molecule, it is not suitable for determining the structures of molecules that adopt significantly "new" structures. The phase problem therefore still presents a significant roadblock in macromolecular crystallography.

Despite its enormous success, particularly in structural biology where over 100,000 protein structures have been solved and recorded in the protein databank [6], x-ray crystallography still presents some difficulties. Aside from the phase problem, there are also experimental difficulties, chief among these being crystal growth and radiation damage. With modern synchrotron x-ray sources and x-ray detectors, and robotic sample handing, successful structure determination still requires the preparation of high quality crystals of at least 10 microns in size in order to collect suitable diffraction data. Preparation of such crystals is not necessarily possible for many protein molecules or complexes of interest. Furthermore, small crystals require an intense incident x-ray beam in order to produce measurable diffraction at large scattering angles (i.e. at high resolution). There is a delicate balancing act however, since the ionizing effect of the x-rays means that increasing the x-ray intensity increases the chances of alterations in the atomic structure, thus distorting the very quantity that one is trying to image. The result is that the resolution of the image of the molecule can be limited by radiation damage to a degree that makes the image uninformative. These restrictions can be severe in some cases and have limited the application of x-ray crystallography. For example, membrane-bound proteins, despite their importance as drug targets, are difficult to crystallize and are radiation-sensitive.

II. X-RAY FREE-ELECTRON LASERS

For the last 30 years, the primary source of x-rays for macromolecular crystallography has been synchrotrons. Synchrotrons move electrons in a circular storage ring, typically hundreds of meters in diameter, and the accelerating charged particles emit x-ray radiation (Fig. 2a). Third generation storage rings are dedicated accelerators for generating x-rays that use insertion devices to increase the brilliance of the x-rays produced. Recently, a new generation of x-ray source has been developed that promises to overcome many of the difficulties of x-ray crystallography. These sources are called x-ray free-electron lasers (XFELs) and are revolutionising the field of x-ray crystallography, and x-ray diffractive imaging in general. The main characteristics of XFELs are extraordinary high brightness – 6 to 8 orders of magnitude brighter than third generation x-ray synchrotron sources and extremely brief pulses – about 3 orders of magnitude briefer than those of synchrotron sources. These characteristics help to overcome many of the limitations of conventional protein x-ray crystallography, as is explained below.

In an XFEL, electrons are accelerated to close to the speed of light in a linear accelerator (linac) and passed through an undulator – a series of permanent magnets that alternate in polarity. As a result of the Lorentz force, the electrons pass though the undulator in a sinusoidal path. The relativistic electrons see a contracted undulator spacing and with an electron energy of the order of $10\,\mathrm{GeV}$, the accelerating electrons emit x-rays with a wavelength of the order of $1\,\text{Å}$. Coherence and amplification of the emitted x-ray radiation are achieved through a process known as microbunching. This is a result of an interaction between the electrons and the field of the emitted x-rays. This interaction produces a second longitudinal Lorentz force that results in the electrons evolving into a series of narrow, planar sheets, or microbunches, with spacing equal to the emitted x-ray wavelength. The emissions from the microbunches are in phase and add coherently to give an intense radiation output of very high spatial coherence.

XFELs are large and expensive facilities. "Soft x-ray," or low energy, or long wavelength ($\sim$100 Å) XFELs have been constructed in Hamburg in Germany, Trieste in Italy, and Daresbury in the UK. XFELs suitable for high-resolution molecular imaging are "hard x-ray" sources with wavelengths

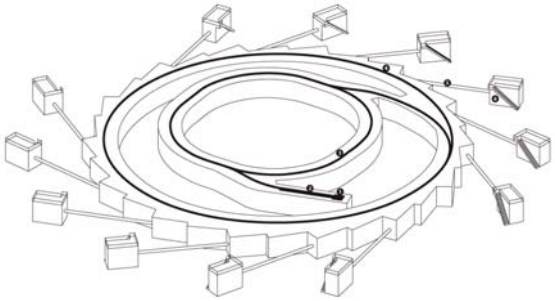

(a)

(b)

Fig. 2: (a) A synchrotron based on a circular storage ring. (b) The LCLS in Stanford, California.

of $\sim$1 Å. The first hard x-ray XFEL was the Linac Coherent Light Source (LCLS) at the SLAC National Accelerator Laboratory in Stanford, California in the USA, which began operation in 2009 (Fig. 2b). The second operational hard x-ray XFEL is the Spring-8 Angstrom Compact Free Electron Laser (SACLA) at the Spring-8 facility in Harama in Japan, which came online in 2012. Other hard x-ray XFELs are under construction or in the planning stages. The European XFEL at DESY in Hamburg, Germany, is due to begin operation in 2017. Other notable hard XFELs nearing completion are the PAL-XFEL in South Korea and the SwissFEL in Switzerland. Hard x-ray XFELs are between 1 and 3 km in length with price tags of the order of US\$0.5 $-$1.5 B.

The unique properties of XFELs have required development of a dedicated class of x-ray detectors to suitably capture the results of diffraction experiments. Detectors in current use are based on a p-n junction CCD (pnCCD) and are made up of ASIC modules. The commonly used CSPAD detector is made up of 64 ASICs with a total of 2.3 Mpixels [7]. The ASIC modules are arranged in 4 quadrants with a central gap to allow

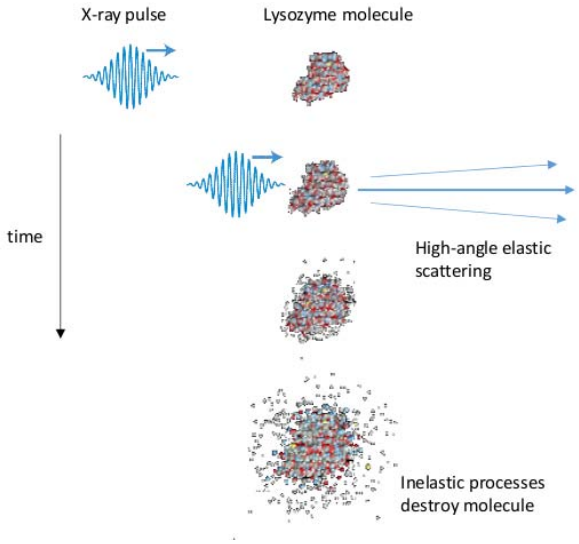

Fig. 3: The diffraction-before-destruction paradigm.

the intense, unscattered x-ray photons to pass through. The detector readout time must exceed the x-ray pulse repetition rate with allowance for component reset. New detectors are being developed with improved linearity, gain homogeneity, dynamic range, and readout times to accommodate the higher pulse repetition rates of upcoming XFELs [8].

III. Crystallography with XFELs

The high brightness and short pulse-length of XFEL sources offer the potential to overcome two of the primary limitations of protein x-ray crystallography using synchrotron sources. First, the high incident x-ray brightness means that a stronger diffracted signal is produced for a given crystal size, and thus measurable diffraction can be obtained from crystals of smaller size. The result is that diffraction signals can be measured from so-called nanocrystals, with dimensions of the order of 100 nm, or two orders of magnitude smaller in linear dimensions than the smallest crystals required for synchrotron sources. The advantage here is that molecules for which it is difficult, or impossible, to prepare micron size crystals, can often be easily induced to form nanocrystals. Hence, this feature in many cases can overcome difficulties associated with crystal preparation, opening up x-ray imaging to molecules that have otherwise been inaccessible.

Second, the short pulse duration of XFELs can, somewhat paradoxically, address problems associated with radiation damage. For every elastically scattered x-ray photon that contributes to the desired, measured diffraction, there may be ten or more inelastically scattered or absorbed photons that deposit energy in the specimen. The high brightness of the XFEL pulse means that a large amount of energy is deposited. The inelastically scattered or absorbed photons eject photoelectrons from the atoms, which is followed by a cascade

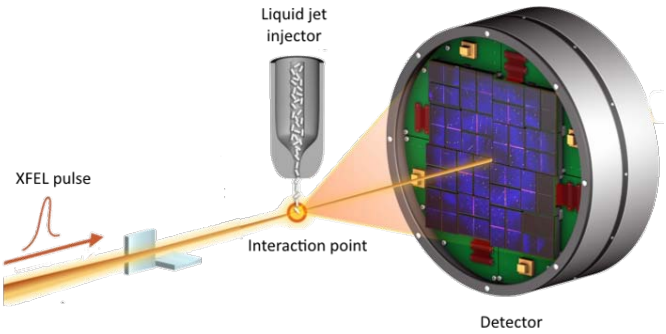

Fig. 4: The serial femtosecond nano-crystallography setup [10].

of lower energy electrons due to secondary impacts. An important effect is that removal of the electrons from the atoms leaves behind positively charged ions that repel each other, leading to destruction of the sample in a "Coulomb explosion." However, the duration of a typical XFEL pulse is of the order of $100\,\mathrm{fs}$, which is shorter than the time taken for initiation of the Coulomb explosion. Thus, although the intense x-ray pulse of an XFEL destroys the specimen, the diffracted pulse encodes information on the pristine specimen, or molecule, on a length scale that is smaller than the resolution required. Therefore, by recording the diffracted signal, an image can be obtained of the intact molecule (Fig. 3). This strategy is known as "diffraction-before-destruction," and was a key early recognition of the potential for imaging with XFELs [9].

In summary then, XFELs potentially address the difficulties of crystal preparation and radiation damage that plague conventional protein x-ray crystallography. They also have the potential to allow for exotic solutions to the phase problem, which is described further below.

In a typical XFEL experiment, x-ray pulses are focused and delivered to the specimen in an evacuated chamber. Typical focal spot sizes are of the order of one micron across. In the case of the LCLS, the x-ray pulses are delivered with a maximum pulse repetition rate of $120\,\mathrm{Hz}$. It is then necessary to deliver the specimen, often in the form of nanocrystals, into the focal spot so that they interact with the x-ray pulses. Great ingenuity has been required in developing suitable specimen, or particle, delivery techniques. Since most biological samples require an aqueous environment, the most common delivery technique has been a fine water jet that carries hydrated nanocrystals. The most common injector for proteins is a liquid microjet [11]. The sample is pumped through a narrow glass capillary which is surrounded by another capillary of larger radius. A high pressure gas, typically helium, is pumped through the sheath between the two capillaries, and at the exit of the injector the gas flow focuses and stabilises the sample stream. The diameter of the focused liquid stream is typically a few microns and typical flow rates are of the order of $10\,\mathrm{m\,s^{-1}}$. This is known as a gas dynamic virtual nozzle (GDVN). The coaxial gas flow surrounding the liquid jet also prevents sample freezing in the vacuum chamber. Other sample injection techniques include aerosol injectors that are suitable when sample hydration is not critical, and lipid cubic phase (LCP) injectors that use a viscous gel-like medium suitable for membrane proteins. So-called fixed-targets are also being developed, in which protein nanocrystals are deposited on an ultra-thin, silicon nitride membrane [12]. An array of such membranes is positioned and scanned through the x-ray beam. Advantages of fixed targets are that background scattering from the water jet is eliminated and scanning of the membrane can be synchronised with the x-ray pulses.

The nanocrystals in an XFEL experiment are delivered to the x-ray interaction region in a serial fashion and intercept the x-ray pulses at random, and in random, unknown orientations (Fig. 4). Structure determination using this approach is called serial femtosecond nanocrystallography (SFX) [13], and is in contrast to conventional crystallography in which crystals are mounted on a goniometer which is used to control their orientation [1]. SFX requires new processing methods to deal with this kind of data. The key steps to reduce the diffraction data to estimates of the intensities of the Bragg reflections are hit-finding, indexing, and merging/integration.

A typical diffraction experiment, which may run for up to 5 days, can generate up to 10^7 diffraction patterns, occupying tens to hundreds of TB of storage. Many of the XFEL pulses will not intercept a crystal, while others may intercept two or more crystals. Some diffraction patterns may have excessive background diffraction due to the water jet or other spurious diffraction. Given the volume of data, fast methods are needed to extract the potentially useful diffraction patterns, referred to as "hits." The initial step of identifying hits is referred to as "hit-finding." Typical hit-finding involves defining masks on the detector to avoid regions where spurious diffraction is likely to occur, estimation and subtraction of background diffraction, and using thresholds on diffraction intensity and feature sizes to identify likely Bragg peaks. This information is then used to identify and extract diffraction patterns that are likely to be useful. A popular program for conducting this processing is Cheetah [14]. The x-ray focal spot size, jet diameter, particle concentration and pulse repetition rate together determine the likelihood that an x-ray pulse intersects a particle(s) and the rate at which diffraction patterns are produced. The fraction of x-ray pulses that result in a hit is referred to as the "hit rate." Typical hit rates with optimised conditions may be of the order of $5-10\,\%$. A high hit rate is desirable since this optimises throughput, and minimises the time and cost of use of the facility.

The next step is to extract from the hits, those patterns that contain useable information. In the case of nanocrystal specimens, the diffraction patterns contain sharp Bragg reflections and

these are used to identify suitable patterns that originate from a single nanocrystal. An example of such a diffraction pattern is shown in Fig. 5a. This process is referred to as "indexing," and it jointly determines good quality diffraction patterns that are due to one nanocrystal, as well as the orientation of the crystal in the incident x-ray beam. The fraction of hits that can be indexed is referred to as the "indexing rate." Typically, indexing rates are in the range $20-50\,\%$. Therefore, typically $1-5\,\%$ of the original diffraction patterns provide useful data.

For a particular orientation of a nanocrystal, the measured diffraction pattern represents a particular two-dimensional section through the three-dimensional Fourier transform of the particle. The orientation of the nanocrystal determined by the indexing allows the position of this 2D section to be determined. The objective is to use these sections from many crystals in different orientations to "fill up" Fourier space. Since the diffracting particles (nanocrystals) are small, the diffraction patterns are weak and the signal-to-noise is low. Therefore, patterns from many crystals in similar orientations are also needed to average out the effects of noise. The process of assembling all the 2D diffraction patterns to produce a 3D dataset is referred to as "merging" of the data. For typical protein nanocrystals, of the order of 20,000 indexed diffraction patterns are needed to merge the data, which requires the measurement of 200,000 to 500,000 patterns in the experiment. A representation of a 3D merged dataset is shown in Fig. 5b. A popular program used to conduct indexing and merging is CrystFEL [15].

A difficulty in an XFEL experiment is that there are various parameters, aside from the crystal orientation, that are unknown. The dimensions, and therefore the scattering strength, of an individual nanocrystal are unknown. The intensity of the incident x-ray pulse in the self-amplifying process described above for generation of the x-ray pulse (referred to as self-amplified stimulated emission, or SASE) is a random process and varies from pulse to pulse (or from "shot-to-shot") in an unknown manner. The relative scale of each recorded diffraction pattern is therefore unknown. Also, the degree to which the intensity of each Bragg reflection is recorded on the detector depends critically on parameters such as the precise orientation of the crystal, the temporal bandwidth of the x-ray pulse, and the overall dimensions of the nanocrystal, all of which are also unknown from shot to shot. This is referred to as the "partiality" of each Bragg reflection and results in a variable scale between Bragg reflections within each diffraction pattern. Fortunately, these parameters sample their distributions during an experiment, and because a large number of patterns are available, these random effects can be averaged out in a process called "Monte Carlo averaging" developed by [16]. The procedure involves, for each diffraction pattern, integrating the intensity in a small region around each Bragg reflection and then averaging over the patterns. The result is an unbiased estimate of the intensity of each Bragg reflection.

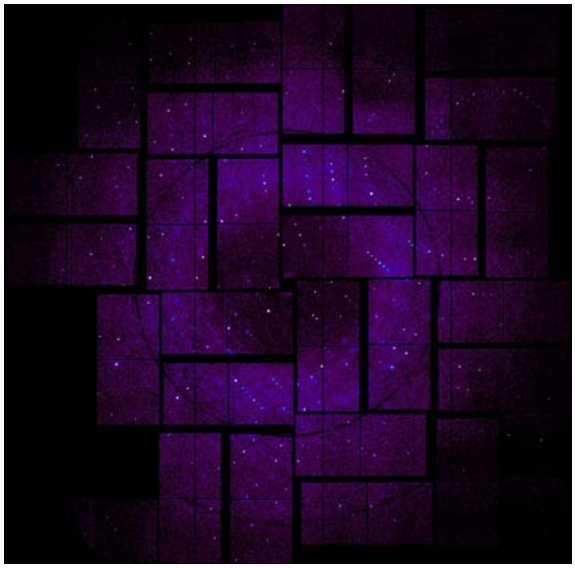

(a)

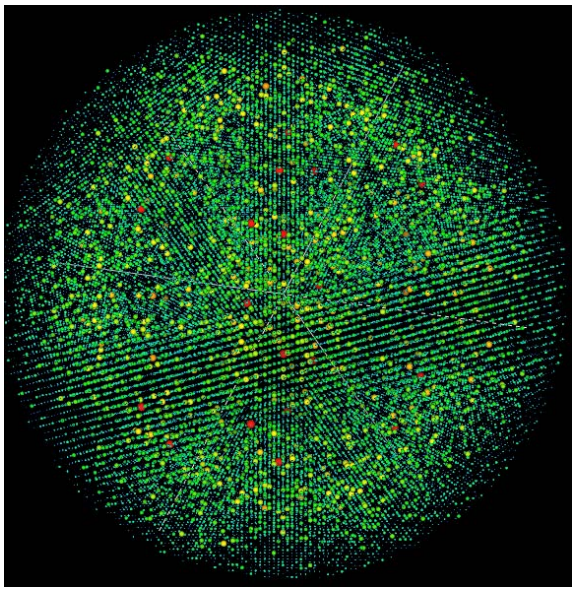

(b)

Fig. 5: (a) An example single XFEL diffraction pattern from a lysozyme nanocrystal (Image: Anton Barty/DESY). (b) A representation of the Bragg intensities in 3D Fourier space. A projection of 3D Fourier space is shown. Each point shown represents a Bragg sample in Fourier space. The size and colour of the blob at each point represents the amplitude of the measured Bragg reflection at that point.

The intensities of the Bragg reflections thus obtained mimic the data that are obtained in a conventional crystallography experiment using a synchrotron. From this point on, conventional crystallographic processing and analysis is conducted [1], often using the molecular replacement method for phasing, to determine the molecular structure. SFX as described has proved successful with now over 30 structures determined and

deposited in the Protein Data Bank (e.g. [17], [18]).

IV. Future prospects

XFELs and their applications are in their infancy. They have many applications in science and technology that go beyond the focus in this article, which has been on imaging biological macromolecules using serial nanocrystallography. In the context of bioimaging, four prospects worth mentioning are upcoming XFELs, new phasing methods, imaging molecular dynamics, and single particle imaging. New XFEL sources such as the European XFEL in Hamburg will have much higher pulse repetition rates, up to 27 kHz, which offer new opportunities in the study of dynamic processes, and in throughput, but also challenges in detector design and data processing. XFELs offer opportunities for new approaches to phasing in crystallography, which could dramatically change the field by avoiding the restrictions of, in particular, the molecular replacement method. New phasing methods being explored are based on using information resulting from radiation damage [19], [20], using diffraction between the Bragg reflections that is observed using very small nanocrystals [21], [22], and using diffuse diffraction that results from small amounts of lattice disorder that is often present in protein crystals [23]. XFEL imaging is ideally suited to studying dynamic processes by controlled manipulation of specimens and reactions, synchronised with the XFEL pulses. Dynamic imaging has already started to bear fruit [24], [25]. Finally, the ultimate goal of molecular imaging is imaging without crystals: so-called single particle imaging (SPI). Even more intense XFEL sources, coupled with the diffract-before-destroy paradigm, offer prospects for SPI, although problems with particle location, background scattering and very low signal levels make this very challenging. Without the sharp diffraction features from crystals, for example, detection and orientation of diffraction patterns as described above is very difficult. Progress, however, is being made with SPI [26].

Acknowledgement

This work was supported by a New Zealand Marsden Grant.

References

[1] J. Drenth, “Principles of protein X-ray crystallography,” Springer Science & Business Media, 2007.

[2] R. P. Millane,“Phase retrieval in crystallography and optics,” J. Opt. Soc. Am. A, vol. 7, pp. 394–411, 1990.

[3] M. M. Woolfson and F. Hai-fu, “Physical and non-physical methods of solving crystal structures,” Cambridge University Press, 1995.

[4] G. Scapin, “Molecular replacement then and now,” Acta Crystallogr. Sect. D, vol. 69, pp. 2266–2275, 2013.

[5] S. Gao *et al.*, “Monomer/Oligomer Quasi-Racemic Protein Crystallography,” J. Am. Chem. Soc., vol. 138, pp. 14497-14502, 2016.

[6] “PDB Current Holdings Breakdown,” 2016. [Online]. Available: http://www.rcsb.org/pdb/statistics/holdings.do [Accessed: 2016-04-11].

[7] S. Herrmann *et al.*, “CSPAD-140k: a versatile detector for LCLS experiments,” Nucl. Instrum. Methods Phys Res. A, vol. 718, pp. 550-553, 2013.

[8] P. Denes, “Two-dimensional imaging detectors for structural biology with X-ray lasers,” Phil. Trans. R. Soc. B, vol. 369, doi: 10.1098/rstb.2013.0334., 2014.

[9] R. Neutze, R. Wouts, D. van der Spoel, E. Weckert and J. Hajdu, “Potential for biomolecular imaging with femtosecond X-ray pulses,” Nature, vol. 406, pp. 752–757, 2000.

[10] S. Boutet *et al.*, “High-Resolution Protein Structure Determination by Serial Femtosecond Crystallography, Science, vol. 337, pp. 362–364, 2012.

[11] U. Weierstall, “Liquid sample delivery techniques for serial femtosecond crystallography,” Philos. Trans. R. Soc. Lond. B Biol. Sci., vol. 369, doi:10.1098/rstb.2013.0337, 2014.

[12] M. S. Hunter *et al.*, “Fixed-target protein serial microcrystallography with an x-ray free electron laser,” Sci Rep., vol. 4, doi: 10.1038/srep0602, 2014.

[13] H. N. Chapman *et al.*, “Femtosecond X-ray protein nanocrystallography,” Nature, vol. 470, pp. 73–77, 2011.

[14] A. Barty, R. A. Kirian, N. C. Maia, M. Hantke, C. H. Yoon, T. A. White and H. N. Chapman, “Cheetah: software for high-throughput reduction and analysis of serial femtosecond X-ray diffraction data,” J. Appl. Cryst., vol. 47, pp. 1118–1131., 2014.

[15] T. A. White, V. Mariani, W. Brehm, O. Yefanov, A. Barty, K. R. Beyerlein, F. Chervinskii, L. Galli, C. Gati, T. Nakane, A. Tolstikova, K. Yamashita, C. H. Yoon, K. Diederichs and H. N. Chapman “Recent developments in CrystFEL,” J. Appl. Cryst., vol. 49, pp. 680–689, 2016.

[16] R. A. Kirian, X. Wang, U. Weierstall, K. E. Schmidt, J. C. Spence, M. Hunter, P. Fromme, T. White, H. N. Chapman and J. Holton, “Femtosecond protein nanocrystallography-data analysis methods,” Opt. Express, vol. 18, pp. 5713–5723, 2010.

[17] L. Redecke *et al.*, “Natively inhibited Trypanosoma brucei Cathespin B structure determined by using an x-ray laser,” Science, vol. 339, pp. 227-230, 2012.

[18] Y. Kang *et al.*, “Crystal structure of rhodopsin bound to arrestin by femtosecond x-ray laser,” Nature, vol. 523, pp. 561-567, 2015.

[19] S. Banumathi, P. H. Zwart, U. A. Ramagopal, M. Dauter and Z. Dauter “Structural effects of radiation damage and its potential for phasing,” Acta Crystallogr. Sect. D, vol 60(6), 1085-1093.

[20] R. B. Ravelli, H. K. S. Leiros, B. Pan, M. Caffrey and S. McSweeney, “Specific radiation damage can be used to solve macromolecular crystal structures,” Structure, vol. 11, pp. 217–224, 2003.

[21] J. C. Spence *et al.*, “Phasing of coherent femtosecond X-ray diffraction from size-varying nanocrystals,” Optics Express, vol. 19, pp. 2866–2873, 2011.

[22] R. P. Millane and J. P. Chen, “Aspects of direct phasing in femtosecond nanocrystallography,” Phil. Trans. R. Soc. B, vol. 369, 20130498, 2014.

[23] K. Ayyer *et al.*, “Macromolecular diffractive imaging using imperfect crystals,” Nature, vol. 530, pp. 202–206, 2016.

[24] C. Kupitz *et al.*, “Serial time-resolved crystallography of photosystem II using a femtosecond laser,” Nature, vol. 513, pp. 261–265, 2014.

[25] J. Tenboer *et al.*, “Time-resolved serial crystallography captures high-resolution intermediates of photoactive yellow protein,” Science, vol. 346, pp. 1242–1246, 2014.

[26] T. Ekeberg *et al.*, “Three-dimensional reconstruction of the giant mimivirus particle with an x-ray free-electron laser,” Phys. Rev. Lett., vol 114, 098102, 2015.

SkyPi: All-sky camera network for meteorite tracking

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

***Abstract*—We describe the hardware and software for SkyPi – an all-sky camera network for transient event detection in the night sky. The hardware design is presented as well as the operating software and some preliminary results.**

I. Overview

This paper describes the hardware and software for a network of all-sky cameras for meteorite tracking. These cameras will be deployed at rural schools across central otago, and the images will be used to provide science activities for school pupils. When bright meteorites are located in multiple cameras, triangulation will be used to determine where such meteorites could have landed, and meteorite hunts will be arranged!

Meteorites have been successfully recovered [3] from similar camera networks, include the Desert Fireball Network [1] over the Nullabor Plain in central Australia. This network consists of 32 cameras. A sample image from the DFN is shown in Figure 1.

In addition to measurements on meteorites, the data from these camers will be mined to provide data on sky darkness, as well as statistics about cloud coverage. This information will be used to inform local-government about opportunites for dark-sky tourism, as well as light-pollution.

The following sections describe the hardware, software and we conclude with some preliminary results.

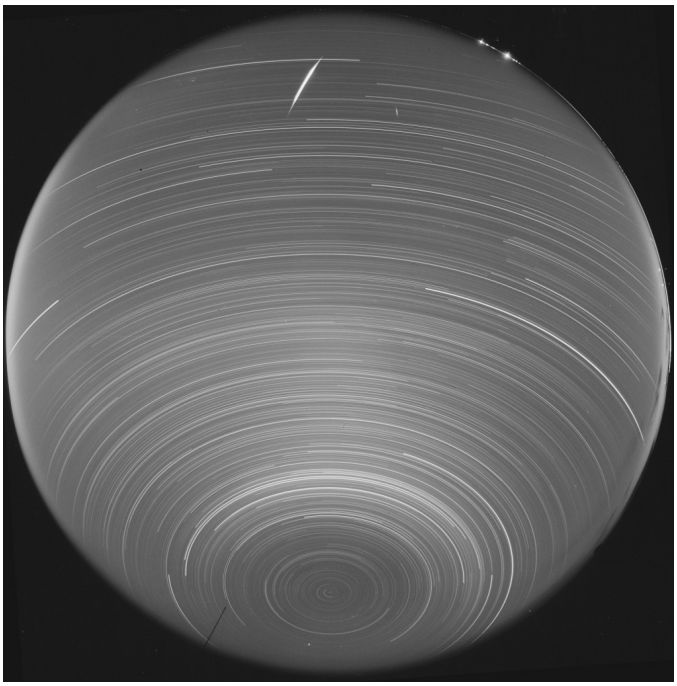

Fig. 1. An image from the Desert Fireball Network, showing the Mason Gully Meteorite.

II. Hardware

The primary focus of the camera hardware is that it be low-cost and low-power so that deployment can be widespread. The main features are that it have sufficient computational power to perform some of the image analysis necessary to reduce the bandwidth required for transmission of images from the camera network to processing center at the University of Otago. A full schematic of the hardware is shown in Figure 2.

A 24V DC power supply was chosen so that the cameras could be situated outdoors, some distance from a power point, and any voltage drop wouldn't be significant. The 24V supply is converted to 5V using a DC-DC switching converter that can operated from a DC supply between 9V and 36V.

A. Host processor

The host processor is a Raspberry Pi [5] 3. It features a quad-core 64-bit ARM-7 CPU with 1GB of RAM, wireless networking and a camera interface. The camera interface is particularly important, as it provides high-bandwidth image data directly to the CPU. The system operates the Raspbian (Debian-derived Linux) operating system.

In normal operation, the processor will be connected to the internet, either by wired ethernet, or using the WiFi built-in to the Raspberry Pi 3. The operating software for the host processor is described in Section III.

B. Camera

The host processor camera interface is a Camera Serial Interface Type 2 (CSI-2), which connects directly to the Broadcom BCM2835 SOC. It provides a bandwidth of up to 4 Gbps between the camera sensor and the host memory.

The camera is CMOS (Omnivision OV5647 []) sensor board camera capable of 2592 x 1944 pixel static images. The lens chosen is an Jiangxi Hongxin Optical Co., Ltd, M12 mount, 180 FOV lens. An API is avilable to control the acquisition of image data from the sensor, including contol of the analog

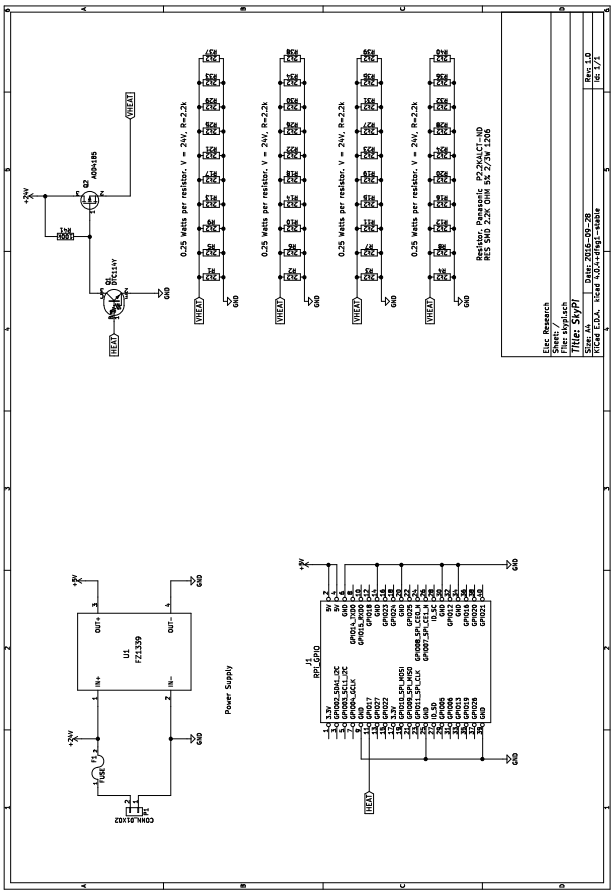

Fig. 2. Schematic Design for the SkyPi hardware

gains in the CMOS sensor readout. Experimentation with different settings yielded an optimal exposure time of two-seconds. Longer than this, and the noise in the sensor can become large (and warm areas of the chip show up in averaged images). Unfortunately cooling of the sensor to reduce noise is not feasible, as dew-heating (described in Section II-C) is required.

C. Dew heater

Early prototype cameras experienced significant problems with condensation on the inside of the camera dome. Several heating options were explored to mitigate this, including PTC resistive heaters.

For simplicity, an array of 40 1206 SMD resistors (each capable of dissapating 0.6 W) is driven directly from the 24V DC supply, these resistors are placed on the top layer of the skypi pcb. The heater supply (VHEAT) is controlled by the host using a digital transistor driving a FET. This control allows the heater to be switched off when not needed. At this stage, there isn't a temperature sensor on the board, and it is expected that the heater will operated while imaging is happening (during hours of darkness).

Fig. 3. Wide-angle (fisheye) lens used for SkyPi. This lens has a field of view exceeding 180 degrees, and is mounted using an M12x0.5 thread.

Fig. 4. First prototype skypi showing resistive heater, and insulated enclosure.

Adding temperature and humidity would help minimize power consumption, and also reduce excessive camera temperature when operating in hot environments.

III. Software

The software is written in the Python programming language, using the opencv image processing toolkit. It consists of image stacking routines to continuously accumulate star-trail images of the night sky. Once images as stacked into ten minute star-trail images, these are saved to disk, and automatically analysed to locate potential meteorite trails. Finally images are uploaded to a central server where they are archived, for meteorite triangulation and post-processing.

To minimize power and data transfer, the position of the sun is simulated, and the camera only begins operation when the sun is more than ten degrees below the horizon. This ensures that the skies are reasonably dark.

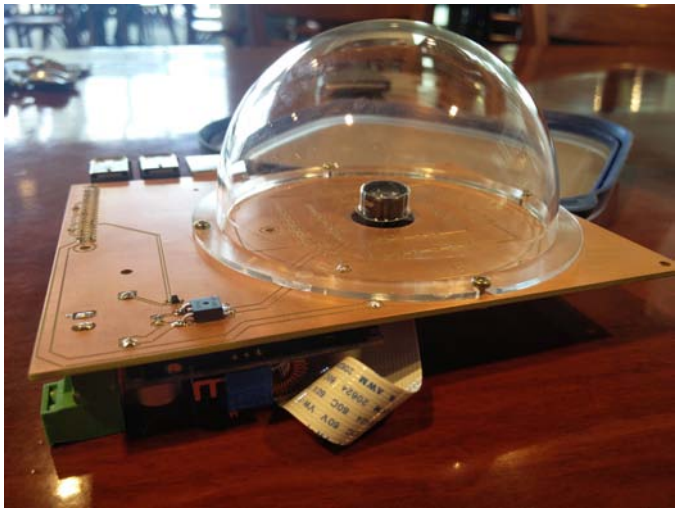

Fig. 5. Current prototype, without resistors installed, but showing all other components mounted on the main printed circuit board.

Fig. 6. Meteorite observed with a SkyPi prototype from Dunedin NZ.

A. Image Stacking

The process is multi-threaded. The image stacking thread is responsible for accumulating images without any breaks – this ensures that no meteors are missed.

The image stacking algorithm accumulates the sum of N two-second exposures, where N is calculated to achieve the desired length of star trail (in minutes). A ten-minute star trail image consists of the arithmentic mean of 300 exposures. Figure 6 shows an example image accumulated with this technique.

B. Meteorite Detection

A separate thread in the operating code processes the stacked images, and attempts to locate meteorites. This thread uses the OpenCV image processing library [2] via it's Python API. The initial algorithm relies on the fact that meteors are closely approximated by straight lines (over angles less than 20 degrees).

The star-trail image is converted from a 32-bit floating point representation, to a greyscale 8-bit image. Then a probablistic Hough line transform [4] is used to locate any line segments in the image longer than the known length of the star trails (about 40 pixels for a ten-minute image). If a line segment is located, then the parameters of the longest line in the image are appended to a file.

C. Data storage

Data (detection data, and star-trail images) are stored locally in the host processors filesystem. At regular intervals (every five minutes) a script is triggered that uploads star trail images to a central repository using the rsync protocol. After successful upload, the local copies of the images are deleted.

Every morning, a script examines the detection data, and locates images where meteors have occurred. This data will be manually processed to examine candidate meteors that have been observed from multiple sites, and to locate possible ground falls. Automation in this area is being explored.

IV. Conclusions

Some improvements are planned for future versions of the SkyPi. For remote operation, solar power would be an option, but minimizing power consumption would require some design changes. These would include a supervisor processor with real-time clock and battery voltage sensing so allow a very low-power shutdown mode during daylight hours. In addition, a temperature and humidity would help minimize power consumption by calculating how much heater power would be necessary to avoid condensation.

Improvements in operating software to help deal with light-polluted urban environments are planned.

The initial network is planned to include seven rural schools from Central Otago. It is envisaged that the network will operate for four months, and during this pilot phase, the design will be refined as necessary. Future expansion of the network beyond central otago is under consideration. This will require a more distributed data storage architecture. And we expect to transition to the use of an Amazon S3 storage container for temporary storage in the cloud, and this data automatically downloaded for processing, or processed using EC2 compute instances. This has the advantage of finer-gained control of access security, as well as scalability to potentially thousands of cameras.

Finally, the motion of meteorites through the atmosphere is complex. Modelling the paths from the observed data is an exciting avenue for further research.

Acknowledgements

This author would like to acknowledge support from the University of Otago, members of staff from the Otago Museum and Planetarium, and by volunteers from the Astronomical Society.

REFERENCES

[1] Philip A Bland. The desert fireball network. *Astronomy & Geophysics*, 45(5):5–20, 2004.

[2] Gary Bradski and Adrian Kaehler. *Learning OpenCV: Computer vision with the OpenCV library*. " O'Reilly Media, Inc.", 2008.

[3] Kathryn A Dyl, Gretchen K Benedix, Phil A Bland, Jon M Friedrich, Pavel Spurnỳ, Martin C Towner, Mary Claire O'Keefe, Kieren Howard, Richard Greenwood, Robert J Macke, et al. Characterization of mason gully (h5): The second recovered fall from the desert fireball network. *Meteoritics & Planetary Science*, 51(3):596–613, 2016.

[4] Nahum Kiryati, Yuval Eldar, and Alfred M Bruckstein. A probabilistic hough transform. *Pattern recognition*, 24(4):303–316, 1991.

[5] Eben Upton and Gareth Halfacree. *Raspberry Pi user guide*. John Wiley & Sons, 2014.

Vector Green's Functions for Electrodynamics Applications

Malcolm Morrison, Colin Fox
Electronics Group
University of Otago
Email: morma584@student.otago.ac.nz

***Abstract*—The use of scalar Green's functions is commonplace in electrodynamics, but many useful systems require computation of one or more vector quantities. Vector Green's functions have been used in electrodynamics for many decades, but are inconsistent and poorly understood. We recast vector Green's functions in the language of distribution theory to match their scalar counterparts. We use vector Green's functions to calculate the electrodynamic vector potential under physical boundary conditions.**

I. Introduction

Green's functions have simplified the solving of inhomogeneous, linear, scalar boundary value problems (BVPs) which are common in many fields of study, e.g. quantum physics [1], many body simulations [2], and electrodynamics [3]. If one can find the Green's function for a BVP, then one can construct the solution to an arbitrary region through use of the boundary element method [4].

While Green's functions are very old, it was only the development of distribution theory and the notion of a generalised function that allowed a full understanding of the uses and behaviour of Green's functions [5]. Modern distribution theory deals with scalar functions and functionals. Despite the development of the theory of vector distributions shortly thereafter [6], it passed largely unnoticed.

Vector Green's functions have been used for the better part of a century, but their use has not generally been rigorously justified. However the selections of the functions are inconsistent [7] [8]. This can be put on a secure footing by a connection to distribution theory and generalised functions such as that which allowed scalar Green's functions to become so ubiquitous. Moreover, vector Green's functions are often constructed such that they satisfy radiation conditions in order to guarantee uniqueness of the solutions to the BVPs. While this does yield useful solutions, radiation conditions are just a small fraction of possible physically relevant boundary conditions. Thus, only developing theory for radiation conditions at the boundary limits possible applications.

In this paper we connect the existing vector distribution theory to create a general theory for vector Green's functions. We then develop vector Green's functions for specific cases for the electrodynamic vector potential with simple examples, finding analogies to common scalar boundary conditions in order to guarantee the uniqueness of the electric and magnetic fields.

II. Vector Green's Functions

A. Background

The following are some useful definitions from Ref [6].

Definition II.1. A *test function* is an infinitely differentiable function with compact support. This is also often referred to as a *bump function*.

Definition II.2. A *functional*, $< f, \cdot >$, is defined such that for each function $\phi(x)$, it assigns a real number. We will use functionals of the form:

$$< f, \phi >_{\Omega} \stackrel{def}{=} \int_{\Omega} f(\mathbf{x})\phi^*(\mathbf{x})dV, \tag{1}$$

for a given function $f(\mathbf{x})$ in the region Ω with the star denoting the complex conjugate.

Definition II.3. A *distribution* is a continuous linear functional on the space of test functions. The distribution generated by the function $f(\mathbf{x})$ is denoted f.

Now let $f(\mathbf{x})$ be differentiable and $f^{(i)}(\mathbf{x})$ locally integrable, using the superscript (i) to denote a single differentiation in the x_i variable. The distribution generated by $f^{(i)}(\mathbf{x})$, found through integration by parts is:

$$< f^{(i)}, \phi >_{\Omega} = - < f, \phi^{(i)} >_{\Omega} \tag{2}$$

Definition II.4. A *test vector* is defined with test functions $t_1(\mathbf{x})$, ..., $t_n(\mathbf{x})$ as

$$\mathbf{T}(\mathbf{x}) = \hat{e}_1 t_1(\mathbf{x}) + ... + \hat{e}_n t_n(\mathbf{x}), \tag{3}$$

where $\hat{e}_i$ is the ith basis unit vector.

Definition II.5. A vector $\mathbf{S}$ whose components are distributions s_1, ..., s_n is called a *vector distribution*, i.e.,

$$\mathbf{S} = \hat{e}_1 s_1 + ... + \hat{e}_n s_n. \tag{4}$$

We can now extend our functionals to deal with vectors by:

$$\begin{aligned} < \mathbf{S}, \mathbf{T} > &= < s_1, t_1 > + < s_2, t_2 > + < s_3, t_3 > \\ &= \int \mathbf{S} \cdot \mathbf{T}^* dV \end{aligned} \tag{5}$$

The identities for the gradient, divergence, and curl follow from equation (2):

$$< \nabla s, \mathbf{T} > = - < s, \nabla \cdot \mathbf{T} >, \tag{6}$$

$$< \nabla \cdot \mathbf{S}, t > = - < \mathbf{S}, \nabla t >, \tag{7}$$

$$< \nabla \times \mathbf{S}, \mathbf{T} > = - < \mathbf{S}, \nabla \times \mathbf{T} > . \tag{8}$$

B. Scalar Green's functions

The following is a brief summary of scalar Green's functions [5].

Any linear BVP in a region Ω can be expressed as

$$Lu(\mathbf{x}) = f(\mathbf{x}), \quad B_{(i)}(u) = c_i \tag{9}$$

where L is a linear partial differential operator, $f(\mathbf{x})$ a function describing the inhomogeneity, and $B_{(i)} = c_j$ is the j'th boundary condition (BC) over the boundary $\partial\Omega$.

Definition II.6. The *Green's function*, $g(\mathbf{x}|\boldsymbol{\xi})$, is the solution with compact support to the related homogeneous partial differential equation (PDE),

$$Lg(\mathbf{x}|\boldsymbol{\xi}) = \delta(\mathbf{x} - \boldsymbol{\xi}), \quad B_{(i)}(u) = 0, \tag{10}$$

where $\boldsymbol{\xi} = (x_1', ..., x_n')$.

If the boundary is at infinity, then the solution to the BVP is simply

$$u(\mathbf{x}) = \int_{\Omega} g^*(\mathbf{x}|\boldsymbol{\xi}) f(\boldsymbol{\xi}) d^n\boldsymbol{\xi} \tag{11}$$

due to the compact support condition, where n is the number of dimensions. Otherwise we must look at the adjoint problem:

$$L^* w(\mathbf{x}|\boldsymbol{\xi}) = \delta(\mathbf{x} - \boldsymbol{\xi}), \quad B^*_{(i)}(w) = 0, \tag{12}$$

where $w(\mathbf{x}|\boldsymbol{\xi})$ is the adjoint Green's function and $B^*_{(i)}$ are the adjoint BCs.

Then, substituting (12) and (9) into (1) we integrate by parts to find

$$< Lu, w >_{\Omega} = < u, L^* w >_{\Omega} + \Big[J(u, w) \Big]_{\Omega}, \tag{13}$$

where J is referred to as the bilinear concommitant. Since $L^* w(\mathbf{x}|\boldsymbol{\xi}) = \delta(\mathbf{x} - \boldsymbol{\xi})$ and $Lu(\mathbf{x}) = f(\mathbf{x})$, we find that

$$u(\mathbf{x}) = \int_{\Omega} f(\boldsymbol{\xi}) w^*(\mathbf{x}|\boldsymbol{\xi}) d^n\boldsymbol{\xi} + \Big[J(u, w) \Big]_{\Omega} \tag{14}$$

We can choose the adjoint boundary conditions however we like, so it makes sense to choose them such that $\Big[J(u, w) \Big]_{\Omega}$ is as simple as possible. Ideally this can be made to vanish.

1) Electrodynamics Application: These techniques have been very useful in solving Maxwell's equations for the scalar potential, Φ. Maxwell's equations for a static charge density, $\rho(\mathbf{x})$, in free space with no magnetic fields can be rewritten as Poisson's equation

$$\nabla^2 \Phi(\mathbf{x}) = -\epsilon_0^{-1} \rho(\mathbf{x}). \tag{15}$$

In free space the boundary conditions vanish and we can find the solution to this through equation (11) as the system is self adjoint. The Green's function is defined by

$$\nabla^2 g(\mathbf{x}|\boldsymbol{\xi}) = \delta(\mathbf{x} - \boldsymbol{\xi}), \tag{16}$$

which can be solved in 3-D to yield:

$$g(\mathbf{x}|\boldsymbol{\xi}) = -\frac{1}{4\pi|\mathbf{x} - \boldsymbol{\xi}|}, \tag{17}$$

where $|\mathbf{x} - \boldsymbol{\xi}| = ((x_1 - x_1')^2 + (x_2 - x_2')^2 + (x_3 - x_3')^2)^{\frac{1}{2}}$. This yields the well known electrostatic potential for a point charge

$$\Phi(\mathbf{x}) = \frac{1}{4\pi\epsilon_0} \int \frac{\rho(\boldsymbol{\xi})}{|\mathbf{x} - \boldsymbol{\xi}|} d^3\boldsymbol{\xi}. \tag{18}$$

The free space Green's function can be used for problems where the boundaries do not extend to infinity by finding uniqueness conditions based on the boundary of a region. This allows some complex boundaries and charge distributions to be solved for through the method of images and boundary element calculations [4].

2) Vector Green's function: A vector BVP over the region Ω can be described in a similar way to the scalar problem. In general they will take the form:

$$\mathbf{L}\mathbf{u}(\mathbf{x}) = \mathbf{f}(\mathbf{x}), \quad \mathbf{b}_{(i)}(\mathbf{u}) = \mathbf{c}_i \tag{19}$$

for some vector linear differential operator $\mathbf{L}$ and boundary conditions $\mathbf{b}_{(i)}$ on the boundary $\partial\Omega$. We start our development of vector Green's functions by defining its adjoint, $\mathbf{W}_1(\mathbf{x}|\boldsymbol{\xi})$, as the solution to the homogeneous adjoint BVP:

$$\mathbf{L}^* \mathbf{W}_1(\mathbf{x}|\boldsymbol{\xi})(\mathbf{x}) = \delta(\mathbf{x} - \boldsymbol{\xi}) \hat{e}_1, \quad \mathbf{b}^*_{(i)}(\mathbf{W}_1) = 0. \tag{20}$$

Using our vector distribution theory it can be shown that

$$< \mathbf{L}\mathbf{u}, \mathbf{W}_1 >_{\Omega} = < \mathbf{u}, \mathbf{L}^* \mathbf{W}_1 >_{\Omega} + \oint_{\partial\Omega} \mathbf{C}_1 \cdot \hat{\mathbf{n}} \, ds \tag{21}$$

Where $\mathbf{L}^*$ is the adjoint linear operator and $\mathbf{C}_1$ is the vector bilinear concommitant. By using (19) and (20), this can be rewritten as

$$\mathbf{u} \cdot \hat{e}_1 = -\int_{\Omega} \mathbf{W}_1^* \cdot \mathbf{f} \, d^n\boldsymbol{\xi} - \oint_{\partial\Omega} \mathbf{C}_1 \cdot \hat{\mathbf{n}} \, ds. \tag{22}$$

Evidently, this Green's function yields the $\hat{e}_1$ component of the solution. The same process can be done for all basis vectors, and hence the full solution can be expressed as:

$$\begin{aligned} \mathbf{u} &= \sum_{j=1}^{n} \big(\mathbf{u} \cdot \hat{e}_1 \big) \\ &= \sum_{j=1}^{n} \Big(-\int_{\Omega} \mathbf{W}_j^* \cdot \mathbf{f} \, d^n\boldsymbol{\xi} - \oint_{\partial\Omega} \mathbf{C}_j \cdot \hat{\mathbf{n}} \, ds \Big), \end{aligned} \tag{23}$$

where the subscript j refers to the fact that the Green's function and bilinear concommitant are generated by the directional delta function $\delta(\mathbf{x} - \boldsymbol{\xi})\hat{e}_j$. This can be rewritten

by collecting the $\mathbf{W}_j$ and $\mathbf{C}_j$'s into matrices, where each j denotes a row. Thus equation (23) can be written simply as:

$$\mathbf{u} = -\int_{\Omega} \mathbb{W}^* \mathbf{J}\, d^n\boldsymbol{\xi} - \oint_{\partial\Omega} \mathbb{C}\hat{\mathbf{n}}\, ds, \tag{24}$$

with the block letters denoting rank 2 tensors (a.k.a. matrices/dyadics).

The BVPs that shall be covered in this paper are self adjoint. This means that the BVP is the same as its adjoint and hence the Green's function is also self adjoint so we can write $\mathbb{W} = \mathbb{G}$, where $\mathbb{G}$ is the vector Green's function, constructed from $\mathbf{G}_i$'s and the adjoint of $\mathbb{W}$.

III. Specific Solutions for Electrodynamics

A. Free Space Magnetostatic Vector Potential

The 3-D static vector potential generated by a current distribution $\mathbf{J}(\mathbf{x})$ is governed by the Vector Poisson's equation in the Lorentz gauge [3]:

$$\nabla^2 \mathbf{A}(\mathbf{x}) = -\mu_0\, \mathbf{J}(\mathbf{x}), \tag{25}$$

where the region is unbounded. We choose the solutions such that $\mathbf{A}$, and it's derivatives go to zero as $|\mathbf{x}| \to \infty$, hence the bilinear concomitant vanishes. Substituting this into equation (22), we obtain:

$$\mathbf{A} \cdot \hat{x} = -\int \mu_o\, \mathbf{G}_x^*(\mathbf{x}|\boldsymbol{\xi}) \cdot \mathbf{J}(\boldsymbol{\xi})\, d^3\boldsymbol{\xi}. \tag{26}$$

We find $\mathbf{G}_x$ by solving the PDE $\nabla^2 \mathbf{G}_x(\mathbf{x}|\boldsymbol{\xi}) = \delta(\mathbf{x}-\boldsymbol{\xi})\hat{x}$, which can be done through separation of variables. However, because the function must go to zero at infinity, the components orthogonal to the current source have only the trivial solution. Thus,

$$\mathbf{G}_x(\mathbf{x}|\boldsymbol{\xi}) = -\frac{1}{4\pi|\mathbf{x}-\boldsymbol{\xi}|}\hat{x}. \tag{27}$$

By repeating this for $\hat{y}$ and $\hat{z}$,

$$\mathbb{G}(\mathbf{x}|\boldsymbol{\xi}) = g(\mathbf{x}|\boldsymbol{\xi})\, \mathbb{I}, \tag{28}$$

where $\mathbb{I}$ is the identity matrix and $g(\mathbf{x}|\boldsymbol{\xi})$ is the free space scalar potential Green's function from (17).

The vector potential for free space is then given by:

$$\mathbf{A}(\mathbf{x}) = -\mu_o \int \mathbb{G}^*(\mathbf{x}|\boldsymbol{\xi})\, \mathbf{J}(\boldsymbol{\xi})\, d^3\boldsymbol{\xi}. \tag{29}$$

This equation can be found in most sources on dyadic Greens functions [9], but is often assumed with, at best, only a loose justification. We have showed that the approach is justified in this case.

Recalling that $\mathbb{G}^* = \mathbb{G}$ and that it is diagonal, we obtain the well known equation for vector potential:

$$\mathbf{A}(\mathbf{x}) = \frac{\mu_o}{4\pi} \int \frac{\mathbf{J}(\boldsymbol{\xi})}{|\mathbf{x}-\boldsymbol{\xi}|}\, d^3\boldsymbol{\xi}. \tag{30}$$

B. Half-Space Magnetostatic Vector Potential

A current source sits somewhere in region Ω_1 above a boundary $\partial\Omega$ defined by the plane $z = 0$. The regions $\Omega_1 = \{z > 0\}$ is vacuum and $\Omega_2 = \{z < 0\}$ is linear with respective permativities and permiabilities ϵ_2 and μ_2. There are no charge distributions.

First we need find conditions for the uniqueness of the magnetic field. Let us define two solutions to an equivalent version of equation (25),

$$\nabla \times \nabla \times \mathbf{A}(\mathbf{x}) = -\mu_0\, \mathbf{J}(\mathbf{x}), \tag{31}$$

as $\mathbf{A}_1$ and $\mathbf{A}_2$, for some $\mathbf{J}(\mathbf{x})$. The difference $\mathcal{A} = \mathbf{A}_1 - \mathbf{A}_2$ then satisfies the equation:

$$\nabla \times \nabla \times \mathcal{A} = 0. \tag{32}$$

Taking the dot product between this with $\mathcal{A}$ then integrating over a region Ω gives:

$$\int_{\Omega} (\mathcal{A} \cdot \nabla \times \nabla \times \mathcal{A})\, dV = \tag{33}$$
$$\int_{\Omega} \Big(\nabla \times \mathcal{A}) \cdot (\nabla \times \mathcal{A}) - (\nabla \cdot \mathcal{A} \times \nabla \times \mathcal{A}) \Big)\, dV.$$

The LHS vanishes and we apply the divergence theorem to obtain:

$$\int_{\Omega} (\nabla \times \mathcal{A})^2\, dV = \oint_{\partial\Omega} (\mathcal{A} \times \nabla \times \mathcal{A}) \cdot \hat{\mathbf{n}}\, ds. \tag{34}$$

Since $(\nabla \times \mathcal{A}(\mathbf{x}))^2 > 0\ \forall \mathbf{x}$ then $\nabla \times \mathcal{A}(\mathbf{x})$, and hence the difference in magnetic fields between the two solutions will be zero only if the surface integral is zero. Therefore any magnetic field produced by an $\mathbf{A}$ that satisfies our BVP will be unique. There are several obvious cases where this will be true, which can be connected to scalar analogies as presented in [4]:

1) Dirichlet-esque boundary condition: Where $\mathbf{A}$, or $\hat{\mathbf{n}} \times \mathbf{A}$ is well defined along on $\partial\Omega$.
2) Neumann-esque boundary condition: Where $\nabla \times \mathbf{A}$, or $\hat{\mathbf{n}} \times \nabla \times \mathbf{A}$ is well defined on $\partial\Omega$.
3) Mixed or Robin-esque boundary condition: Where a linear combination of the above four terms is well defined on $\partial\Omega$.

Our boundary condition is informed by the fact that no current can cross $\partial\Omega$. Using Lorentz's force law [3] for some arbitrary vector $\mathbf{v}$,

$$\hat{e}_z \cdot \big[\mathbf{E} + \mathbf{v} \times \mathbf{B}\big]_{z=0} = 0. \tag{35}$$

But since $\mathbf{v}$ is arbitrary this only holds when $[\hat{e}_z \times \mathbf{B}]_{z=0} = 0$ and $[E_z]_{z=0} = 0$, with the latter trivially satisfied. Rewriting in terms of the vector potential, the BC can be expressed

$$\hat{\mathbf{n}} \times \nabla \times \mathbf{A} = 0, \tag{36}$$

a BC of the second type listed above.

In addition to this we require that our fields and their derivatives tend to zero as $|\mathbf{x}| \to \infty$.

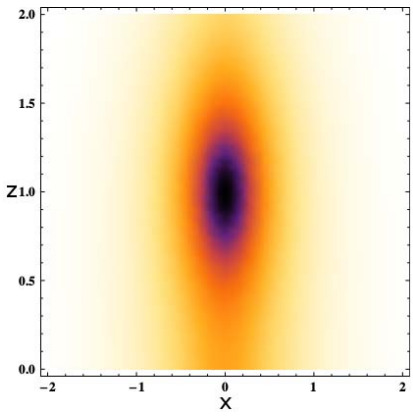

(a) A_x

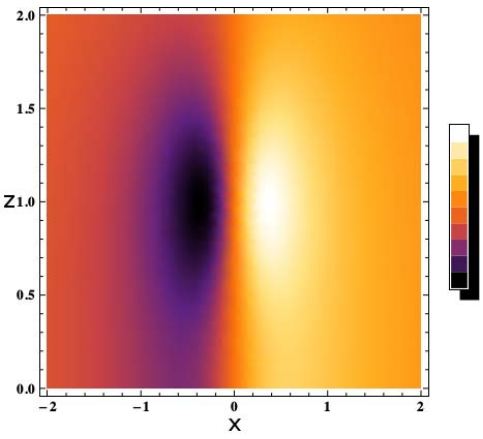

(b) A_y

Fig. 1: Vector potential at $y = 0.5$ generated by a small loop of current with $\rho_0 = 0.1$, centered at $z = 1$ in a half space $z > 0$. The scale and constants are chosen arbitrarily.

No looking at the vector Green's functions, $\mathbf{G}_i$, we firstly note that they are vector potentials themselves and therefore also restricted by condition (36). We assume that the tensor $\mathbb{G}$ is diagonal and substitute its components into (36) to obtain simplified BC's

$$\frac{\partial}{\partial z}\mathbf{G}_x\Big|_{z=0} = 0, \quad \frac{\partial}{\partial z}\mathbf{G}_y\Big|_{z=0} = 0, \tag{37}$$

$$\frac{\partial}{\partial x}\mathbf{G}_z\Big|_{z=0} = 0, \quad \frac{\partial}{\partial y}\mathbf{G}_z\Big|_{z=0} = 0. \tag{38}$$

The Green's functions here are written as vectors, but since two of their components are zero, we can treat them as scalar functions with well known solutions.

We shall find the vector Green's functions using the method of images, whereby the boundary conditions can be met by superimposing the free space solution of a delta source with free space solutions of delta sources located outside the region of interest [4]. Condition (38) is a homogeneous Dirchelet condition, which for an infinite plane the solution is given by superimposing an image delta source of opposite sign the same distance from the boundary, but on the other side as the main source. The conditions in (37) are homogeneous Neumann conditions, satisfied in the same way as above, only with the same sign as the original source [4]. Thus the Green's tensor takes the form:

$$\mathbb{G} = -\frac{1}{4\pi|\mathbf{x}-\boldsymbol{\xi}|}\mathbb{I} - \frac{1}{4\pi|\mathbf{x}-\boldsymbol{\xi}_r|}\begin{bmatrix}1&0&0\\0&1&0\\0&0&-1\end{bmatrix} \tag{39}$$

where $\boldsymbol{\xi} = (x', y', z')$ and $\boldsymbol{\xi}_r = (x', y', -z')$ with the r denoting reflection. The solution for $\mathbf{A}$ is provided by equation (24) with the surface integral vanishing under our conditions, assuming there is no time varying charge distribution.

An example use of this is an imperfect magnetic dipole, defined by the current distribution

$$\mathbf{J}(\boldsymbol{\xi}) = I\delta(\rho' - \rho_0)\delta(z' - 1)\hat{\boldsymbol{\phi}}', \tag{40}$$

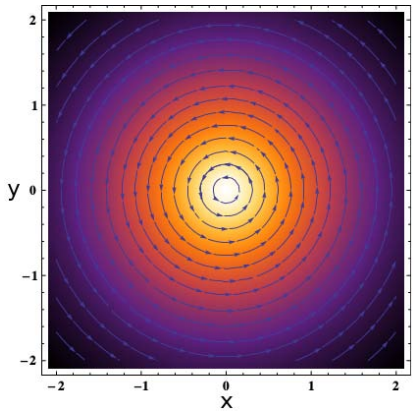

Fig. 2: A cross section of $\mathbf{A}$ at $z = 0.5$.

using cylindrical coordinates. Figures 1 and 2 were found using Mathematica's numerical integration package and plotting with arbitrary values.

C. Half-Space Harmonic Vector Potential

Now consider the Helmholtz equation [3],

$$(\nabla \times \nabla \times +k^2)\mathbf{A}(\mathbf{x}) = -\mu_0\,\mathbf{J}(\mathbf{x}), \tag{41}$$

with the same BC's as in the previous section.

Once again we must find uniqueness conditions, which we do by mirroring [10]. We start by defining $\mathcal{A} = \mathbf{A}_1 - \mathbf{A}_2$, where the $\mathbf{A}_i$ are solutions to (41) generated by the Green's function $\mathbf{G}_a^i$, defined by

$$(\nabla \times \nabla \times +k^2)\mathbf{G}_a^i = \delta(\mathbf{x} - \boldsymbol{\xi})\hat{\mathbf{a}}, \tag{42}$$

for an arbitrary unit vector $\hat{\mathbf{a}}$. The difference $\mathcal{A}$ satisfies the equation

$$(\nabla \times \nabla \times +k^2)\mathcal{A} = 0. \tag{43}$$

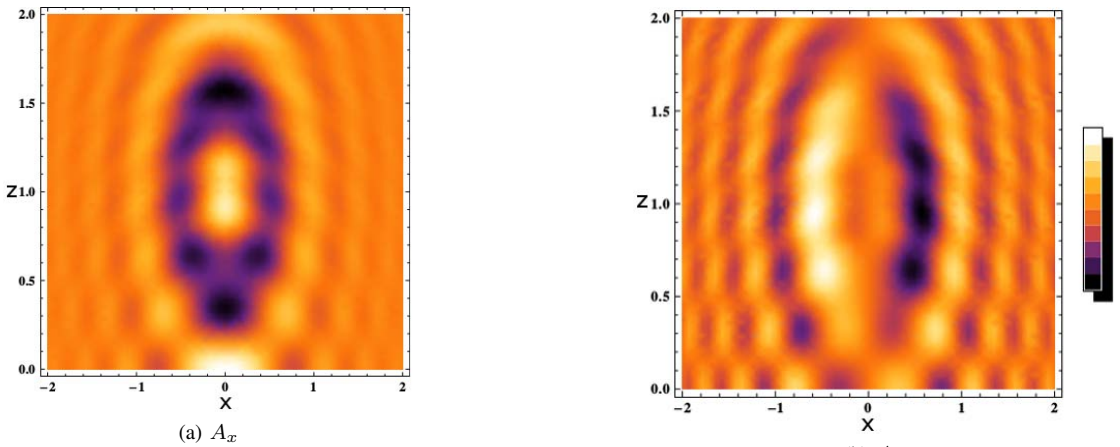

(a) A_x (b) A_y

Fig. 3: Vector potential at $y = 0.5$ generated by a small loop of harmonically oscillating current with $\rho_0 = 0.1$, centered at $z = 1$ in a half space $z > 0$. The scale and constants are chosen arbitrarily.

Green's vector identity gives us

$$\int_\Omega \left(\mathbf{G}_a^i \cdot (\nabla \times \nabla \times +k^2)\mathcal{A} - \mathcal{A} \cdot (\nabla \times \nabla \times +k^2)\mathbf{G}_a^i\right) dV' = \oint_{\partial\Omega} (\mathcal{A} \times \nabla \times \mathbf{G}_a^i - \mathbf{G}_a^i \times \nabla \times \mathcal{A}) \cdot \hat{\mathbf{n}}\, ds'. \quad (44)$$

We can straight away see that a vector Dirchlet condition of $\hat{\mathbf{n}} \times \mathbf{A}(\mathbf{x}) = -\mathbf{f}(\mathbf{x})$, yields

$$\mathcal{A} \cdot \hat{\mathbf{a}} = \oint_{\partial\Omega} \mathbf{f} \cdot (\nabla \times \mathcal{A})\, ds', \quad (45)$$

which is a constant, and therefore creates unique $\mathbf{B}$ and $\mathbf{E}$ fields. Likewise, a vector Neumann condition of $\hat{\mathbf{n}} \times \nabla \times \mathbf{A}(\mathbf{x}) = -\mathbf{f}(\mathbf{x})$, yeilds a constant value of $\mathcal{A} \cdot \hat{\mathbf{a}}$ and creates unique $\mathbf{B}$ and $\mathbf{E}$ fields.

The Robin condition is not as clear. One might try a boundary condition of the form

$$w(\mathbf{x})\mathbf{A}(\mathbf{x}) + \nabla \times \mathbf{A}(\mathbf{x}) = -\mathbf{f}(\mathbf{x}), \text{ or} \quad (46)$$

$$w(\mathbf{x})\hat{\mathbf{n}} \times \mathbf{A}(\mathbf{x}) + \hat{\mathbf{n}} \times \nabla \times \mathbf{A}(\mathbf{x}) = -\hat{\mathbf{n}} \times \mathbf{f}(\mathbf{x}). \quad (47)$$

Following the above approach we obtain by using triple cross product identities that

$$\mathcal{A} \cdot \hat{\mathbf{a}} = \oint_{\partial\Omega} \left((2w\mathbf{G}_a^i + \mathbf{f}) \times \mathcal{A}\right) \cdot \hat{\mathbf{n}}\, ds', \quad (48)$$

which is not a constant vector as $\mathbf{G}_a^i$ depends on both $\boldsymbol{\xi}$ and $\mathbf{x}$.

On deeper examination we find that there are in fact two Robin-esque conditions. They are:

$$w(\mathbf{x})\mathbf{A}(\mathbf{x}) + \hat{\mathbf{n}} \times \nabla \times \mathbf{A}(\mathbf{x}) = -\mathbf{f}(\mathbf{x}) \text{ and} \quad (49)$$

$$w(\mathbf{x})\hat{\mathbf{n}} \times \mathbf{A}(\mathbf{x}) + \nabla \times \mathbf{A}(\mathbf{x}) = -\mathbf{f}(\mathbf{x}). \quad (50)$$

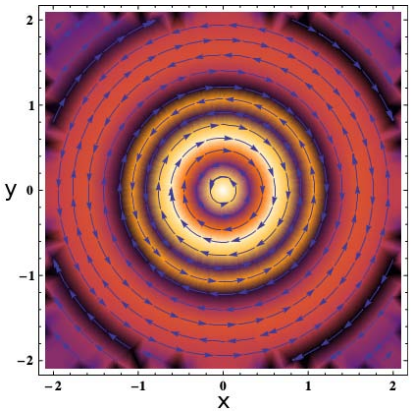

Fig. 4: A cross section of $\mathbf{A}$ at $z = 0.5$. Scale and constants are chosen arbitrarily.

Substituting equation (50) into (44) yields

$$\mathcal{A} \cdot \hat{\mathbf{a}} = \oint_{\partial\Omega} \hat{\mathbf{n}} \cdot \mathbf{f} \times \mathcal{A}\, ds', \quad (51)$$

which is a constant and creates unique $\mathbf{B}$ and $\mathbf{E}$ fields. The proof for boundary condition (49) is similar.

We have the BC's

$$\hat{\mathbf{n}} \times \nabla \times \mathbf{A} = 0, \text{ and} \quad (52)$$

$$\left[\frac{\partial}{\partial t} A_z\right]_{z=0} = 0. \quad (53)$$

We now know (52) guarantees the uniqueness of the electric and magnetic fields. Condition (53) is a further restriction for when we add the time dependence of $\mathbf{A}(\mathbf{x}, t) = \mathbf{A}(\mathbf{x}) \exp(i\omega t)$. This can only happen when $A_z(\mathbf{x}, t)\big|_{z=0} = 0$.

Because of our uniqueness theorems, we can once again use the method of images. The Green's functions are generated in the same way as in the previous section, only with the free

space scalar Green's function of Helmholtz equation. That is, our Green's tensor is given by:

$$\mathbb{G} = -\frac{e^{-ik|\mathbf{x}-\boldsymbol{\xi}|}}{4\pi|\mathbf{x}-\boldsymbol{\xi}|}\mathbb{I} - \frac{e^{-ik|\mathbf{x}-\boldsymbol{\xi}_r|}}{4\pi|\mathbf{x}-\boldsymbol{\xi}_r|}\begin{bmatrix}1&0&0\\0&1&0\\0&0&-1\end{bmatrix}. \tag{54}$$

This already satisfies BC (53) as any A_z generated by this will be zero at the boundary. Once again, the $\mathbf{A}$ field is generated by equation (24) and the surface integral vanishes under our conditions, assuming there is no time varying charge distribution.

Now we compute the field of an imperfect oscillating magnetic dipole, defined by the current distribution

$$\mathbf{J}(\boldsymbol{\xi}, t) = Ie^{i\omega t}\delta(\rho' - \rho_0)\delta(z' - 1)\hat{\boldsymbol{\phi}}', \tag{55}$$

using cylindrical coordinates (ρ, ϕ, z). The fields are shown in figure 3 and 4 were found again using Mathematica's numerical integration package and plotting with arbitrary values.

IV. Conclusion

In this paper we have combined the extension of distribution theory to vector distributions from reference [6] with the formalism of scalar Green's functions. We have applied this theory to Electrodynamics and have re-derived more generally several commonly used formulae, as well as discovering several uniqueness conditions for vector potentials.

Future work involves deriving the fields in the lower half-spaces for conducting, dielectric, and paramagnetic materials, as well as extending the finite boundary method to vectors in order to deal with inhomogeneities.

Acknowledgements

This work is funded as a part of the New Zealand National Science Challenge.

References

[1] E. N. Economou, *Green's functions in quantum physics*. Springer, 1984, vol. 3.

[2] M. Schmutz, "Real-time Green's functions in many body problems," *Zeitschrift für Physik B Condensed Matter*, vol. 30, no. 1, pp. 97–106, 1978.

[3] J. D. Jackson, *Classical electrodynamics*. Wiley, 1999.

[4] P. K. Banerjee and R. Butterfield, *Boundary element methods in engineering science*. McGraw-Hill London, 1981, vol. 17.

[5] V. V. Sergeevich, "Generalized functions in mathematical physics," *Moscow Izdatel Nauka*, vol. 1, 1976.

[6] R. J. Gaonon, "Distribution theory of vector fields," 1970.

[7] J. Stratton and L. Chu, "Diffraction theory of electromagnetic waves," *Physical Review*, vol. 56, no. 1, p. 99, 1939.

[8] J. R. Mentzer, *Scattering and diffraction of radio waves*. London, 1955, vol. 7.

[9] C.-T. Tai, *Dyadic Green functions in electromagnetic theory*. Institute of Electrical & Electronics Engineers (IEEE), 1994.

[10] L. Kannenberg, "Uniqueness of solutions to helmholtz's equation with linear boundary conditions," *American Journal of Physics*, vol. 57, no. 1, pp. 60–63, 1989.

Fractional Behaviour of Rechargeable Batteries

Rahat Hasan
School of Science and Engineering
University of Waikato
New Zealand

Jonathan Scott
School of Science and Engineering
University of Waikato
New Zealand

***Abstract*—For decades authors have preferred to model batteries with either Thevenin-style models using RLC, or Randles-style by adding a Warburg element. These are claimed to model accurately. We present convincing empirical evidence suggesting that a fractional-derivative (constant-phase element) model is required. Our data shows that existing state-of-the-art models may be overly complicated, requiring numerical rather than physical considerations to find parameters.**

I. Introduction

When subjected to a step change in current, batteries exhibit a step change in output voltage owing to their internal resistance. Following the step voltage change there additionally follows a gradual decay curve. This is usually attributed to chemical diffusion processes within the cell. Similarly, when the load current returns to zero, the terminal voltage does not immediately return to the steady-state, open-circuit voltage of the cell, but again exhibits a slow recovery.

Figure 1 shows such a recovery curve measured on a 900mAh nickel-metal hydride (NiMH) battery. The battery was cycled carefully to start in the 50-70 percent state of charge (SoC) range. The battery was connected to an E5270B and a constant current of 90mA was drawn for a period of 1 minute. This represents a discharge of only one-sixth of 1 percent of Q, the total capacity of the battery, drawn at the so-called 10C rate. In other words, only a small amount of the battery's capacity was drawn, and at a very modest rate. In spite of this, a significant change in terminal voltage is observed. As steady-state, open-circuit voltage is the most reliable indicator of a cell's state of charge, considerable effort has been put into understanding and modelling this recovery phenomenon.

The authors of [1], [2], [3], [4] modelled this characteristic using RC networks. These works were inspired by Randles original 1947 model [5], but disregard the fractional nature at which his work hints. Figure 2 shows a typical 2^{nd} order RC model where U_{OC} and U_t represent the open circuit and the terminal voltage respectively, and of course R_o represents the Ohmic series resistance. It is claimed that the first RC network of R_c and C_c represents the effects due to mass transport and the second RC network of R_d and C_d represents the double layer effect, after [5]. In the next section we will demonstrate that this entire class of model is inappropriate.

II. Application of Modified Swingler Method

In [6] Swingler proposed a modification of Gardener's method for resolving summed exponential functions. He

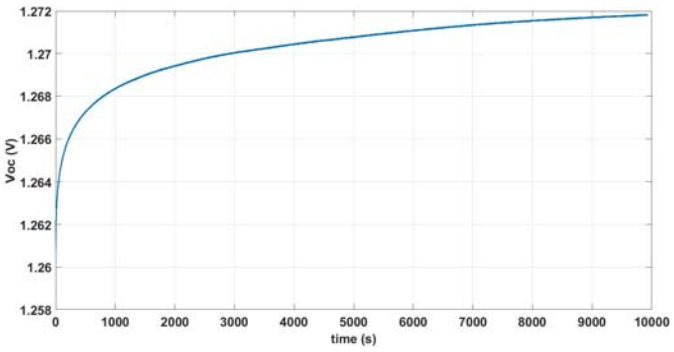

Fig. 1. The recovery curve of a 900mAh NiMH battery immediately after being subjected to a load of 90mA for 60 seconds beginning at a little over 50% SoC.

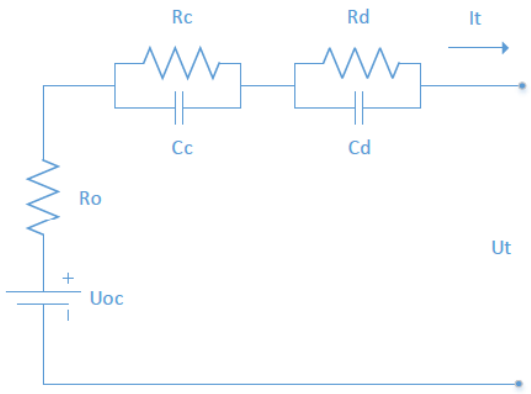

Fig. 2. A typical second-order RC battery equivalent-circuit model, reproduced from [3].

observed that a function $f(x)$ made up by summing a number of exponential decay terms could be processed to yield a series of delta functions whose amplitudes and delays betrayed the amplitudes and decay time constants of the constituent exponential functions. The execution of Swingler's process proved to be less simple than promised, but a modified algorithm was put forward in [7] that gives good results. This technique can be applied to the recovery part of a battery voltage waveform, and ought to identify the multiplicity of reactive elements required in a battery equivalent-circuit model, as each will give rise to a single decay time constant. We applied this algorithm to the recovery

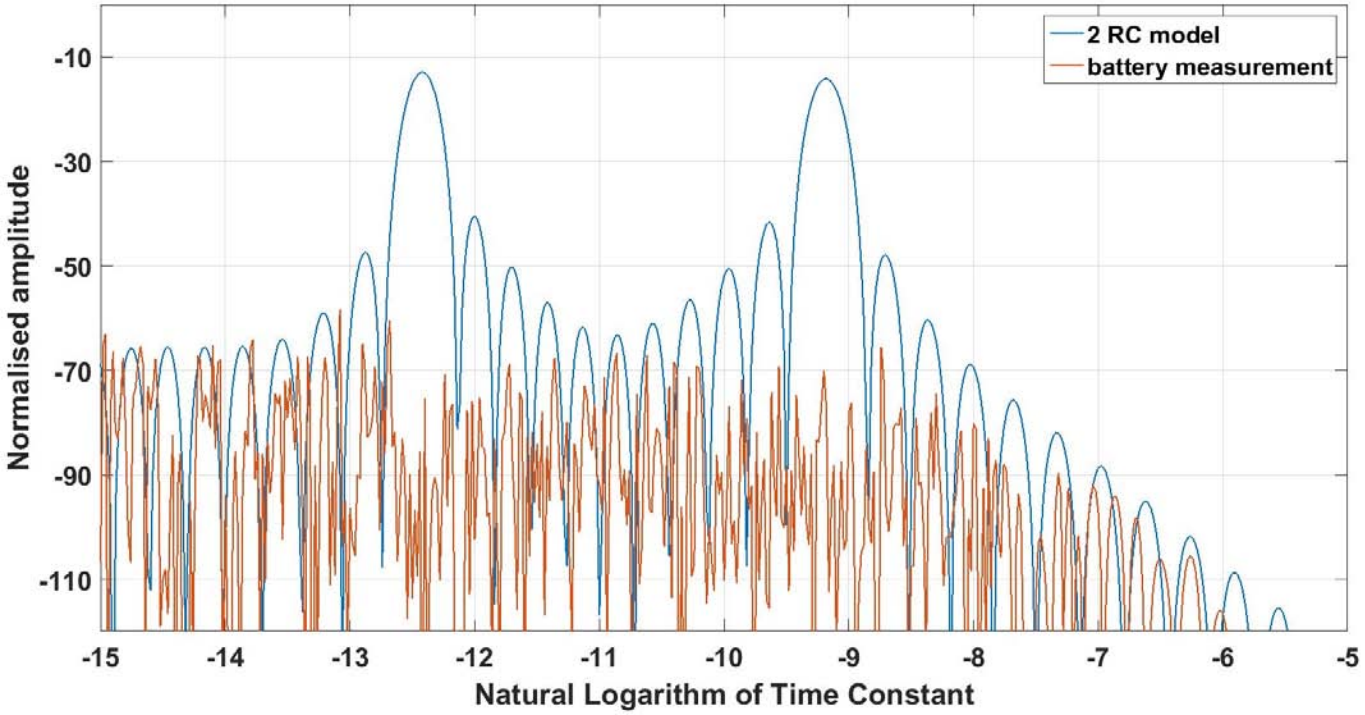

Fig. 3. Output of exponential-function analysis applied to the recovery curve of figure 1 and to a battery recovery curve generated using a two-RC battery model for comparison.

curve shown in figure 1. The result is shown in figure 3. The most important observation is that there is no evidence of any small number of exponential functions. Output of a two-RC model was analysed for comparison, and clearly shows 2 peaks 50dB above the noise floor. This observation suggests that RC models are not appropriate. A Constant-Phase Element (CPE) has a time-domain function that can only be approximated with an infinite series of exponentials, and is not expected to show any peaks on a Swingler-style analysis.

III. Fractional-Order Models

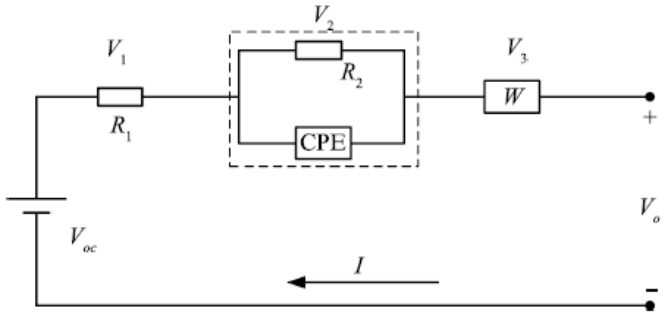

Fig. 4. Fractional equivalent circuit model reproduced from [12].

The idea of modelling batteries with fractional system was first introduced by the authors of [8] in 2006. The authors

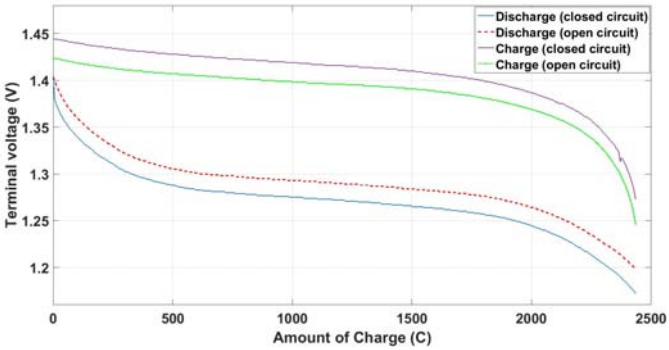

Fig. 5. Cycle test of 900mAh NiMH battery

claimed to be able to estimate the state of charge of lead-acid batteries within 5% error using a mathematical model based on limited frequency band of 2mHz-200Hz. One of the main drawbacks of this model is that it does not have any physical justification or any compact equivalent circuit. It is purely mathematical model with no clear electronic equivalent. A similar mathematical model involving complex algorithm was later proposed in [9] in 2010. This algorithm is specific to the cranking capability of a lead-acid battery. The best that can be said about this work is that it tends to confirm that batteries are fractional in their nature.

Other authors employed Randles battery model with varying degrees of success [10], [11]. In both papers, the authors measured impedance of a lead-acid battery over a certain

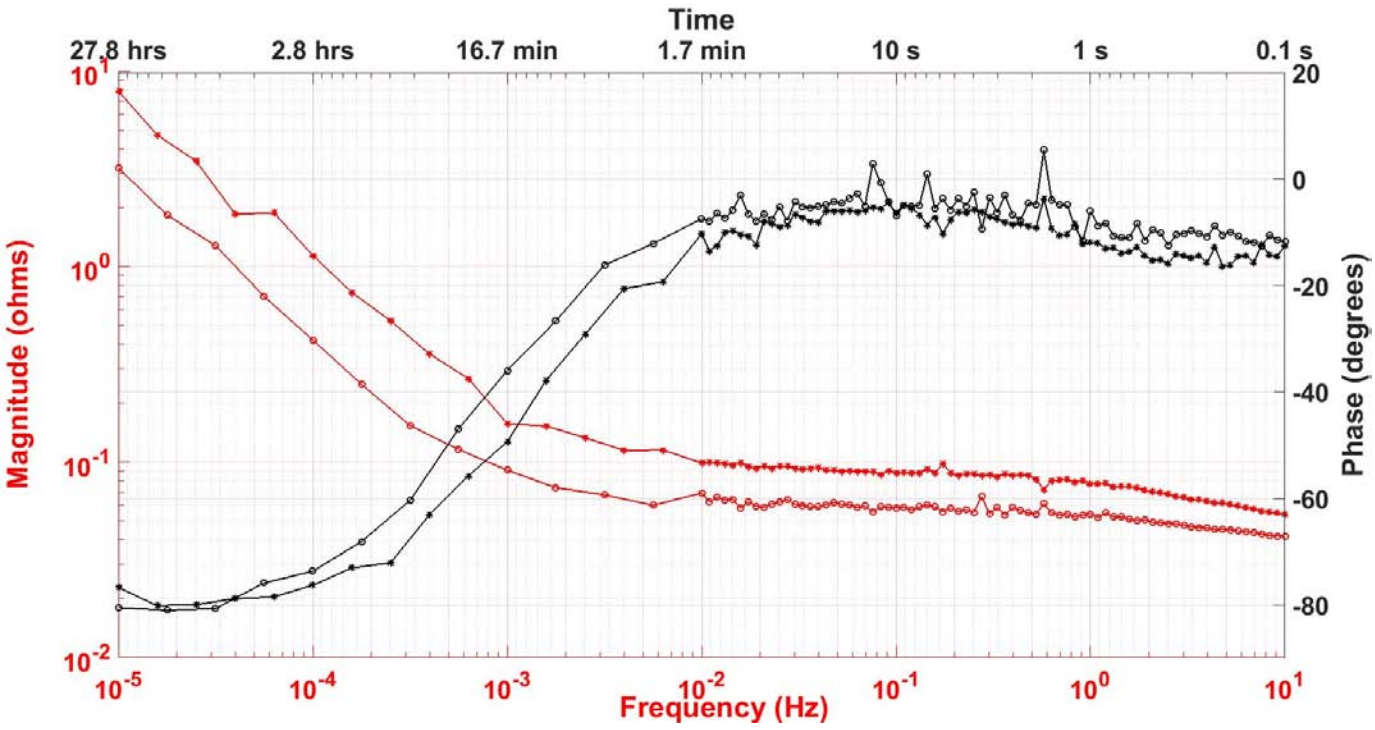

Fig. 6. The bode plot of the magnitude and phase of impedance of two NiMH batteries of 900mAh and 2400mAh capacity.

band of frequencies and used the frequency response to fit the parameters of a first-order Randles model.

As this manuscript was being prepared, Yan Ma et alia in [12] proposed a fractional battery model with a constant phase element (CPE) and a Warburg element as shown in figure 4. V_{oc} represents the open circuit voltage, V_o is the battery terminal voltage, R_1 represents the ohmic series resistance and W denotes the Warburg element. At first this work seems very powerful. In figures 1 and 5 of [12], the authors plot EIS data measured on a 26650 Lithium-ion battery on real/imaginary axes, but nowhere do they state the range of frequencies used in the measurement, nor do the plots show data points or variations with noise. In extracting their model parameters they eventually resort to a numerical fitting process. The model is then tested by having it predict very similar time-voltage data as that to which it was fitted. Finally, they note that the model predicts with "most errors below 20mV" which is claimed to represent only about 1% error in SoC, yet publically-available plots show in the linear region that Li-ion batteries have more like 3mV per percent of SoC.

We contend that any model of the complexity proposed in [12] can be fitted to a set of data and subsequently used to predict similar data. This does *not* suffice to verify the appropriateness of the model, especially if that model is overly complex, perhaps with too many degrees of freedom. We will now show measurement that suggest a simpler fractional-order model is appropriate.

IV. Cycle test

A 900mAh NiMH battery was cycled in order to determine the full capacity of the battery. This is important as we want to be certain that the range of the SoC stays within 50% during impedance measurement. The steps followed to obtain figure 5 are listed below:

- A current pulse of 0.18A (0.2 C) was generated to charge or discharge the battery for a period of 1 minute using Agilent E5270. Agilent E5270B Precision IV Analyzer contains SMUs (Source/Monitor Units) for voltage/current sourcing and voltage/current measurement as low as 0.1 fA.
- The battery was then allowed to rest for 2 minutes for the recovery voltage to settle down after every 0.3% SOC charging and discharging.
- The battery was idled for 12 hours in between charge and discharge.
- Open circuit and under-load terminal voltages of the battery were measured..

V. Impedance measurement

We measured the impedance of two NiMH batteries against frequency from 10μHz to 10Hz using a Solartron 1260A analyser with a fixed dc offset corresponding to 50% SoC. The batteries were rated at 900mAh and 2400mAh. Stimulus levels were chosen to ensure that cells did not deviate more than 10% from 50% SoC for even the lowest stimulus frequencies, where current flowed in one direction for periods approaching 14 hours. Figure 6 depicts the results.

For both cells, the magnitude of impedance is relatively flat at higher frequency, but increases below 10mHz, while phase shifts from about 0 degrees to settle at about -80 degrees. Such a Bode plot is characteristic of a single CPE corresponding to a derivative of order 0.89.

The phase traces in figure 6 show a deviation of 10–15 degrees as frequency increases above 1Hz. We have not conducted any analysis as to what might cause this so far, chiefly as we are interested in modelling SoC, and the response in this frequency range is not really of interest.

The frequency responses of NiMH obtained at higher frequencies are noisy. In order to confirm that the noise were not generated from the measurement system, a resistor of value comparable to the magnitude of battery impedance was measured using Solartron 1260. The results reproduced shown in figure 7 concludes that the noise in the impedance measurement data for batteries were not generated from Solartron 1260.

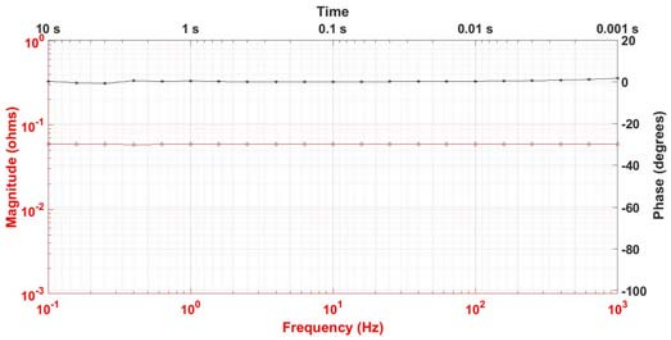

Fig. 7. The bode plot of the magnitude and phase of impedance of a pure resistor in the Solarton measurement setup.

VI. Discussion and Conclusion

We have shown that the impedance characteristic of NiMH batteries corresponds to that of a single CPE in series with a fixed resistor. The authors of [13] and [14] have adopted a similar approach to model the impedance characteristics of electrode-electrolyte interface and implantable electrode respectively with CPE. Since a CPE is defined by two parameters, its magnitude and the order of the derivative function relating current and voltage, this should lead to a battery model with greatly reduced parameter set.

References

[1] J. Chiasson and B. Vairamohan, "Estimating the state of charge of a battery", *IEEE Transactions on Control Systems Technology*, volume 13, Issue 3, May 2005, pp465–470.

[2] Tiezhou Wu, Lunan Liu, Qing Xiao, Quan Cao, and Xieyang Wang, "Research on SoC Estimation Based on Second-order RC Model", *Telkomnika*, vol. 10, no. 7, Nov 2012, pp. 1667–1672.

[3] ZeCheng, JikaoLv, YanliLiu, and ZhihaoYan, "Estimation of State of Charge for Lithium-Ion Battery Based on Finite Difference Extended Kalman Filter, *Journal of Applied Mathematics*, 2014, Vol 2014, pp. 1–10.

[4] Wu Guoliang, Lu Rengui, Zhu Chunbo, and C.C.Chan, "State of Charge Estimation for NiMH Battery Based on Electromotive Force Method", *IEEE Vehicle Power and Propulsion Conference (VPPC)*, September 3–5, 2008, Harbin, China.

[5] Randles, J. E. B., "Kinetics of rapid electrode reactions", *Discussions of the Faraday Society*, 1:11, 1947. doi:10.1039/df9470100011

[6] D. N. Swingler, "A Differential Technique for the Fourier Transform Processing of Multicomponent Exponential Functions", *IEEE Transactions on Biomedical Engineering*, vol. BME-24(4), 1977, pp408–410.

[7] Rahat Hasan and Jonathan Scott, "Application of Swingler's method for analysis of multicomponent exponentials with special attention to non-equispaced data", *2016 IEEE 12th International Colloquium on Signal Processing & Its Applications (CSPA)*, 2016, pp. 12-15.

[8] Jocelyn Sabatier, Mohamed Aoun, Alain Oustaloup, Gilles Gregoire, Frank Ragot and Patrick Roy, "Fractional system identification for lead acid battery state of charge estimation", *Signal Processing - Fractional calculus applications in signals and systems*, Volume 86 Issue 10, October 2006 , pp. 2645 - 2657.

[9] Mikal Cugnet, Jocelyn Sabatier, Stphane Laruelle, Sylvie Grugeon, Bernard Sahut, Alain Oustaloup, and Jean-Marie Tarascon, "On Lead-Acid-Battery Resistance and Cranking-Capability Estimation", *IEEE TRANSACTIONS ON INDUSTRIAL ELECTRONICS*, VOL. 57, NO. 3, MARCH 2010, pp. 909-917.

[10] Luiz Carlos Stevanatto, Valner Joo Brusamarello, and Stanislav Tairov, "Parameter Identication and Analysis of Uncertainties in Measurements of LeadAcid Batteries",*IEEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT*, VOL. 63, NO. 4, April 2014, pp. 761-768.

[11] Stanislav Tairov and Luiz Carlos Stevanatto, "Impedance Measurements for Battery State of Health Monitoring", *2nd International Conference on Control, Instrumentation and Automation (ICCIA)*, 2011, pp. 79-83.

[12] Yan Ma, Xiuwen Zhou, and Hong Chen, "Fractional Modeling and SOC Estimation of Lithium-ion Battery", *IEEE/CAA Journal of Automatica Sinica*, Vol 3, No. 3, July 2016, pp. 281-287.

[13] Mark H. Jones and Jonathan Scott, "Scaling of Electrode-Electrolyte Interface Model Parameters in Phosphate Buffered Saline", *IEEE Transactions on Biomedical Circuits and Systems*, November 2013, pp. 1-8.

[14] J. Scott and P. Single, "Compact nonlinear model of an implantable electrode array for spinal cord stimulation (SCS)", *IEEE Transactions on Biomedical Circuits and Systems*, vol. 8, no. 3, pp. 382-390, Jun 2014.

A MATLAB ToolBox for the Bit-Stream Implementation

Faiz Rasool
Dept. of Engineering
International College of Auckland
Auckland, New Zealand
faiz@ica.ac.nz

Syed Adeel Ali
APIMATIC
Auckland, New Zealand
Adeel05@gmail.com

Muhammad Waseem Soomro
School of Professional Engineering
Manukau Institute of Technology
mwaseem@manukau.ac.nz

Abstract* — The *supremacy of bit-streams over dedicated VLSI technology in the implementation of parallel, real-time systems, regarding cost and system structure, has been reported in recent research activities. The aim of this paper is to demonstrate the development and the implementation of a library of bit- stream elements, in a GUI environment like the Simulink This gives the advantage of using the bit- stream elements in conjunction with a large number of existing continuous time Simulink elements. Thus, all sorts of new or existing system models, supported by Simulink can be tested under bit- stream environment. The focus of this study is on the bit-stream implementation of building blocks for Fuzzy Systems.

Keywords — Fuzzy Systems, Neural Networks, Bit Streams, MATLAB ToolBox

I. INTRODUCTION

It has not been so long since the distributed computing began to replace the conventional computing in numerous applications due to speed and performance [1]. The application areas of distributed techniques range from stochastic modeling of the clinical epidemiology to the cognitive model building of the human mind, from forecasting the climate in the current century to deducing the structure of the Milky Way galaxy. In this study, distributed techniques will be discussed in the context of Fuzzy and Neural Systems.

The goal is to implement fuzzy systems and neural networks on hardware because of its advantage of speed and size over microcontroller and PC. In analogue domain, excellent results are available in the implementation of fuzzy systems [2, 3, 4]. Some hybrid and digital technique based results are also reported for fuzzy systems implementation [5, 6]. Both analogue and digital techniques have been applied for neural network implementations as well. Mead [7] used VLSI (Very Large Scale Integrated Circuit) technique to build neural networks for character recognition. Usually, the digital implementation is done using the pulse code encoding of analogue values [8, 9, 10].

Patel used FPGAs to present a novel technique, the bit-stream representation of the analogue signals, for the hardware implementation of the distributed systems such as Fuzzy and Artificial Neural Networks (ANNs). He designed and constructed various bit-stream functional elements using VHDL [11]. These elements are modular and can be connected to the construction of complex bit-stream structures. In the earlier research stages, Patel implemented cascade controllers [15] by using single bit stream instead of a multi-bit word for analogue signal representation. In [12], he used the same technique to present an alternative approach for the parallel execution of algorithms. In [14], he applied this bit-stream technique on a D.C. servo mechanism and proved it as a feasible alternative for controller construction for multi- dimensional systems. Later on, he came up with a fuzzy solution by implementing a fuzzy controller [13] for the same servo motor system and demonstrated the success of the uniformly weighted bit-streams technique. The implementation on Artificial Neural Networks has also been presented by Patel in [15] using the similar technique.

The success and the potential of the uniformly weighted bit-streams technique presented by Patel, for the parallel, real time systems, is the motivation for this study. It is already discussed earlier that bit-stream elements have been designed by Patel in VHDL. Although there is no issue with the performance of the bit-stream functional elements, the limitation was that the elements require complete systems modeled in VHDL to run at . Also, the existing systems are quite difficult to understand and modify. Therefore, an environment was required, where bit-stream elements could be modeled and simulated with ease. That environment could rapidly build complex systems and modify conveniently when needed. Simulink is the MATLAB tool for model-based design which can be used to build new models or modify the existing ones accordingly. It supports both linear and nonlinear systems and can model them in sampled time, continuous time or a combination of both. More importantly, the complexity of the system is not an issue as it is in VHDL system design. Therefore, Simulink was chosen for the bit-stream modeling. It has the advantage of having numerous existing models which can be re-used with bit-stream elements. For the development, the powerful programming language of MATLAB was selected. Another attraction was the availability of the translation of the Simulink models into several hardware description languages by MATLAB. Thus, models can be designed, simulated as well as implemented on FPGAs with ease. Contribution of the current work can precisely describe as:

- A Simulink library of bit-stream functional elements is built including the basic arithmetic elements as well as the advanced elements for fuzzy and neural network systems. The elements must have their inputs and outputs of bit-stream type
- The functionality of compatible is verified with the existing VHDL bit-stream elements.
- The performance of the bit-stream elements is tested in parallel, real-time systems modeled in Simulink.

This paper is organized into six sections. Section 2 presents the conversion of the analogue signal into bit streams. In Section 3, construction of primitive and arithmetic elements is discussed. Section 4 talks about the design procedure for the building block of fuzzy systems. In Section 5, a bit-streams implemented example from Artificial Neural networks are discussed, and results are compared with already existing results. The paper is concluded in Section 6.

II. Signal Representation and Bit-Stream Generation

Bipolar analog signals can be converted into digital signals by using either multi-bit or single bit conversion. The multi-bit conversion uses a binary number system with each bit having weight 2m−1. In single bit conversion system, each bit is assigned a weight by its distance from preset reference point [16]. Later approach, called bit stream representation of analog signals, is adopted in this research. In this single bit representation approach, weights of positive 'Q+' and negative 'Q−' logic levels or quanta are used to represent the magnitude of any analog signal:

$$Vi(k) = A(k) - B(k) / N \qquad (2.1)$$

where Vi(k) = analogue input value at kth instant, A(k) = summation of positive quanta, B(k) = summation of negative quanta, N = size of the frame under consideration.

Here, it is worth mentioning that if some Q+ and Q− are same in a bit-stream, then this will have net result zero or Z. Similarly, some positive quanta Q+ result in an increase in magnitude and vice versa (see [17] and references therein for detail). This fact is elaborated in Figure 1.

Various arithmetic and fuzzy blocks are developed in this work. However, all the blocks accept input signals in the form of bitstream only. It requires a bit-stream generator which can convert any analog signal in the form of bit-streams. It is implemented in Simulink and consist of:

1. A circuit containing summer, amplifier, integrator, and a relay.
2. 2. A D-Flip-Flop
3. 3. A feedback path with a function to convert D-Flip-Flop's outputs 0 and 1 into −1 and 1, respectively

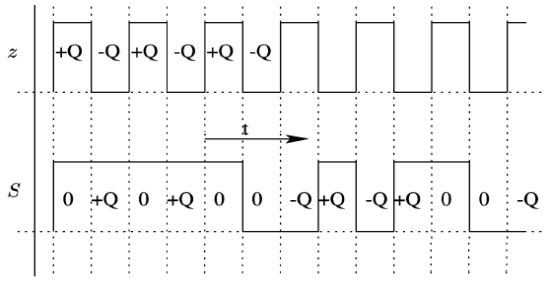

Figure 1: The Z along with a Positive Signal

This generator works for the analogue input range between −1 V to +1 V. For an input of +1 V; a bit-stream is generated containing a series of '1's. Similarly, a series comprised of all '0's will be the result of −1 V as the input. For an input of 0 V, the zero valued bit-stream with an alternating sequence of '1's and '0's will be generated. In Figure 2, an analog sinusoid and its equivalent bit-stream pattern are shown.

It can be seen that as the value of sinusoidal is increasing then, bits with 1ts are increasing, and when the value is maximum, 1, then 1ts bits are maximum. Similarly, as value is decreasing then 0ts bits are increasing up to the maximum.

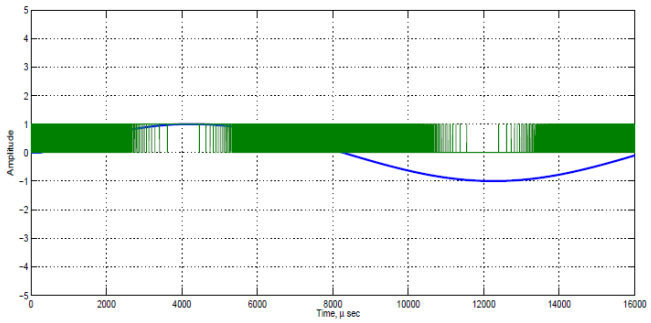

Figure 2: Analog Signal to Bit-Stream Conversion

III. Arithmetic Elements

The backbone of bit-streams based MATLAB toolbox for complex fuzzy and neural networks are basic arithmetic function elements. In this section, firstly general technique is discussed then primitives are designed. In the end, an example from arithmetic elements is presented.

3.1. General Technique

All the elements, despite different functionality, follow same construction technique. By Equation 2.1 any signal can be represented using bit-streams as

$$\begin{aligned} V_i(k) &= \frac{A(k) - B(k)}{N} = \frac{D_N(k)}{N} \\ V_i(k) &= \frac{A_z - B_z + D_N(k)}{N} \\ V_i(k) &= \frac{(A_z + D_N(k)/2) - (B_z - D_N(k)/2)}{N} \end{aligned} \qquad (3.1)$$

Where A_z and B_z are positive and negative quanta of zero or Z bit-stream. Zero valued components can be decoupled from any bit-stream "BS" as:

$$\begin{aligned} h &= BS \wedge \bar{Z} \\ l &= \bar{BS} \wedge Z \end{aligned} \tag{3.2}$$

After this, we can present procedure to design arithmetic components:

1. Separation of the non-zero components from the bit-stream using equation 3.2
2. Application of element specific operation using a truth table
3. Either
   - Re-insertion of the operated components into a zero valued bit-stream, or
   - Generation of output bit-stream using binary to the bit-stream generator in case multi-bit words is produced after step 2.

3.2. Primitive Elements

i. Z-Generator

A zero value bit-stream generator or 'Z' generator can be obtained from a clock FB. It is essential components of all the arithmetic elements as discussed earlier. It waveform and construction can be shown from Figure 3.

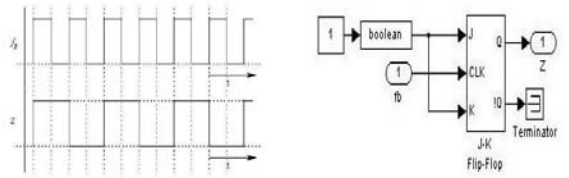

(a) Wave Form (b) Z-generator's construction

Figure 3: A Z-generator with waveform and Construction Diagram

ii. Saturating Summer and Accumulator (SAC-SAT)

The SAC-SAT block is used in functional arithmetic elements before the 3rd step discussed in general technique ([17] and references therein for detail). Saturating Summer and Accumulator (SAC-SAT) is a register used to buffer quanta for the functional elements. The function of SAC is to sum and accumulate the quanta while the SAT comes into action after reaching its maximum capacity to prevent an overflow condition. A positive quantum increases the value of SAC while a negative quantum decreases it while SAT overcomes the saturation states by bounded it to upper or lower limit respectively as shown in Figure 4.

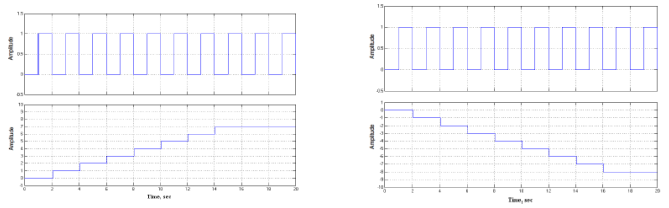

(a) Conceptual behavior of 4-bit SAC-SAT for a constant 1 (b) Conceptual behavior of 4-bit SAC-SAT for negative constant −1

Figure 4: Conceptual behavior of 4-bit SAC-SAT

iii. Saturating Summing and Accumulating Out-putter (SAC-SAT-out)

One of the basic requirements of this study is that all the functional elements must take inputs of bit-stream type but, there are some instances where processing elements output multi-bit words. To transmit to other elements, these words must be converted to bit-streams and hence a binary to the bit-stream generator is required. Block diagram of a binary to the bit-Stream generator is shown in Figure 5.

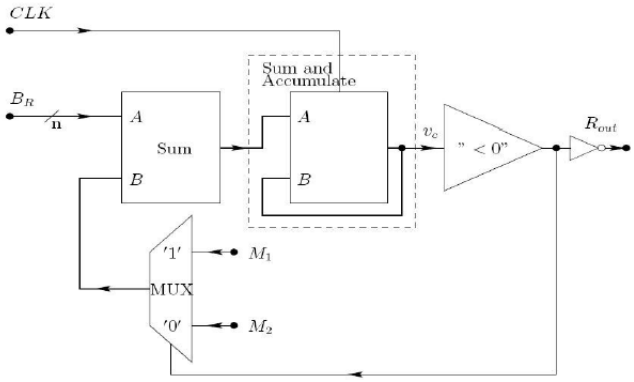

Figure 5: Binary to Bit-Stream Generator

3.3. Construction of Arithmetic Elements

Various arithmetic elements are designed in this toolbox such as an adder, subtracter, integrator, differentiator, and multiplier, etc. Due to space limitation, the only adder is discussed here.

i. Adder

The summation is the most familiar arithmetic operation. In case of bit-streams, the summation of two bit- stream signals means the addition of the quanta present in both bit-streams i.e. if a bit-stream S1, having 2 positive quanta in a frame of 8 bit wide is added with another bitstream S2, having 1 positive quanta in a frame of identical width then the resulting bit-streams S3 will contain 2 + 1 = 3 quanta within the dimensions of a similar frame. From Equation 3.1, the addition of bit-stream signals can be given as follows:

$$
\begin{aligned}
V_i(k) &= \frac{D_N(k)}{N} \\
V_3(k) &= V_1(k) + V_2(k) \\
V_3(k) &= \frac{D_{N1} + D_{N2}}{N} \\
V_3(k) &= \frac{A_z - B_z + D_1 + D_2}{N} \\
V_3(k) &= \frac{(A_z + D_1/2 + D_2/2) - (B_z - D_1/2 - D_2/2)}{N} \qquad (3.3)
\end{aligned}
$$

Where ($Az + D_1/2 + D_2/2$) are positive quanta and ($B_z - D_1/2 - D_2/2$) are negative quanta. $D_1/2$ and $D_2/2$ represent aggregate difference over an N sized frame.

The summation can be easily performed using a logical OR operation as far as quanta from both the signals do not arrive at the same time. Otherwise, one of the two quanta will be lost. Since it is very much probable that two quanta arrive at the same instant, a logical OR-based adder could not be used. One solution is to store the additional quanta in some buffer and then place it in the next available slot. In this way, a lossless adder can be made. It is worth mentioning that adder can be constructed by using either SAC-SAT or binary to bit-stream generator. We will only discuss the first type of adder here.

ii. Adder using SAC-SAT

As discussed before, a buffer is required to construct a lossless adder. Recall that our primitive element SAC-SAT can work as a buffer, so it could be used here to accumulate the quanta, but since the accumulated Quanta are also required to be ejected as output bit-stream, a SAC-SAT-out is a better choice and is adopted to construct the adder of Type 1. The structure of this bit-stream adder can be divided into two sections: Input section and output section. A SAC-SAT based adder is shown in Figure 6.

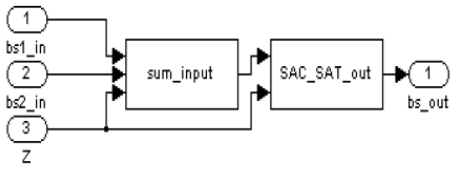

Figure 6: SAC-SAT based adder

iii. Input Section

The task of the input section is to determine the number of quanta arrived at the input. This task is performed using the truth table shown in Figure 7.The Column labeled as Z contains the states of zero valued bit-stream, S1, and S2 labeled columns depict the states of the input bitstreams that are to be summed and the last column with the heading Incident Net Quanta shows the computed incident quanta processed using the first three columns. The operation can be summarized as

- If both the inputs were similar in phase and shape with Z, then there are no quanta at the input. So a '0' is fed to the output section.
- If Z and S1 are at '0' logic and S2 at logic '1' then one positive quanta has arrived at the input, asserted by S2
- Similarly, if S1 and S2 are at logic '0' while Z is at logic '– 1' then one negative quanta has been asserted by each of the bitstreams at the input and hence −2Q is sent to the output section.

Table 3.1: Binary to Bit-Stream Generator

	Z	S_1	S_2	Incident Net Quanta
1	0	0	0	0
2	0	0	1	+1Q
3	0	1	0	+1Q
4	0	1	1	+2Q
5	1	0	0	-2Q
6	1	0	1	-1Q
7	1	1	0	-1Q
8	1	1	1	0

iv. Output Section

The output section is comprised of the primitive, SAC-SAT-out. Its task is to accumulate the quanta provided by the input section and to insert the quanta at appropriate slots. Its working has been described earlier using a basic example of single bit-stream. Now presenting another example that shows the behavior of the SAC-SAT-out i.e. the output section of the adder during the addition of two bit-streams. Figure 8 shows this addition with two bit-streams, S1 and S2 along with zero references, Z. The table down the waveforms consists of three rows. First shows quanta being fed at the input, second depicts the states of SAC before and after quanta insertion, and the last row shows the output bit-stream, S0. All the steps of the operation are based on the methodology described earlier with an example, therefore, are not analyzed in detail here, but a flowchart presenting an overview of the operation is shown in Figure 9 to describe how it has been coded. This flowchart shows a single iteration of the output section.

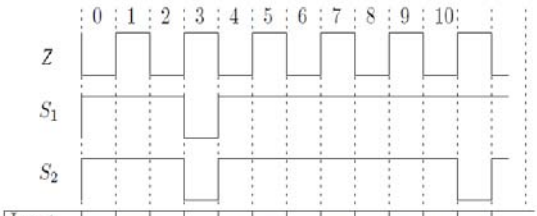

Input Quanta	+2	0	+2	-2	+2	0	+2	0	+2	0	+2	-1
SAC	2,1	1,1	3,2	0,0	2,1	1,1	3,2	2,2	4,3	3,3	5,4	3,3
S_o	1	1	1	1	1	1	1	1	1	1	1	1

Figure 8: Bit-stream addition using Adder

Figure 10 presents the Simulink model of the summation of two sinusoids, S1 and S2. Some components along with the adder can be seen including 2 Analogue-to-Digital- Converters, 2 signal generators, 1 pulse generator, 1 Z generator, 1 value element, 1 scope with a multiplexer and the Adder of Type 1. The specifications for this model are:

$$\begin{aligned} S_1 &= \frac{60}{128}cos(200\pi t) \\ S_2 &= \frac{60}{128}cos(200\pi t + \frac{\pi}{2}) \\ f_b &= 40KHz \\ N &= 256 \end{aligned} \tag{3.4}$$

Figure 11 shows the output waveforms of the above simulation along with the VHDL simulation results of [11] for comparison. It can be clearly seen that the waveform for both simulations is identical as well as the theoretical peak value of t=0.56. Although, the actual inputs, and the outputs are bit-streams, but the plot presents analogue values. This has been done by using the "value element" whose task is to add the quanta of the past N clocks first and then to normalize the count value by N = 2. So the bit-streams range from -128/128 to 127/128 for N= 256, R= 8.

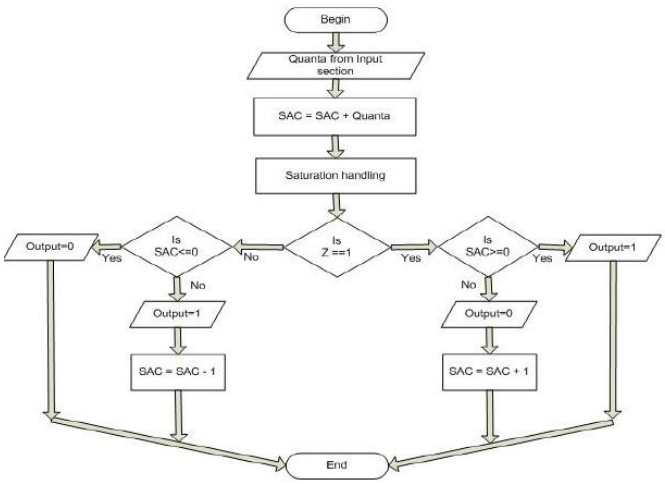

Figure 9: Output Flowchart for Adder

IV. FUZZY LOGIC TOOLBOX

Along with the functional elements discussed in the previous section, a few more components are required to design bit-stream fuzzy systems. Particularly they are Min/Max operators, membership function generator, and a defuzzification element. The elements must satisfy the basic requirement of inputs and outputs of bit-stream type. Only membership function generator is discussed here due to space limitation.

i. Bit-Stream Membership Function Generator

For the bit-stream implementation, a membership generator has been presented in Figure 12 (for single membership value) and Figure 13 (for both positive and negative membership values).

The process is an addition to two offsets. First one is the modified form of the primitive, Binary to Bit-Stream Generator and the second one is a constant with the value 1. Figure 14 shows the concept of this generator. The dotted line, labeled as 1, is the normal bit-stream trace with slope = 2. For a membership curve with slope = 2 and having a breakaway point at +32, the input needs to be kept at 0 until the breakaway point comes, which is correspondent to the point at +64 of line 1. So, adding an offset of −(64 + 128) = −192 results the dashed line, labeled as 3 and by further adding an offset of +128, the desired curve is obtained finally and presented as line 2. The second offset is a permanently set as 1, while the first offset is based on the input parameters: M_1, M_2, and P as shown in the model. The SAC is incremented or decremented at every bit-clock by P (a multi-bit word) using the Truth Table presented in Figure 15, while M_1 and M_2 have the same characteristics as used in other blocks. The first offset value can be calculated by choosing appropriate values of the input parameters in the following equations.

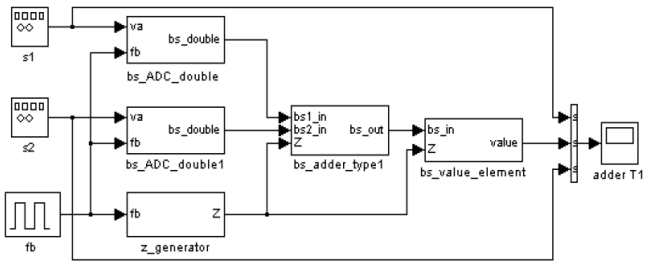

Figure 10: Simulink model for bit-stream summation using adder

$$\begin{aligned} offset \times 2 &= N\frac{M_1 + M_2}{M_1 - M_2} \\ \frac{P}{\|M_1 - M_2\|} &= 2 \end{aligned} \tag{4.1}$$

To obtain a negative membership value, the same procedure is used except that the resulting membership value from Linear Membership Function Generator is multiplied by −1 and then the second offset i.e. 1 is added as shown in Figure 13.

Example

Presenting a negative linear bit-stream membership function with a breakaway point at +48/128, slope = 2 and N = 256. For the offset value, notice that +48 is correspondent to +96 on line 1 of Figure 14. Hence, the offset = − (96 + 128) = − 224. The waveform has been plotted in Figure 16.

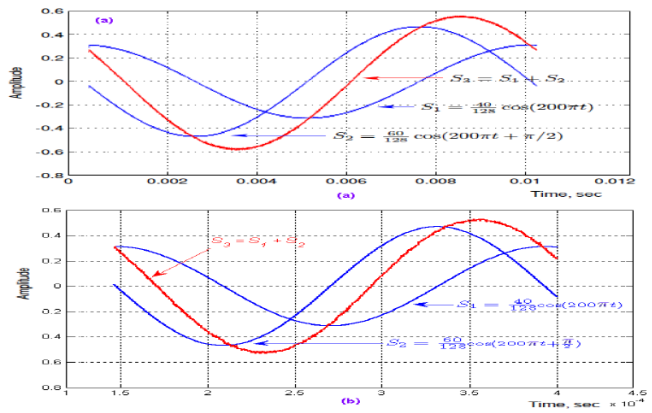

Figure 11: Comparison between Adder simulations for (a) VHDL (b) MATLAB

$$
\begin{aligned}
offset \times 2 &= N\frac{M_1 + M_2}{M_1 - M_2} \\
-192 \times 2 &= 256\frac{M_1 + M_2}{M_1 - M_2} \\
\frac{M_1}{M_2} &= \frac{1}{5} \\
\frac{P}{\|M_1 - M_2\|} &= 2 \\
\frac{P}{\|1 - 5\|} &= 2 \\
P &= 8
\end{aligned} \tag{4.2}
$$

V. Artificial Neural Networks Example

In this section, an academic example from neural networks, called XOR pattern matching, is presented. This example is implemented using bit-streams based MATLAB toolbox and results are compared with VHDL based results. It is assumed that reader has basic concepts of neural networks.

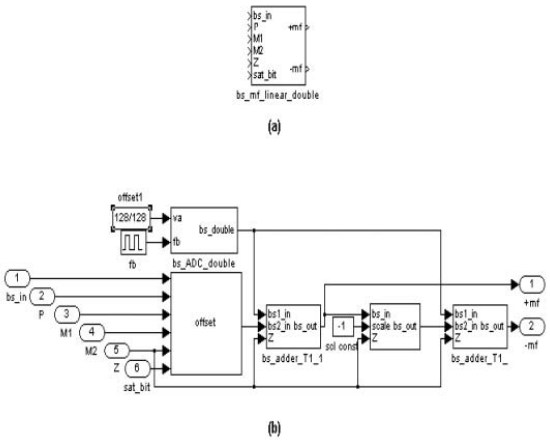

Figure 12: Linear Membership Function Generator (a) Compact (b) Expanded

5.1. XOR Pattern Matching

The categorization of a pattern set having an exclusive-OR behavior is one of the typical and simple ex- examples that shows the potential of any network to categorize a non-linearly distinguishable set of patterns. The network has to be multi-layered, and the XOR set of patterns can be presented in a Uni-Polar or a Bi-Polar form. Patel [11] has successfully implemented both Bi-Polar and Uni-Polar XOR classifiers using bit-streams in VHDL. In this section, the bit-stream implementation of Bi-Polar classifiers will be presented only using the bit-stream components developed in MATLAB along with the comparison with the VHDL implementation.

Bi-Polar Exclusive-OR

Figure 17 shows the data for the Bi-Polar classification, where I1and I2 are the bipolar inputs while Y is the output.

A symmetrical saturating linear transfer function is used as the neuronal squashing function, which makes the derivative of the transfer function as two-valued: zero or a constant.

If a pattern p is presented, then the output from the neurons of the hidden layer is given by:

$$o_i^h = f^h(w_{ij}P_j + b_i) \tag{5.1}$$

where i = 1, 2 · · · , wij =connection weigh between input p_j and neuron i, b_i = i^{th} neuron bias, $f^h(\,)$ = squashing function of the hidden layer.

Similarly, the output neuron's equation is:

$$o_i^h = f^0(w_j o_j^h + b) \tag{5.2}$$

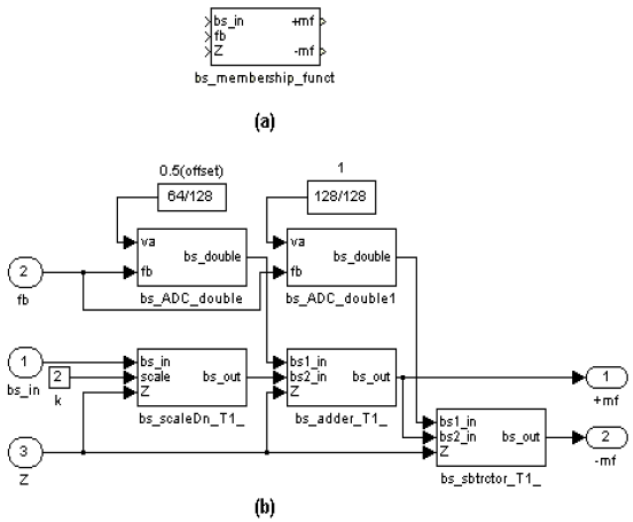

Figure 13: Bi-Valued Linear Membership Function Generator (a) Compact (b) Expanded

where $o^h j$ = activation of the hidden (j) neuron. The Neural Network Toolbox of MATLAB is used to train the network with the data of the Figure 17 which resulted in:

$$w^h = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, b^h = \begin{bmatrix} -1 \\ -1 \end{bmatrix}$$
$$w^0 = \begin{bmatrix} 1 & 1 \end{bmatrix}, b^0 = \begin{bmatrix} 1 \end{bmatrix} \quad (5.3)$$

w^h is the weight and b^h is the bias to the neurons in the hidden layer, while w^o and b^o play the identical roles for the output layer neurons. The training values are fractional numbers, very close to +1 or −1 and the problem space symmetry suggests that there is no problem in rounding the training value. Hence they are rounded to the values shown in Equation 5.3. Although the inputs in Figure 17 are
presented as binary, they could have any value from $-\infty < I_1, I_2 < \infty$ because the network is linear. Normalization makes the inputs to be constrained to $-1 < I_1, I_2 < 1$. The nonlinear saturating squashing functions are aligned to these boundaries.

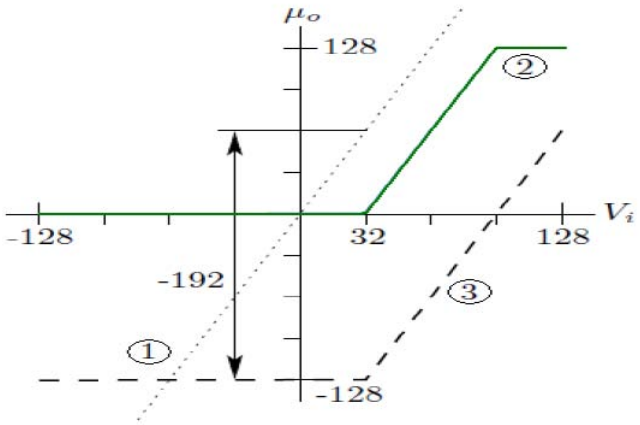

Figure 14: Transformation of a bit-stream into membership bit-stream

	Z	S_i	Addition in SAC
1	0	0	0
2	0	1	P
3	1	0	$-P$
4	1	1	0

Figure 15: Truth Table for Addition in SAC

A bit-stream system with N = 256, is limited to the inputs and outputs between a range from −127 to +128, conveniently written as ±128. Anything beyond the limits is saturated. This property of the bit-streams is used to implement the squashing function in an artificial neuron of the Bit-Stream Artificial Neural Network (BSANN). Therefore, the inputs for the BSANN are scaled to conform to:

$$-1 < I_1^{BS}, I_2^{BS} < 1$$

Figure 18 shows a Bi-Polar XOR architecture. It can be observed that three inputs are being fed to all the artificial neuron models, each of which performs the operation (wp + b). Although, the output can never exceed ±128 or±1, due to the numerical representation of the floating points, the operation (wp + b) can reach up to very large magnitudes. For example, if three numbers 60,90 and −50 are added in the bit-stream environment, then there will be two possibilities. Either to add in two steps like (60 + 90) − 50 or to add concurrently like (60 + 90- 50). The first possibility will result in 60 + 90 = 127 − 50 = 87, which is not desired, while the second possibility will result in 60 + 90 − 50 = 100, which is correct. One solution to avoid these saturation issues is to scale down the inputs by some factor in a way that their sum remains within the saturation limits, and then to scale up the summation result by the same factor at the time of addition. Since the inputs and the outputs have their maximum values as ±128, the scaling should be applied to the bias terms and the weights. The scaling factor depends on upon the maximum value, the operation (wp + b) can attain. The weights given in Equation 5.3 define the range for the BSANN XOR Network as $-3 \leq (wp + b) \leq 1$, which means the hidden layer of the BSANN would try to achieve −3 × 128. Since it is not possible, so the weights and the bias need to be reduced by a factor. This scaling can be expressed as:

$$\frac{1}{k}(wp + b) = \frac{w}{k}p + \frac{b}{k} \quad (5.4)$$

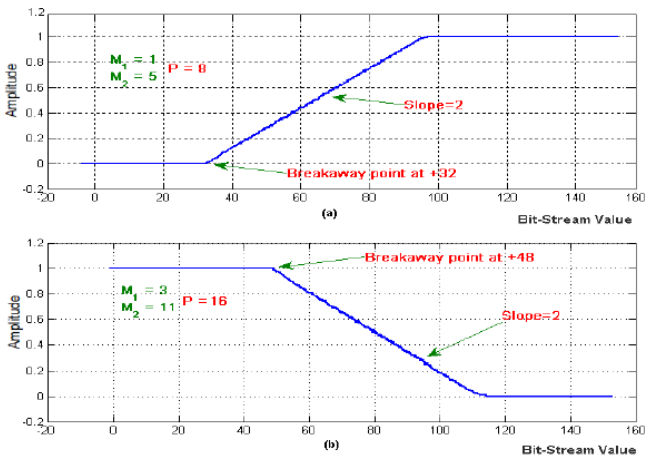

Figure 16: Liner Membership Functions (a) Positive (b) Negative

I_1	I_2	Y
-1	-1	-1
-1	1	1
1	-1	1
1	1	-1

Figure 17: Bi-Polar XOR classification

The new neuronal squashing function f '( ) with the effects of scaling can be equated with the un-scaled squashing function by:

$$f(wp+b) \leftrightarrow f'(\frac{w}{k}p+\frac{b}{k})$$

Now, a±128 saturating squashing function will be characterized with:

$$f[\frac{1}{k}(wp+b)] = \begin{cases} 128/k & \frac{1}{k}(wp+b) > 128/k; \\ -128/k & \frac{1}{k}(wp+b) < 128/k; \\ \frac{1}{k}(wp+b) & \text{for all other values} \end{cases} \quad (5.5)$$

Before discussing actual arithmetic elements, it is important to look at some primitive elements. Primitives elements are not actual bit-stream toolbox building blocks by themselves but extensively used in their construction.

The architecture in Figure 18 portray this implementation, and it can be seen that all the weights and the bias terms are scaled down by a factor of 3 using the composite scalar. However, observations show that in the case of Bi-Polar XOR, this scaling is not mandatory and an efficient network, having less number of components can be made without scaling. Hence, this reduced network has been modeled using Simulink and presented in Figure 19.

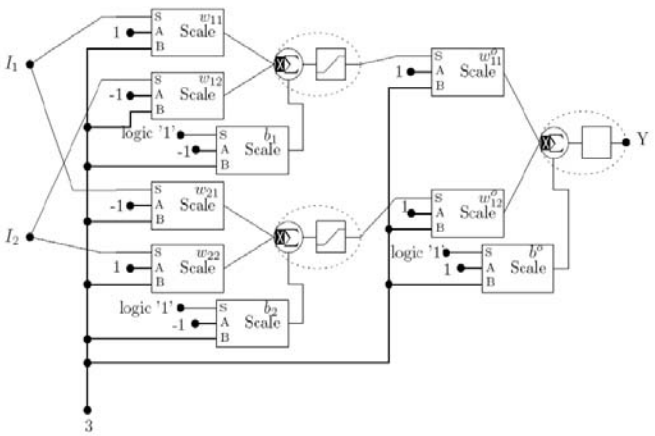

Figure 18: Conceptual Model for Bit-Stream Bi-Polar XOR

Figure 20 shows the response of this implementation along with the outcomes of the VHDL implementation of [11]. The plot Y satisfies the requirements of Bi-Polar XOR truth table.

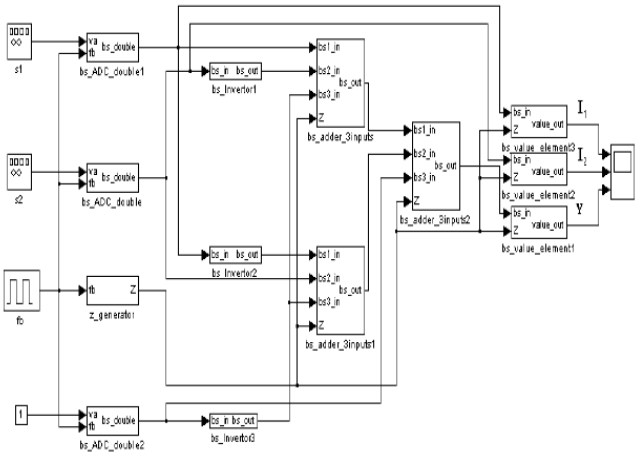

Figure 19: Simulink Model for Bit-Stream Bi-Polar XOR

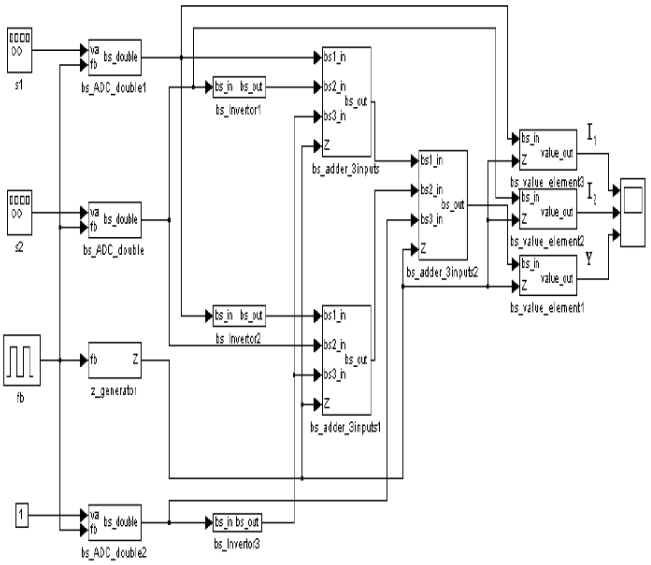

Figure 20: Bi-Polar XOR Responses (a) MATLAB (b) VHDL

VI. CONCLUSIONS

This paper presents a methodology for the simplified modeling of real-time, parallel systems in the bit-stream environment. Bit-stream solutions have been getting popularity for the hardware implementation of parallel systems as an alternative of the dedicated VLSI, especially when the project's cost or the complex structure of the systems are under consideration. A VHDL version of bit-stream components has already shown the effectiveness of the technology by implementing several systems, but the limitation of VHDL in building large and complex system models creates a hurdle to test the VHDL bit-stream components with complex systems. This limitation led to the foundation of this paper. Here, the bitstream elements have been developed in MATLAB, which provides an exceptional environment for complex system modeling. Hence, bit-stream behavior can be tested with new systems as well as with a large collection of existing system models regardless of the complexity.

REFERENCES

[1] M. S. Friedrichs, P. Eastman, V. Vaidyanathan, M. Houston, S. Legrand, A. L. Beberg, D. L. Ensign, C. M. Bruns, and V. S. Pande, "Accelerating molecular dynamic simulation on graphics processing units", Journal of Computational Chemistry, 2009

[2] S. Guo, L. Peters, and H. Surmann, "Design and application of an analog fuzzy logic controller", IEEE Transactions on Fuzzy Systems, Vol. 4, Issue, 4, pages. 429–438, 2002

[3] L. Peters, S. Guo and R. Camposano. "A novel analog fuzzy controller for intelligent sensors," Fuzzy Sets and Systems, vol. 70, pages. 235–247, 1995.

[4] T. Yamakawa, "A fuzzy inference engine in nonlinear analog mode and its application to a fuzzy logic control." vol. 4, pages. 496–522, 1993.

[5] G. Aranguren, L. A. L. Nozal, X. Basogain, J. L. Martin, and J. L. Arroyabe, "Hardware implementation of a pipeline fuzzy controller and software tools", Fuzzy Sets and Systems Featured Issue: Selected papers from ACIDCA 2000, vol. 128, Pages. 61–79, 2002.

[6] J. L. Huertas, S. Sanchez-Solano, A. Barriga, and I. Baturone, "A hardware implementation of fuzzy controllers using analog-digital VLSI techniques," Computers and Electrical Engineering, vol. 20, pages. 409–419, 1994

[7] C. Mead, "Analog VLSI and Neural Systems", Addison-Wesley Longman Publishing Co., Inc, Boston, MA, USA, 1989.

[8] M. Chiaberge, E. M. Sologuren, and L. M. Reyneri, "A pulse stream system for low-power neuro-fuzzy Computation", IEEE Transaction on Circuit and Systems I: Fundamental Theory and Applications Vol. 42, Pages. 946–954, 1995.

[9] L. M. Reyneri, "A performance analysis of pulse stream neural and fuzzy computing systems," Vol. 42, Pages. 642–660, 1995.

[10] H. Hikawa. "Frequency-based multilayer neural network with on-chip learning and enhanced neuron characteristics." IEEE Transaction on Neural Networks, vol. 10. Pages: 545–553, 1999.

[11] N. D. Patel " Bit-streams: applications in control", Ph.D. Thesis ECE, University Of Auckland, 2006.

[12] N. Patel, G. Coghill, and S. K. Nguang, " Digital realization of analogue computing elements using bit streams", In Proc. 3rd IEEE International Workshop on System-on- Chip for Real-Time Applications, Pages. 76–80, 2003.

[13] N. Patel, S. K. Nguang, and G. Coghill, " Bit-streams: Fuzzy controller implementation for motor control", In Proc. IEEE International Conference on Industrial Technology, Pages. 15–17, 2006.

[14] N. D. Patel, S. K. Nguang, G. G. Coghill, and A. K. Swain, "On-line implementation of servo controllers using bit-streams", In Proc TENCON 2005 IEEE Region 10, Pages. 1–6, 2005.

[15] N. D. Patel, S. K. Nguang, and G. G. Coghill. "Neural network implementation using bit streams",
IEEE Transactions on Neural Networks, Vol. 18, Pages. 1488–1504, 2007

[16] A. Ohta, T. Isshiki, and H. Kunieda , "New FPGA Architecture for bit-serial pipeline data path", in Proc. IEEE Symposium on FPGAs for Custom Computing Machines, pp. 58–67, 1998.

[17] S. A. Ali, "Bit-Streams-ToolBox For MATLAB", Master's Thesis ECE, University Of Auckland, 2009.

[18] W. Mcculloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity", Bulletin of Mathematical Biology, vol. 5, Issue. 6, Pages. 115-133, 1943

Transient Array Radio Telescope: Calibration and Aperture Synthesis

Max Scheel, Timothy C. A. Molteno and Phill C. Brown
Department of Physics University of Otago
Box 56, Dunedin, NEW ZEALAND
Email: max@max.ac.nz

***Abstract*—The Transient Array Radio Telescope (TART) is a 24-element aperture synthesis array radio telescope. It is designed as a test bench for imaging algorithm development, as well as a survey-instrument for transient events. We present an update for the telescope hardware, describe software for traditional radio-interferomic aperture-synthesis imaging and also explain the process of calibration. While refining the calibration is under ongoing development, the number of the imaged sources that are consistent with known sources provides some confidence in the telescopes conceptional design and operation.**

I. Overview

Interferometric measurements in radio astronomy lead back to early observations of the sun by Ryle and Vonberg in 1946 [1]. Since then several multi million dollar antenna array based telescopes such as VLA, ALMA, LOFAR and MWA have been built [2] [3] [4] [5].
These telescopes rely on an indirect synthesis process to transform the calibrated observation data from the antenna array into an image.
The calibration procedure of traditional radio telescope projects such as ALMA or VLA typically involves various calibration measurements and synthesis imaging of spatially close-by calibrator sources before, after or even during each observation of science interest.
The calibration process is essential for the imaging process and becomes more difficult with increasing number of antennas, bandwidth and larger field of view. Precursor projects to the SKA such as MWA were designed to further develop calibration techniques [6]. While the Transient Array Radio Telescope (TART) is a cost-effective downscaled radio telescope, calibration remains non-trivial as an essential part within the imaging process.
The TART is a 24-element aperture synthesis array radio telescope. It is designed as a test bench for imaging algorithm development, as well as a survey-instrument for transient events. TART operates in the L1 GPS band at $1.575\,\mathrm{GHz}$. Both hardware and software designs are open-source, released under the GNU GPLv3 license and are available at the TART website [7]. An initial design and a more detailed overview over the TART is given in [8].
Advances in hardware and calibration have been made and result in ability to apply traditional synthesis imaging techniques. In this paper we want to give an overview and a brief update of the instrumental setup and the data acquisition process. We than give an introduction in the synthesis imaging process in section IV, that only with appropriate calibration, described in section V can generate consistent results.

Fig. 1: Prototype of the TART deployed at the rural test site on Signal Hill, Dunedin. Telescope tiles with antennas on the right. The base station module as well as solar power system is placed within the metal box 15 meters away from the tiles.

II. Instrumentation

This section shall give a brief description of the components of the TART and their functionality. The TART prototype as deployed on our rural test site and depicted in figure 1 comprises of 24 active ceramic patch antennas, located on four identical tiles. A simplified system diagram can be found in figure 2. Each tile hosts six active antennas as well as a radio hub module. A solar powered base station module is located 15 meters away from the tiles.

Since the first prototype, of which further details can be found in [8], advances have been made that simplified the build of the telescope and reduced the number different printed circuit boards down to three:

A. Base Station Module

The latest version of the base station module, as shown in figure 3 is host of a power supply, a local oscillator, as well as an FPGA and a miniature ARM-based computer. At the current stage the Spartan-6 FPGA on a Papilio Pro development board is solely acquiring and temporarily storing snapshots of raw antenna data samples [9], however further capabilities are in development allowing further functionality such as the commonly used cross-correlation of the incoming data samples.
Connected to the FPGA, via Serial Peripheral Interface (SPI),

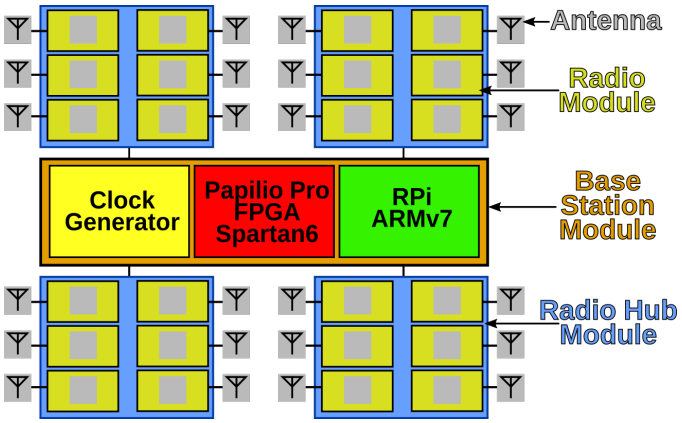

Fig. 2: System diagram of the telescope. The telescope consists of one base station module, four radio hub modules each of which is housing six radio modules and 24 antennas.

the base station module also hosts a Raspberry Pi Model 2 (RPi2). The SPI communication allows the configuration and control of the FPGA from the RPi2 as well as the transfer of the acquired raw data back to the RPi2 for further processing and storage.
The base station power supply not only provides power to all board components 3.3 V and /5 V but also distributes 24 V to each of the radio hub modules.

B. Radio Hub Module

As shown in the system diagram in figure 2, the TART consists of four radio hub modules. The purpose of the radio hub module is to provide power and a common clock signal to each radio and transfer raw data streams back to the base station module. The received clock signal from the base station is fed into a jitter cleaner and the resulting clock signal, as well as power, is distributed to six radio modules. The raw data stream of each radio module is converted to a differential signal before leading back to the base station.
Each radio hub module is connected to the base station module via two Cat 6 Ethernet cables. With a total of eight twisted pairs it is possible for the base station module to supply 24 V power and a common clock to the radio hub module, as well as receive six streams of antenna samples as differential pairs.

C. Radio Module

The radio module hosts an off-the-shelf MAX2769B radio chip which amplifies [10], mixes and filters the received electric signal and produces a stream of binary samples with a sampling frequency of $f_s = 16.368\,\mathrm{MHz}$, at a centre frequency of $f_c = 4.092\,\mathrm{MHz}$ and passband 3dB bandwidth $B = 2.5\,\mathrm{MHz}$.

III. Data Acquisition

At Each of the radios sampling frequency of 16.368 MHz leads to a total incoming raw data rate of 392.832 MHz. To get an initial handle on the data rate we only consider snapshots of consecutive samples at reduced coverage (uptime).

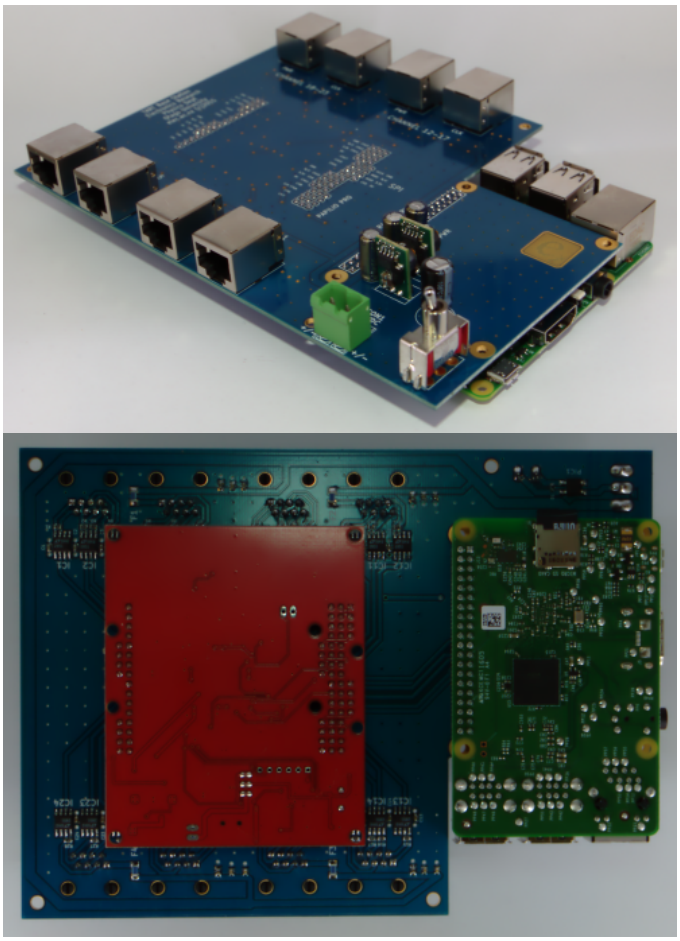

Fig. 3: Top and bottom view of base station module. The top side holds a total of eight RJ45 connectors to provide clock and power as well as receive antenna data. The bottom side (below) accommodates a Papilio Pro FPGA development board (left, red) for synchronous data acquisition and a Raspberry Pi Model 2 (right, green) for basic data processing and storage.

Once a minute the RPi2 issues the command to the FPGA to acquire a one second snapshot of 2^{24} subsequent raw data samples for each antenna. The samples are buffered in the SDRAM on the Papilio Pro. The acquired raw data is subsequently transferred via SPI to the RPi 2, where it is stored on disk and regularly uploaded to mass storage for further offline processing.

IV. Synthesis Imaging

In this section the basic principles of synthesis imaging are described. Conventional radio interferometric imaging is based on the inverse Fourier transform of the visibility function $V(u,v)$ to recover the sky brightness $I(l,m)$. This approach is based on the Van Cittert-Zernike theorem, which relates the complex visibility,

$$V(u,v,w) = \int \frac{A(l,m)I(l,m)}{n} \exp(i2\pi(lu+mv+(n-1)w))dldm, \quad (1)$$

to the sky brightness $I(l,m)$ and the primary beam response of an interferometer $A(l,m)$ [11]. Typically and mainly for computational simplicity a small field assumption, i.e. $n \approx 1$

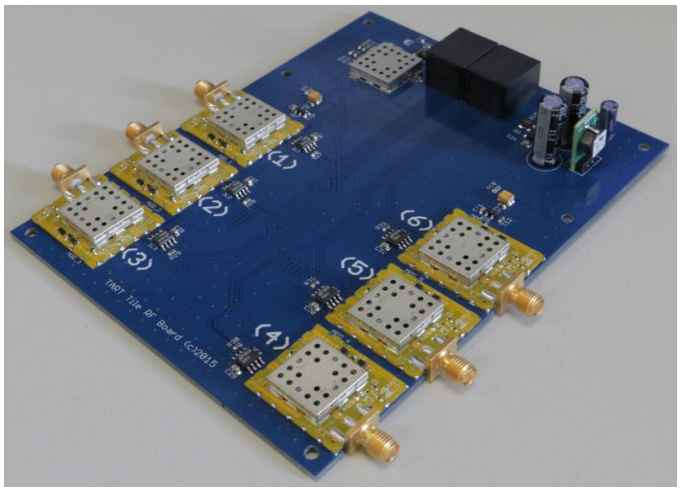

Fig. 4: Top side of the radio hub module. At the top of the board is shielded jitter cleaner as well as a DC conversion block (right). The down converted DC power and the cleaned up clock signal is provided to six radio modules (left and bottom). Generated raw data streams are converted in the differential transceivers and lead back to the base station module as differential pairs.

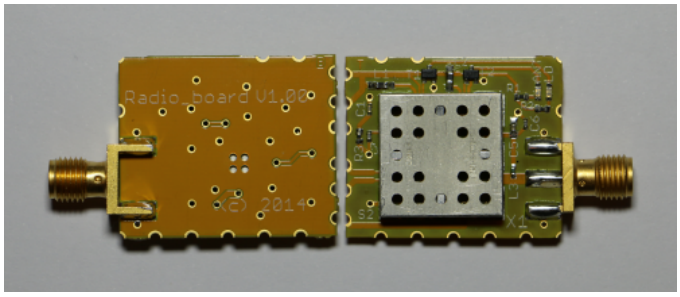

Fig. 5: Front and back side of a radio module. The radio module comprises of two main components: A shielded MAX2769B radio chip and an SMA connector to attach a powered antenna with integrated low noise amplifier.

is made. Ignoring the primary beam response this reduces 1 to a two dimensional Fourier transform

$$V(u,v) = \int I(l,m)\exp(i2\pi(lu+mv)dldm. \qquad (2)$$

In practice the visibility is measured by correlating data from pairs of antennas. The distance and orientation of pairs antennas with respect to observation phase centre are defined as baselines and are measured in wavelengths in the uv plane. For N antennas there are $N\cdot(N-1)/2$ possible pairs for correlation. For $V(u,v)$ measured for a discrete set of N points (u_i,v_i) the inverse transform can be approximated by

$$I_D(l,m) = \sum_i V(u_i,v_i)\exp(i2\pi(u_i l + v_i m), \qquad (3)$$

where $I_D(l,m)$ is called dirty image. The dirty image

$$I_D(l,m) = I(l,m) * B(l,m) \qquad (4)$$

is a convolution of the true sky brightness $I(l,m)$ with the dirty beam $B(l,m)$. The dirty beam, also known as point spread function, is the response of the telescope to a point source at the phase centre [12]. Therefore the dirty beam is calculated by setting $V(u_i,v_i)=1$ and Fourier transforming the sparsely covered uv plane. To perform the Fourier transform V is gridded on a regular rectangular grid. The transform is then carried out using FFTs. The full width at half maximum (FWHM) of the resulting dirty beam is commonly used as a standard measure for the resolution of the telescope.
As an example figure 8a depicts an image generated with aperture synthesis. The image covers a large region of $\pm 60\,^\circ$ around the zenith (straight up). Without calibration the TART is unable image the strong artificial point source that was positioned, in the far field, straight above the telescope. A first attempt for the TART that leads to a more promising image as shown in 9 is presented in the next section.

V. Calibration

This section describes the currently used basic calibration process of the TART. While calibration in the future will become increasingly more complex we start off with two major parts. One essential component in order to be able to produce meaningful results with synthesis imaging is the knowledge of the baseline lengths i.e. determine the spacial frequency components u_i and v_i of the measured visibility $V(u_i,v_i)$ in the uv plane. The baseline lengths and orientations can be derived from the position of each antenna.

A. Baseline calibration

By design each tile of the TART is identical. Hence variations in antenna positions and differences of antenna positions on a tile are small (below $1\,\mathrm{cm}$) and assumed to be negligible. An additional assumption is that all tiles are in the same plane, hence the tiles can only have offsets towards East or North and can be rotated by an angle α_i.
As a first step for baseline calibration distances between pairs of antennas of different tiles are measured with a tape measure. Secondly we create a tile based model of the array with a fixed position and orientation of the first tile (0). We than vary offsets and rotations of the remaining three tiles $(1,2,3)$ with a Nelder-Mead method [13]. Calculating antenna positions for the model, we derive a sum of squares of baseline length differences

$$\mathrm{score} = \sum (b_{ij}^{\mathrm{measured}} - b_{ij}^{\mathrm{model}})^2, \qquad (5)$$

for all measured baseline lengths $b_{ij}^{\mathrm{measured}}$ and model baseline length b_{ij}^{model} and between antenna i and j.
Minimizing the sum of squares leads to a modified solution for the antenna positions, which is shown in figure 6. While tile 1 and 2 remain close to the initial array design solution, there is an substantial offset and rotation apparent for tile 3. This particular tile experienced a drift between the initial tile placement and fixation to the ground.
The presented method yields a calibrated solution for the array geometry, however a global orientation angle of the array left to be determined until after the initial gain calibration.

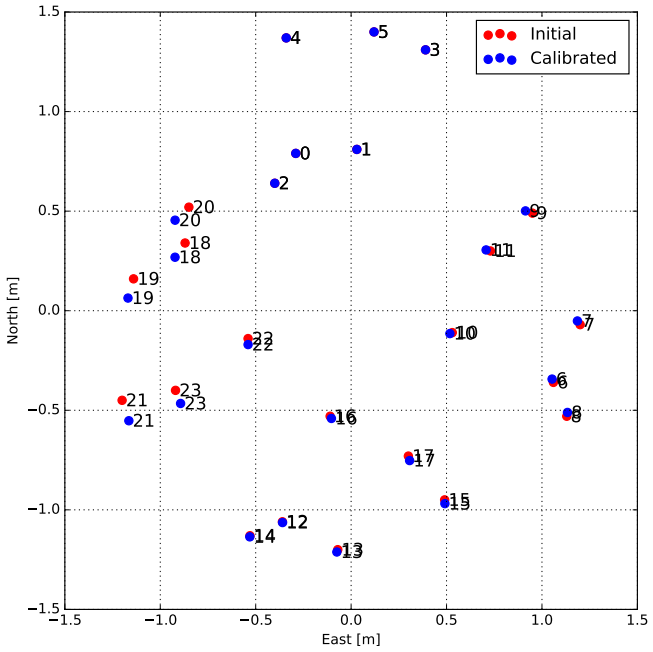

Fig. 6: Calibration of the antenna positions. Shown in red is the initial design configuration. Also shown in blue are the calibrated antenna positions. During the minimisation tile 0 (top) is fixed, hence the calibrated antenna positions remain the same as the initial positions. There is only slight adjusted in rotation and offset of tile 1 (right) and 2 (bottom). However tile 3 (left) undergoes adjustment. This tile experienced the most displacement before the installation of ground anchors.

B. Gain calibration

The observed complex visibility can be defined as

$$\widetilde{V}_{ij} = G_{ij}V_{ij} + \epsilon_{ij} + \eta_{ij}, \tag{6}$$

where G_{ij} is the complex gain for this baseline, V_{ij} is the true visibility, ϵ_{ij} is an additive error and η_{ij} is noise. At this stage we disregard additive offset errors to be part of the noise and focus only on the multiplicative part as an antenna based calibration, where

$$G_{ij} = g_i(t)g_j^*(t) = A_i(t)A_j(t)\exp(\phi_i(t) - \phi_j(t)). \tag{7}$$

Due to an arbitrary instrumental phase we set $\phi_0 = 0$. Without absolute flux density calibration we also set $A_0 = 1$
A point source with flux density S at the phase centre (zenith) has nice properties such that the visibilities simplify to $|V_{ij}| = S$ and $\arg(V_{ij}) = 0$. We therefore utilise a calibration source which acts as a strong point source, position it above the centre of telescope and elevated into far-field of the TART via quad copter. Figure 7 gives an overview of the rural test site from the quad copter at $100\,\mathrm{m}$ elevation.
A band pass filter limits the wide band noise of the BG7TBL RF noise source to $1455 - 1925\,\mathrm{MHz}$ and minimises potential interference [14].

In order to obtain a set of complex antenna gains g_i, a calibration snapshot of raw data is acquired. The complex cross correlator generates complex visibilities from the raw data. A standard BFGS algorithm based minimiser varies the solution for gain amplitudes and phases until the least squares between model visibilities and gain corrected calibration visibility data is minimal. In figure 8 is shown how the determination of the relative gains of the antennas effects the resulting image.
As an additional step the point source is relocated north of the telescope. By obtaining another calibration snapshot we than can solve for the global rotation of the telescope and update antenna position data such that the imaged source appears dear north.

Fig. 7: Rural TART telescope site as seen from the drone during calibration flight. A battery powered noise source is attached to the drone is and acts as a strong point source.

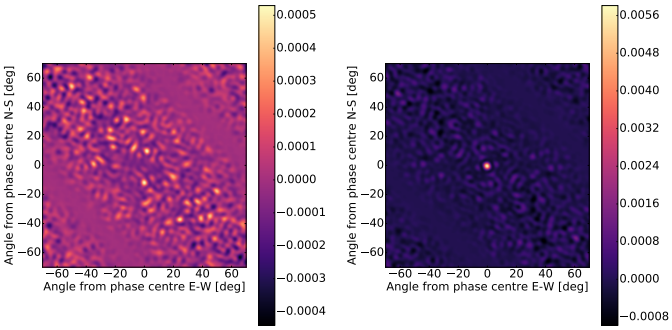

(a) Before calibration. (b) After calibration.

Fig. 8: Synthesized image of the calibration source before and after applying the antenna gain calibration. Applying the gain solution to the measured visibilities causes the point source to appear in the centre of the image and an enormous increase in contrast.

As an example figure 9 and sequence 10 show images with the applyed calibrated gain solution. The generated images show consistent bright spots of emission that propagate across the sky. When also plotting the known position of GPS satellites we can find agreement, within a region around the zenith, between imaged sky brightness and known position of a GPS satellite. Distances are slightly above the order of resolution which leaves room for further calibration improvement. For

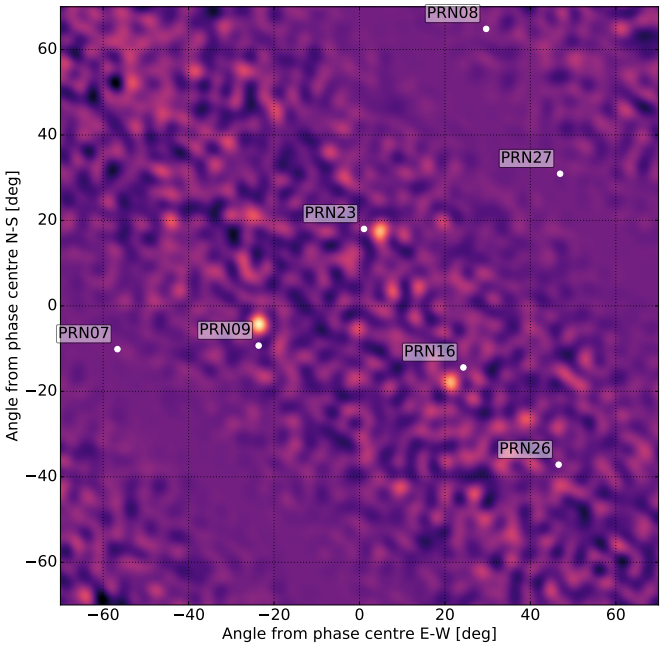

Fig. 9: TART synthesis image and known positions of GPS satellites as white dots. The phase centre of the image is facing straight up (zenith). We can observe qualitative agreement with some GPS satellites, PRN09, PRN23 and PRN16, which are within 30° of the phase centre.

large distances from the zenith the previous made assumption of a flat sky does not hold, and the simplification of the Van Cittert-Zernike theorem is no longer valid.

VI. Conclusion

Current data acquisition modality of the TART telescope allows offline processing of data and development of calibration procedures. While calibration is under ongoing development first calibration results enable traditional synthesis imaging to image sources that are qualitative consistent with known positions of GPS satellites. This provides a degree of confidence in the correct setup, operation and initial data processing pipeline.

Acknowledgment

The TART project is supported by grants from the University of Otago. We would like to thank Chris and Peter from physics workshop for building the tiles, as well as Professor Steven Tingay for helpful discussions.

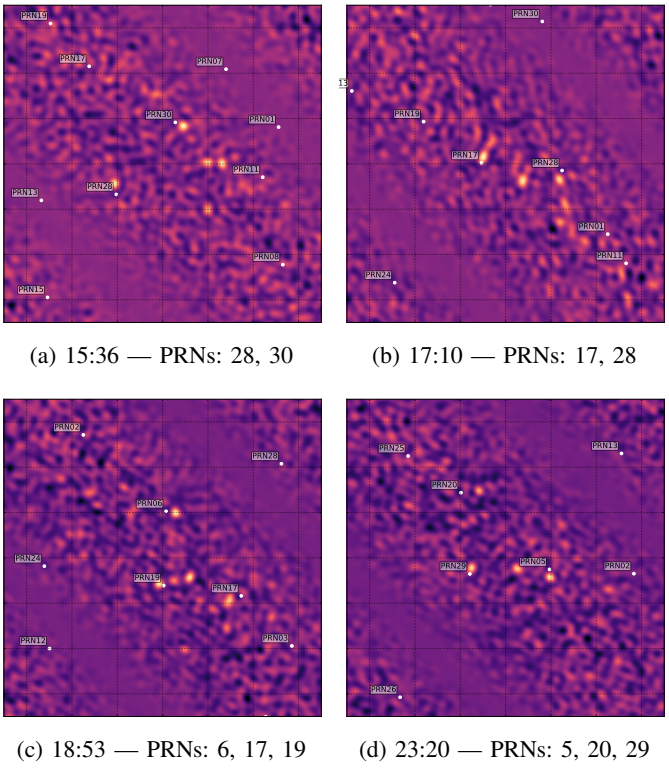

(a) 15:36 — PRNs: 28, 30

(b) 17:10 — PRNs: 17, 28

(c) 18:53 — PRNs: 6, 17, 19

(d) 23:20 — PRNs: 5, 20, 29

Fig. 10: Series of synthesis image snapshots with TART. The known positions of GPS satellites on the 6th May 2016 at different times (UTC) are shown as white dots and labelled with their pseudo-random noise number. While there is obvious emission imaged that mainly coincides with the position of GPS satellites there is also other objects consistently visible some of which can be identified as search and rescue satellites.

References

[1] M. Ryle and D. Vonberg, "Solar radiation on 175 mc./s," *Nature*, vol. 158, no. 339, p. S82, 1946.

[2] J. Condon, W. Cotton, E. Greisen, Q. Yin, R. Perley, G. Taylor, and J. Broderick, "The nrao vla sky survey," *The Astronomical Journal*, vol. 115, no. 5, p. 1693, 1998.

[3] A. Wootten, "Atacama large millimeter array (alma)," in *Astronomical Telescopes and Instrumentation*. International Society for Optics and Photonics, 2003, pp. 110–118.

[4] H. Röttgering, "Lofar, a new low frequency radio telescope," *New Astronomy Reviews*, vol. 47, no. 4, pp. 405–409, 2003.

[5] A. Beardsley, B. Hazelton, M. Morales, R. Capallo, R. Goeke, D. Emrich, C. Lonsdale, W. Arcus, D. Barnes, G. Bernardi *et al.*, "A new layout optimization technique for interferometric arrays, applied to the murchison widefield array," *Monthly Notices of the Royal Astronomical Society*, vol. 425, no. 3, pp. 1781–1788, 2012.

[6] D. A. Mitchell, L. J. Greenhill, R. B. Wayth, R. J. Sault, C. J. Lonsdale, R. J. Cappallo, M. F. Morales, and S. M. Ord, "Real-Time Calibration of the Murchison Widefield Array," *Radio Sci.*, vol. 44, p. 0A01, 2009.

[7] (2016) Transient Array Radio Telescope Project Website. [Online]. Available: https://github.com/tmolteno/TART/

[8] C. Shaw, M. Scheel, and T. Molteno, "Transient Array Radio Telescope: Design," *21st Electronics New Zealand Conference*, 2014.

[9] (2016) Gadget Factory Website - Papilio Pro. [Online]. Available: http://papilio.cc/index.php?n=Papilio.PapilioPro

[10] (2012) MAX2769B Universal GPS Receiver 19-5875; Rev 2; 10/12. [Online]. Available: https://datasheets.maximintegrated.com/en/ds/MAX2769B.pdf

[11] A. Thompson, J. Moran, and G. Swenson, *Interferometry and Synthesis in Radio Astronomy*. Wiley, 2008. [Online]. Available: http://books.google.co.nz/books?id=AwBN5bpuEU0C

[12] R. Sault and T. Oosterloo, "Imaging algorithms in radio interferometry," 2007.

[13] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The computer journal*, vol. 7, no. 4, pp. 308–313, 1965.

[14] (2016) Mini-Circuits Coaxial Bandpass Filter - VBFZ-1690+. [Online]. Available: http://www.minicircuits.com/pdfs/VBFZ-1690+.pdf

On a representation of feedback shift-registers

Daniel Schumayer and Timothy C. A. Molteno
Department of Physics, University of Otago,
730 Cumberland Street, Dunedin 9016
email: daniel.schumayer@otago.ac.nz

***Abstract*—Feedback shift-registers (FSRs) are fundamental building blocks in specialised digital electronics, e.g., cryptographic or counter circuits. FSRs have exponentially large number of possible states and their mathematical description is well developed. Constructing maximal length sequences for linear FSRs can be done algorithmically. However, the analysis of nonlinear FSRs is in its infancy. In the last decade a novel algebraic structure has emerged in the field of Boolean networks, which may provide a useful tool to describe FSRs. Here we explain the main ideas of this new formalism via examples and show its capabilities on a particular NFSR.**

I. Introduction and motivation

Logic circuits can be classified into two types: combinational and sequential. The output of the former type depends only on the instantaneous inputs, while the behaviour of sequential type uses the past inputs as well, i.e. it has "memory". At the heart of sequential digital circuits is a bi-stable component whose state remains stable until an appropriate trigger appears on the inputs. The term *register* designates a synchronous system capable of storing and emitting a binary vector according to external control signals. In sequel we restrict the meaning of a register to storing a single binary value, 0 or 1.

*Shift-register*s are circuits for storing and shifting a number of binary values, possibly with the added complexity of manipulating these values before shifting, however, the meaning of "manipulation" can be numerous. Linear Feedback Shift-Registers (LFSRs) are devices for generating sequences of binary vectors controlled by a feedback function f. The research on LFSRs started in the 1960s [1], [2] and continued actively for many decades. By today it is mature and well-understood subject: synthesis of maximum period LFSRs is straightforward from a primitive minimal polynomial over GF[2] [3], [4]. Statistical properties of sequences of LFSRs have been characterised based on Golomb's postulates [2], [5].

Recently nonlinear feedback shift-registers (NLFSR) have received substantial attention. Their linear complexity exceeds that of LFSRs of the same order [6]. While the cycle structure of LFSRs is explored [1], [2], the analogous investigations for NLFSRs are still searching for answers to fundamental questions: is there a maximal possible length cycle for any number of registers? If the answer is negative, how to decide whether or not an n-register NLFSR has such cycle? If there is, how to construct it and how many feedback functions can generate the maximal cycle? What are the statistical properties of the maximal cycles? What complexity do they have? Some partial results are available; e.g., a special class of NLFSRs with maximal period can be constructed [7], [8].

II. Nomenclature and definitions

Figure 1 represents a typical shift-register. The output of the shift-register is the sequence of states appearing at the output. The circuit is controlled by a clock signal, which influences only the speed of the FSRs but not the states of registers.

Definition 1 (n-stage feedback shift-register)**.** *Let D denote a finite set of possible states of a register, and $f : D^n \to D$ be a function. An n-stage feedback shift-register, induced by the feedback function f, is defined by the transition function $T_f : D^n \to D^n$*

$$T_f : \big(x_1, x_2, \cdots, x_n\big) \mapsto \big(x_2, x_3, \cdots, f(x_1, x_2, \cdots x_n)\big).$$

where $x_j \in D$ for $j = 1, 2, \ldots n$.

This definition is general and assumes little about D, f or T_f. In most cases, however, D is considered to be a finite field with two operations $+$, and $\cdot$. In the Boolean case $D = \{\mathrm{T}, \mathrm{F}\}$, while $+$ and $\cdot$ stand for disjunction and conjunction, respectively. In binary arithmetic $+$ is the modulo-2 addition and it is distinguished by using $\oplus$.

An alternative formulation of the shift-registers focuses solely on the sequence, $S = \{s_k\}_{k=1}^{\infty}$, observed on the output where $s_k \in D$. The first n terms are fixed by initialisation; the remaining terms are determined by the feedback function:

$$s_{n+i} = f(s_i, s_{1+i}, \cdots, s_{n-1+i}) \qquad (i = 1, 2 \ldots).$$

The output sequence is periodic if there are integers $r > 0$ and $0 \le n_0$ such that $s_{n+r} = s_n$ for all $n > n_0$. The smallest such r is called the *period* of the FSR.

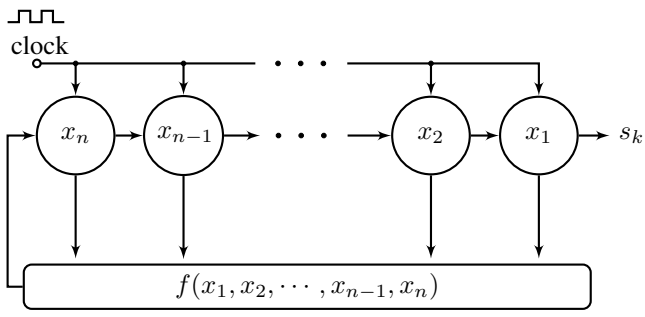

Fig. 1. Schematics of a generic n-stage shift-register is depicted. Each circle represents a register, identified by label k, whose state, $x_k \in D$. All registers are connected to a common clock signal which triggers the shift of register state x_k to x_{k-1} ($k = n, \ldots, 2$), driving x_1 to the output s_k, and also the calculation of the next state of x_n from the feedback function, f.

If T_f is a linear transformation, the feedback shift-register is called linear, otherwise it is non-linear. The recurrence relation for a linear feedback shift-register is particularly simple

$$s_{n+i} = \sum_{j=1}^{n} c_j s_{i+j-1} \qquad (i = 1, 2, \dots), \tag{1}$$

where c_j are fixed constants. It is clear that the initial state $(s_1, s_2, \dots, s_n) = (0, 0, \dots, 0)$ results in $s_{n+1} = 0$, thus S is the constant zero sequence and it doesn't generate a varying sequence. Since the number of different n-tuples from D is $|D|^n$, the period of the longest possible non-recurrent sequence is $|D|^n - 1$. The sequence achieving this period is called *maximal sequence*, or simply *m-sequence*.

Although the construct of f or F provide seemingly quite limited freedom, it is surprising how prolific these functions are. For example, f possesses $|D|^n$ different inputs each of which can be mapped to $|D|$ different symbols, thus there are $|D|^{|D|^n}$ different functions. With a similar argument one may conclude that the number of possible mappings is $(|D|^n)^{|D|^n}$, while the number of one-to-one mappings is $|D|^n!$. Here $|\ |$ is the cardinality of set D and ! stands for the factorial.

Example 1 (Binary linear feedback shift-register)**.** The most important class of shift-registers uses $D = \{1, 0\}$. Let us consider a 2-stage feedback shift-register with feedback functions: $f_1(x_1, x_2) = x_2$, and $f_2(x_1, x_2) = x_1 \vee x_2$, where $\vee$ denotes disjunction.

TABLE I
TABULATED VALUES OF TWO, BINARY FEEDBACK FUNCTIONS.

s_1	s_2	$s_3 = f_1(s_1, s_2)$	$s_3 = f_2(s_1, s_2)$
0	0	0	0
0	1	1	1
1	0	0	1
1	1	1	1

It is apparent that a FSR equipped with f_1 won't generate maximal sequences, since irrespectively of the initialisation its period will be unity. ◁

III. ALGEBRAIC MODELLING

In order to capture the essence of feedback shift-registers we need mathematical tools, preferably tools which allow fast calculation. Definition 1 has set the framework and points directly to linear algebra [9]. However, before one can even attempt to describe the time evolution of a FSR, it is necessary to represent *all* logical functions somehow.

Representing linear logical functions is straightforward as the mapping can naturally be represented by a matrix over the basis formed by the individual primitive logical variables. This structure, however, cannot directly be carried over to nonlinear logical functions. Nevertheless, it is a possibility that, by introducing new variables, the original logical function can be linearised. This is the main idea behind the novel mathematical formalism introduced by Daizhan Cheng and his collaborators nearly decade ago [10], [11].

Let p_1, p_2, $\dots p_n$ be primitive variables and f be formally a function of these variables. By introducing all possible conjunctions of the primitive variables one can create an extended set of variables, P, over which *all* f becomes linear! The price we pay for linearising all functions is exponential expansion in the dimensionality of the basis. Increasing the dimensionality of a vector space isn't new in mathematics or physics, as different tensor products, e.g., Kronecker, Tracy-Singh or Khatri-Rao, are prevalent in linear algebra, statistics, data analysis, quantum mechanics and relativity. However, it seems natural to increase the dimensionality of the space only if it is necessary [10].

Definition 2 (Left semi-tensor product)**.** *Let A be an $m \times n$, and B be a $p \times q$ matrix, while $\alpha = \mathrm{lcm}(n, p)$. The left semi-tensor product of A and B is defined as*

$$A \ltimes B = \left(A \otimes I_{\frac{\alpha}{n}}\right)\left(B \otimes I_{\frac{\alpha}{p}}\right)$$

where $\otimes$ is the Kronecker-product of matrices, I_d is the d-dimensional identity matrix, and the operation between the objects in parentheses is the standard matrix product.

If $n = p$, this semi-tensor product coincides with the standard matrix product of A and B. The semi-tensor product inherits most of its properties from the Kronecker-product:

distributivity

$$(\alpha A + \beta B) \ltimes (\gamma G) = \alpha\gamma A \ltimes G + \beta\gamma B \ltimes G$$
$$\alpha A \ltimes (\gamma G + \delta D) = \alpha\gamma A \ltimes G + \alpha\delta A \ltimes D$$

associativity

$$(A \ltimes B) \ltimes C = A \ltimes (B \ltimes C)$$

where α, β, γ, δ are scalars, while A, B, G and D are matrices. Listing properties of the semi-tensor product would stretch the extension of this article. We rather introduce the necessary results via examples.

Example 2 (Permutation matrix for a vector)**.** Organising the basis we need an operation which can rearrange the elements of vectors or matrices. Let the elements of a vector be labelled by a multi-index $\pi = (\pi_1, \pi_2, \dots, \pi_n)$. We wish to rearrange the elements according to a different multi-index, $\widetilde{\pi} = (\widetilde{\pi}_1, \widetilde{\pi}_2, \dots, \widetilde{\pi}_n)$. The corresponding permutation matrix, $W_{[\pi,\widetilde{\pi}]}$, is defined as an $n \times n$ square matrix with elements

$$w_{i,j} = \begin{cases} 1 & \pi_i = \widetilde{\pi}_j \\ 0 & \text{otherwise} \end{cases} \qquad \text{for } i, j = 1, \dots, n.$$

Let $v = (v_1, v_2, v_3)$ whose elements we wish to rearrange to $\widetilde{v} = (v_3, v_1, v_2)$, i.e. $\pi = (1, 2, 3)$ and $\widetilde{\pi} = (3, 1, 2)$. Then

$$v\, W_{[\pi,\widetilde{\pi}]} = (v_1, v_2, v_3) \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} = (v_3, v_1, v_2) = \widetilde{v}$$

The permutation matrix is defined in an analogous manner for column vectors. ◁

Example 3 (Permutation matrix for a Kronecker-product)**.** Let $u = (u_1, u_2, \dots, u_m)$ and $v = (v_1, v_2, \dots, v_n)$. Direct calculation shows

$$u \otimes v = (u_1v_1, u_1v_2, \dots, u_1v_n, u_2v_1, \dots, u_mv_n)$$
$$v \otimes u = (u_1v_1, u_2v_1, \dots, u_mv_1, u_1v_2, \dots, u_mv_n).$$

As only the order of entries are different, there must be a permutation matrix which shuffles one product into the other. Indicating the dimensions of each vectors is sufficient to uniquely identify the permutation matrix, e.g., $W_{[n,m]}$ for which $(u \otimes v)W_{[n,m]} = v \otimes u$.

Consider two row-vectors, $u = (1, 2, 3)$ and $v = (2, 0)$ and the transformation of the possible Kronecker-products of u and v into each other is

$$(u \otimes v)W_{[2,3]} = (2\,0\,4\,0\,6\,0)\begin{pmatrix} 1&0&0&0&0&0\\ 0&0&0&1&0&0\\ 0&1&0&0&0&0\\ 0&0&0&0&1&0\\ 0&0&1&0&0&0\\ 0&0&0&0&0&1 \end{pmatrix} \tag{2}$$
$$= (2\,4\,6\,0\,0\,0) = v \otimes u$$

$W_{[n,m]}$ can be defined for column vectors too. ◁

Example 4 (Exchange of two vectors)**.** The example above shows that $u \ltimes v$ and $v \ltimes u$ can be transformed into each other, if u and v are vectors of the same length: $v \ltimes u = W_{[2,2]}\, u \ltimes v$. ◁

Example 5 (Exchange of a column vector and a matrix)**.** Here we consider the specific case of exchanging a column vector with a matrix. Let $u \in M_{p\times 1}$ and $A \in M_{m\times n}$ and examine their semi-tensor product, $u \ltimes A$. It can be proven by using Definition 2 that $u \ltimes A = (I_p \otimes A) \ltimes u$. ◁

Theorem 1 (Representation of logical functions)**.** *Any n-ary logical function $f(x_1, x_2, \cdots, x_n)$ has a unique representation in the form of*

$$f(x_1, x_2, \cdots, x_n) = M_f \ltimes x_1 \ltimes x_2 \ltimes \cdots \ltimes x_n,$$

where $x_i \in D$, $i = 1, 2, \dots, n$.

The proof can be found in Ref. [11]. If $k = |D|$, then M_f is a $k \times k^n$ logical matrix and it is called the *structure matrix* of f. Since M_f is a logical matrix, i.e., every column contain one and only one non-zero element whose value is unity. In compact notation: $M_f = \delta_k[c_1, c_2, \dots c_{k^n}]$, where c_i denotes the c_i-th column of the k-dimensional identity matrix. For example, $W_{[2,3]}$ in (2) can be written as $\delta_6[1, 3, 5, 2, 4, 6]$.

Example 6 (Semi-tensor product of Boolean variables)**.** Let us choose a redundant, vector representation $\boldsymbol{p}$ of a Boolean variable $p \in D = \{\mathrm{T}, \mathrm{F}\}$ in the following way

$$p \mapsto \boldsymbol{p} = \begin{pmatrix} p \\ \neg p \end{pmatrix} \text{ thus, e.g., } \mathrm{T} \mapsto \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } \mathrm{F} \mapsto \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Consequently their Boolean product $p \wedge q$ is associated with the semi-tensor product

$$\boldsymbol{p} \ltimes \boldsymbol{q} = \begin{pmatrix} p \\ \neg p \end{pmatrix} \ltimes \begin{pmatrix} q \\ \neg q \end{pmatrix} = \begin{pmatrix} pq \\ p\neg q \\ (\neg p)q \\ (\neg p)(\neg q) \end{pmatrix} \to \begin{pmatrix} 11 \\ 10 \\ 01 \\ 00 \end{pmatrix}.$$

If $p = q = \mathrm{T}$, the entries of $\boldsymbol{p} \ltimes \boldsymbol{q}$ are the binary representation of numbers from 0 to 2^2-1 but in reverse alphabetic order. ◁

Example 7 (Boolean variables from semi-tensor product)**.** One might wonder about the inverse operation of the semi-tensor product, i.e., how to determine the values of individual Boolean variables provided the numerical value of their semi-tensor product is known. Let us assume that $\boldsymbol{x}_1 \ltimes \cdots \ltimes \boldsymbol{x}_n = \delta^j_{2^n}$, thus the index of the only nonzero element is j (in base-10 numeral system) while it corresponds to the binary representation of number $(2^n - j)_{10}$. Therefore conversion of $(2^n - j)_{10}$ into base-2 numeral system provides the value for each Boolean variables. For example, let $x_1 = x_3 = x_4 = \mathrm{F}$ and $x_2 = x_5 = \mathrm{T}$. The semi-tensor product $\boldsymbol{x}_1 \ltimes \cdots \ltimes \boldsymbol{x}_5 = \delta^{23}_{2^5}$, thus $j = 23$, while the corresponding label is $2^5 - 23 = 9 = 01001_2$, and indeed the second and fifth digit of this binary number is 1 while all others are zero. ◁

Example 8 (Representation of negation)**.** In Boolean algebra the number of unary operations is $2^{2^1} = 4$ which are the identity, negation, tautology, and contradiction. Only negation ($\neg$) is non-trivial, and its structure matrix is $M_\neg = \delta_2[2, 1]$. Indeed, by direct substitution

$$\neg p \mapsto M_\neg \ltimes \boldsymbol{p} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \ltimes \begin{pmatrix} p \\ \neg p \end{pmatrix} = \begin{pmatrix} \neg p \\ p \end{pmatrix}.$$

It is clear that $M_\neg \ltimes M_\neg = I_2$, i.e., negation is idempotent. ◁

Example 9 (Representation of conjunction)**.** The unary case is straightforward and uses only standard matrix-vector product, thus it isn't transparent why the semi-tensor product can be useful. Let us consider a binary logical operation: conjunction ($\wedge$). Looking at the representation of a binary product of two Boolean variables in Example 6 one may conclude that $M_\wedge = \delta_2[1, 2, 2, 2]$. Let us check this conclusion by representing $f(p, q) = p \wedge q$ using the new formalism

$$M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{q} = \begin{pmatrix} 1&0&0&0 \\ 0&1&1&1 \end{pmatrix} \ltimes \begin{pmatrix} p \\ \neg p \end{pmatrix} \ltimes \begin{pmatrix} q \\ \neg q \end{pmatrix}$$
$$= \begin{pmatrix} pq \\ p\neg q + (\neg p)q + (\neg p)(\neg q) \end{pmatrix} = \begin{pmatrix} pq \\ \neg(pq) \end{pmatrix}$$

The second component contains all products of p, q and their complements except pq, thus their logical sum must be equal to $\neg(pq)$. In a similar way one may verify that the structure matrix of disjunction is $M_\vee = \delta_2[1, 1, 1, 2]$. ◁

Example 10 (de Morgan's law)**.** The de Morgan's laws are central theorems of propositional logic, one of which states

$$\neg(p \wedge q) = (\neg p) \vee (\neg q).$$

As we have already assigned matrices to the logical operations $\wedge$, $\vee$, and $\neg$, we can check whether the new formalism follows this law or not. Expressing both sides via semi-tensor products:

$$\neg(p \wedge q) \mapsto M_\neg \ltimes (M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{q})$$
$$(\neg p) \vee (\neg q) \mapsto M_\vee \ltimes (M_\neg \ltimes \boldsymbol{p}) \ltimes (M_\neg \ltimes \boldsymbol{q})$$
$$= M_\vee \ltimes M_\neg \ltimes (I_2 \otimes M_\neg) \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}.$$

Substituting $M_\wedge$, $M_\vee$, and $M_\neg$ into these expressions gives

$$M_\neg \ltimes M_\wedge = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

and

$$M_\vee \ltimes M_\neg \ltimes (I_2 \otimes M_\neg) =$$
$$= \begin{pmatrix} 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \ltimes \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \ltimes \left[\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{pmatrix} \right]$$
$$= \begin{pmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

The matrices are equal demonstrating that the new formalism naturally follows de Morgan's law. Similarly one can verify the adherence to the law of negation of disjunction. ◁

Example 11 (Power reduction of Boolean variables)**.** We have shown how logical operators are represented with matrices. As the same variable may appear multiple times in a logical expression, its associated vector also appears in the corresponding algebraic structure. However, the original Boolean expression has a natural basis which doesn't contain the same variable many times, therefore in the algebraic representation there must be a way to eliminate superfluous variables.

As an example, if p is a Boolean variable, then $p \wedge p = p$, i.e., the power of p can be decremented. The corresponding algebraic structure has to capture this property. Let $M_\downarrow$ denote the structure matrix of decrementing the "square" of $\boldsymbol{p}$, or in formula $\boldsymbol{p} \ltimes \boldsymbol{p} = M_\downarrow \ltimes \boldsymbol{p}$. Then any power of $\boldsymbol{p}$ can be reduced to a linear expression in $\boldsymbol{p}$:

$$\ltimes_{j=1}^{K} \boldsymbol{p} = \overbrace{\boldsymbol{p} \ltimes \boldsymbol{p} \ltimes \cdots \ltimes \boldsymbol{p}}^{K \text{ times}} = M_\downarrow \ltimes \overbrace{\boldsymbol{p} \ltimes \cdots \ltimes \boldsymbol{p}}^{(K-1) \text{ times}} = \cdots$$
$$= \left(\ltimes_{j=1}^{K-1} M_\downarrow \right) \ltimes \boldsymbol{p} = M_\downarrow^{K-1} \ltimes \boldsymbol{p},$$

where $M_\downarrow^{K-1}$ stands for the semi-tensor product of $(K-1)$ copies of $M_\downarrow$. The result is general as $\boldsymbol{p}$ could even be a compound variable. For the sake of transparency we attach a leading subscript to $M_\downarrow$ indicating the number primitive Boolean variable the matrix is operating on, e.g., ${}_3M_\downarrow$.

Let us now examine the case of a single variable. The logical statement $p \wedge p = p$ applies directly and we can translate both sides of this equation into the semi-tensor product formalism: $M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{p} = {}_1M_\downarrow \ltimes \boldsymbol{p}$. In details

$$M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{p} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} pp \\ p\neg p \\ (\neg p)p \\ (\neg p)(\neg p) \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} p \\ 0 \\ 0 \\ \neg p \end{pmatrix}$$
$$= \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} p \\ \neg p \end{pmatrix} = {}_1M_\downarrow \ltimes \boldsymbol{p}. \tag{3}$$

Thus we have found that ${}_1M_\downarrow = \delta_4[1, 4]$. Analogously we may examine the product of a compound Boolean variable p with itself. Let us consider $p = \wedge_{i=1}^n p_i$, where p_i are all primitive Boolean variables. The logical $\wedge$ operation is idempotent and commutative, thus $p \wedge p = p$ stands for compound variables as well. Therefore

$$\begin{aligned} \boldsymbol{p} \ltimes \boldsymbol{p} &= (\boldsymbol{p}_1 \ltimes \boldsymbol{p}_2 \ltimes \cdots \boldsymbol{p}_n) \ltimes (\boldsymbol{p}_1 \ltimes \boldsymbol{p}_2 \ltimes \cdots \boldsymbol{p}_n) \\ &= W_{[2,2^n]} \ltimes \boldsymbol{p}_1 \ltimes \boldsymbol{p}_1 \ltimes (\boldsymbol{p}_2 \ltimes \cdots \boldsymbol{p}_n) \ltimes (\boldsymbol{p}_2 \ltimes \cdots \boldsymbol{p}_n) \\ &= W_{[2,2^n]} \ltimes {}_1M_\downarrow \ltimes \boldsymbol{p}_1 \ltimes (\boldsymbol{p}_2 \ltimes \cdots \boldsymbol{p}_n) \ltimes (\boldsymbol{p}_2 \ltimes \cdots \boldsymbol{p}_n) \\ &= W_{[2,2^n]} \ltimes {}_1M_\downarrow \ltimes \boldsymbol{p}_1 \ltimes W_{[2,2^{n-1}]} \ltimes \boldsymbol{p}_2 \ltimes \boldsymbol{p}_2 \ltimes \\ &\qquad (\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n)(\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n) \\ &= W_{[2,2^n]} \ltimes {}_1M_\downarrow \ltimes \boldsymbol{p}_1 \ltimes W_{[2,2^{n-1}]} \ltimes {}_1M_\downarrow \ltimes \boldsymbol{p}_2 \ltimes \\ &\qquad (\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n)(\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n) \\ &= W_{[2,2^n]} \ltimes {}_1M_\downarrow \ltimes (I_2 \otimes W_{[2,2^{n-1}]}) \ltimes (I_2 \otimes {}_1M_\downarrow) \ltimes \\ &\qquad \boldsymbol{p}_1 \ltimes \boldsymbol{p}_2 \ltimes (\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n)(\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n). \end{aligned}$$

For transparency and to emphasise the hierarchy of the expression, the first two terms in the product can be written in an alternative way. Using the fact that $1 = I_1$ we can continue the equation above as

$$= (I_1 \otimes W_{[2,2^n]}) \ltimes (I_1 \otimes {}_1M_\downarrow) \ltimes (I_2 \otimes W_{[2,2^{n-1}]}) \ltimes (I_2 \otimes {}_1M_\downarrow)$$
$$\ltimes \boldsymbol{p}_1 \ltimes \boldsymbol{p}_2 \ltimes (\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n)(\boldsymbol{p}_3 \ltimes \cdots \boldsymbol{p}_n).$$

The structure is now apparent and we indeed confirmed that $\boldsymbol{p} \ltimes \boldsymbol{p} = {}_nM_\downarrow \ltimes \boldsymbol{p}$, where

$${}_nM_\downarrow = \ltimes_{j=1}^{n-1} (I_j \otimes W_{[2,2^{n-j+1}]}) \ltimes (I_j \otimes {}_1M_\downarrow). \tag{4}$$

Constructing ${}_nM_\downarrow$ for a compound variable is thus possible by using matrices discussed earlier: I_j, $W_{[2,2^{n-j+1}]}$, and ${}_1M_\downarrow$. At this stage the need for ${}_nM_\downarrow$ might seem to be non-existent, but this very operation guarantees that any logical operation can be expressed in the basis $\ltimes_{j=1}^n \boldsymbol{p}_i$. ◁

Remark 1 (Explicit form of ${}_nM_\downarrow$)**.** The reduction matrix ${}_nM_\downarrow$ for n primitive variables (4) can also be given explicitly. After introducing the notation $d = 2^n + 1$, one may obtain

$${}_nM_\downarrow = \delta_{2^{2n}}[1, 1+d, 1+2d, 1+3d, \ldots, 2^{2n}]. \tag{5}$$

This form is valid even for $n = 1$, e.g., ${}_1M_\downarrow = \delta_{2^{2\cdot 1}}[1, 1 + 2^1 + 1] = \delta_4[1, 4]$ which is identical to expression (3).

The number of columns in ${}_nM_\downarrow$ is $(2^{2n} - 1)/d = 2^n - 1$. As the first element doesn't contain d, the total number of columns is 2^n. Thus ${}_nM_\downarrow$ is a $2^{2n} \times 2^n$ matrix.

Example 12 (Representation of exclusive OR)**.** Later, when Boolean functions are represent in the semi-tensor formalism, we will need the structure matrix for the exclusive disjunction (XOR) Boolean operation which can be expressed

$$p \oplus q = (p \vee q) \wedge (\neg(p \wedge q)).$$

$$
\begin{aligned}
f(p,q) &= [(\neg p)\wedge q]\vee[p\wedge(\neg q)]\vee[p\wedge q] = z_1\vee z_2\vee z_3\\
&\mapsto M_\vee \ltimes M_\vee \ltimes \boldsymbol{z}_1 \ltimes \boldsymbol{z}_2 \ltimes \boldsymbol{z}_3\\
&= M_\vee \ltimes M_\vee \ltimes \left[M_\wedge \ltimes M_\neg \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}\right] \ltimes \left[M_\wedge \ltimes (I_2\otimes M_\neg) \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}\right] \ltimes \left[M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}\right]\\
&= M_\vee \ltimes M_\vee \ltimes M_\wedge \ltimes M_\neg \ltimes (I_4\otimes M_\wedge) \ltimes \boldsymbol{p} \ltimes \boldsymbol{q} \ltimes (I_2\otimes M_\neg) \ltimes \boldsymbol{p} \ltimes \boldsymbol{q} \ltimes M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}\\
&= M_\vee \ltimes M_\vee \ltimes M_\wedge \ltimes M_\neg \ltimes (I_4\otimes M_\wedge) \ltimes (I_8\otimes M_\neg) \ltimes (I_{16}\otimes M_\wedge) \ltimes \boldsymbol{p} \ltimes \boldsymbol{q} \ltimes \boldsymbol{p} \ltimes \boldsymbol{q} \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}\\
&= M_\vee \ltimes M_\vee \ltimes M_\wedge \ltimes M_\neg \ltimes (I_4\otimes M_\wedge) \ltimes (I_8\otimes M_\neg) \ltimes (I_{16}\otimes M_\wedge) \ltimes {}_2M_\downarrow^3 \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}\\
&= M_f\, \boldsymbol{p} \ltimes \boldsymbol{q}
\end{aligned}
$$

Fig. 2. Step-by-step calculation for Example 13.

The latter form is suitable for semi-tensor product representation:

$$
\begin{aligned}
p\oplus q &\mapsto M_\wedge \ltimes (M_\vee \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}) \ltimes (M_\neg \ltimes (M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}))\\
&= M_\wedge \ltimes M_\vee \ltimes (I_4\otimes M_\neg) \ltimes (I_4\otimes M_\wedge) \ltimes {}_2M_\downarrow \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}\\
&= M_\oplus \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}.
\end{aligned}
$$

Since $M_\wedge$, $M_\vee$, $M_\neg$, and ${}_2M_\downarrow$ have already been determined, one can conclude that $M_\oplus = \delta_2[2,1,1,2]$. ◁

Example 13 (Optimising a binary Boolean function)**.** Consider the binary Boolean function: $f(p,q) = [(\neg p)\wedge q]\vee[p\wedge(\neg q)]\vee[p\wedge q]$. This expression can be simplified as $f(p,q) = [(\neg p)\wedge q]\vee p\vee(q\vee\neg q) = [(\neg p)\wedge q]\vee p = p\vee q$. Could we arrive at the shorter (optimised) Boolean expression using the semi-tensor formalism?

For sake of clarity z_1, z_2, and z_3 are introduced for each terms. The original definition of f then reads as $M_\vee \ltimes (M_\vee \ltimes \boldsymbol{z}_1 \ltimes \boldsymbol{z}_2) \ltimes \boldsymbol{z}_3$. The terms themselves are

$$
\begin{aligned}
z_1 = (\neg p)\wedge q \quad &\mapsto \quad \boldsymbol{z}_1 = M_\wedge \ltimes M_\neg \ltimes \boldsymbol{p} \ltimes \boldsymbol{q},\\
z_2 = p\wedge(\neg q) \quad &\mapsto \quad \boldsymbol{z}_2 = M_\wedge \ltimes \boldsymbol{p} \ltimes M_\neg \ltimes \boldsymbol{q},\\
z_3 = p\wedge q \quad &\mapsto \quad \boldsymbol{z}_3 = M_\wedge \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}.
\end{aligned}
$$

All but $\boldsymbol{z}_2$ is in its canonical form as stated in Theorem 1. If $\boldsymbol{p}$ and $M_\neg$ were swapped, then $\boldsymbol{z}_2$ would conform with the expected order. We can achieve such swap using Example 5,

$$
\boldsymbol{z}_2 = M_\wedge \ltimes \boldsymbol{p} \ltimes M_\neg \ltimes \boldsymbol{q} = M_\wedge \ltimes (I_2\otimes M_\neg) \ltimes \boldsymbol{p} \ltimes \boldsymbol{q}.
$$

Now all expressions have the same structure. We can determine the representation of the original expression of $f(p,q)$ (see Figure 2) and check whether M_f does indeed reflect the simplification in f. Interestingly, if one calculates the semi-tensor product for M_f, it simplifies with no intervention to the expected result: $M_f = \delta_2[1,1,1,2] = M_\vee$. ◁

Example 14 (Maximum length feedback function)**.** Let us consider a feedback functions of $n = 4$ variables:

$$
f(x_1,x_2,x_3,x_4) = 1\oplus x_1\oplus x_2\oplus x_1x_2x_3x_4. \tag{6}
$$

This function generates a maximal length, i.e., $2^4-1 = 15$, sequence. Table II shows the sequence of states started from initial state $s_1 = (0,0,0,1)$. The rows contain states of each register x_1, x_2, x_3, and x_4.

TABLE II
SEQUENCE OF STATES FOR FEEDBACK FUNCTION f.

	states															
x_i	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}	s_{11}	s_{12}	s_{13}	s_{14}	s_{15}	s_{16}
1	0	0	0	1	1	1	0	1	1	0	0	1	0	1	0	0
2	0	0	1	1	1	0	1	1	0	0	1	0	1	0	0	0
3	0	1	1	1	0	1	1	0	0	1	0	1	0	0	0	0
4	1	1	1	0	1	1	0	0	1	0	1	0	0	0	0	1

We are now in a position to examine this FSR using the semi-tensor product formalism. The feedback function is

$$
\begin{aligned}
f(x_1,x_2,x_3,x_4) \mapsto M_\oplus \ltimes \mathbf{1} \ltimes (M_\oplus \ltimes \boldsymbol{x_1} \ltimes\\
(M_\oplus \ltimes \boldsymbol{x_2} \ltimes (M_\wedge \ltimes \boldsymbol{x_1} \ltimes M_\wedge \ltimes \boldsymbol{x_2} \ltimes M_\wedge \ltimes \boldsymbol{x_3} \ltimes \boldsymbol{x_4}))),
\end{aligned}
$$

where $\mathbf{1} = \delta_2^1$ and thus the corresponding structure matrix is

$$
\begin{aligned}
M_f &= M_\oplus \ltimes \mathbf{1} \ltimes M_\oplus \ltimes (I_2\otimes M_\oplus) \ltimes (I_{2^2}\otimes M_\wedge) \ltimes\\
&= (I_{2^3}\otimes M_\wedge) \ltimes (I_{2^4}\otimes M_\wedge) \ltimes {}_2M_\downarrow\\
&= \delta_2[2,1,1,1,2,2,2,2,2,2,2,2,1,1,1,1]
\end{aligned}
$$

This matrix has the expected shape: 2×16, where 2 is the size of the output of a feedback function, while 16 is the number of possible inputs of f from four Boolean variables. ◁

Theorem 1 guarantees that any logical function, even non-linear ones, can be represented in the semi-tensor product formalism, therefore, it seems plausible that the theoretical description of FSRs could benefit from this formalism, especially because linear and non-linear FSRs are treated on equal footing. However, so far we have only focused on the logical feedback function f which determines the next state of the last register only. We wish to associate a single matrix to the entire FSR which describes the time evolution of *all* registers.

The time-evolution of a generic feedback shift-register is given by the following set of equations

$$
\begin{aligned}
x_1(t+1) &= x_2(t)\\
x_2(t+1) &= x_3(t)\\
&\vdots\\
x_{n-1}(t+1) &= x_n(t)\\
x_n(t+1) &= f(x_1(t),x_2(t),\ldots,x_n(t)).
\end{aligned} \tag{7}
$$

In order to render equation (7) into a single equation, we organise the primitive variables into a vector-variable, $z =$

$(x_1, x_2, \ldots, x_n)$. The semi-tensor formalism maps all such n-long Boolean vectors one-to-one to a single 2^n-long vector, $z \mapsto \boldsymbol{z} = \ltimes_{j=1}^{n} \boldsymbol{x}_j$. Equation (7) can thus be written as

$$\boldsymbol{z}(t+1) = T_f \boldsymbol{z}(t),$$

where T_f is the transition function induced by f. Since T_f is a logical function Theorem 1 applies. We can even give an explicit expression for T_f

$$\begin{aligned}
\boldsymbol{z}(t+1) &= \ltimes_{j=1}^{n} \boldsymbol{x}_j(t+1) = \left[\ltimes_{j=1}^{n-1} \boldsymbol{x}_j(t+1)\right] \ltimes \boldsymbol{x}_n(t+1) \\
&= \left[\ltimes_{j=2}^{n} \boldsymbol{x}_j(t)\right] \ltimes f(\boldsymbol{x}_1(t), \boldsymbol{x}_2(t), \ldots, \boldsymbol{x}_n(t)) \\
&= \left[\ltimes_{j=2}^{n} \boldsymbol{x}_j(t)\right] \ltimes M_f \ltimes \left[\ltimes_{j=1}^{n} \boldsymbol{x}_j\right] \\
&= (I_{2^{n-1}} \otimes M_f) \ltimes \left[\ltimes_{j=2}^{n} \boldsymbol{x}_j(t)\right] \ltimes \boldsymbol{x}_1 \ltimes \left[\ltimes_{j=2}^{n} \boldsymbol{x}_j\right] \\
&= (I_{2^{n-1}} \otimes M_f) \ltimes W_{[2,2^{n-1}]} \ltimes \boldsymbol{x}_1 \ltimes \left[\ltimes_{j=2}^{n} \boldsymbol{x}_j(t)\right]^2 \\
&= (I_{2^{n-1}} \otimes M_f) \ltimes W_{[2,2^{n-1}]} \ltimes (I_2 \otimes {}_{n-1}M_{\downarrow}) \ltimes \boldsymbol{z}(t) \\
&= T_f \boldsymbol{z}(t).
\end{aligned}$$

One may read off the expression for T_f:

$$T_f = (I_{2^{n-1}} \otimes M_f) \ltimes W_{[2,2^{n-1}]} \ltimes (I_2 \otimes {}_{n-1}M_{\downarrow})$$

Example 15 (Explicit transition function)**.** In Example 14 we have determined M_f, of a particular 4-stage FSR. Let us construct the corresponding T_f function:

$$T_f = (I_{2^3} \otimes M_f) \ltimes W_{[2,2^3]} \ltimes (I_2 \otimes {}_3M_{\downarrow})$$

All matrices in T_f are known and substituting them results in

$$T_f = \delta_{16}[2, 3, 5, 7, 10, 12, 14, 16, 2, 4, 6, 8, 9, 11, 13, 15] \quad (8)$$

Following on Example 14 let us only determine the state which follows the initial $z(0) = (0, 0, 0, 1) \mapsto \boldsymbol{z}(0) = \delta_{16}^{15}$ state. Calculating the product $T_f \ltimes \boldsymbol{z}(0) = T_f \delta_{16}^{15}$ is easy as it coincides with he standard matrix product: $\boldsymbol{z}(1) = \delta_{13}^{16} \mapsto z(1) = (0, 0, 1, 1)$. Although we haven't included the entire calculation here, it agrees with Table II. ◁

Finally, we state a pleasant property of T_f. As this operator evolves a state in time, the sequence

$$\boldsymbol{z}_0 \to \boldsymbol{z}_1 = T_f \ltimes \boldsymbol{z}_0 \to \cdots \to \boldsymbol{z}_K = \left(\ltimes_{j=1}^{K} T_f\right) \ltimes \boldsymbol{z}_0$$

provides the outputs. Presume that $\boldsymbol{z}_1 = T_f \boldsymbol{z}_0 = \boldsymbol{z}_0$ holds for some initial state $\boldsymbol{z}_0$. How many such $\boldsymbol{z}_0$ exists? It seems plausible that the number of fixed points depends solely on the structure of T_f. The following theorem clarifies this.

Theorem 2 (Number of cycles of length 1)**.** *Let $D = \{0, 1\}$ be the set of possible states of a register, and T_f be the transition function of n such registers. If C_1 denotes the number of cycles of length 1, then*

$$C_1 = \mathrm{Tr}\,(T_f).$$

Example 16 (Cycle of lengths 1 for f)**.** The trace of structure matrix given in (8) is zero. Thus no state is left unchanged under the action of feedback function (6). This conclusion is supported by intuition, since each state is shifted at a time-step, therefore the only possible states which could be fixed points are $(0, 0, 0, 0)$ and $(1, 1, 1, 1)$. These states, however, cannot be fixed states because of the term $1 \oplus x_0$ in f. ◁

IV. Conclusion

Modelling a complex system by segmenting it to simpler blocks is the standard approach in most natural and engineering sciences. Limiting ourselves to very few identical blocks and only two states for each block (Boolean network) leads to astonishingly complex behaviour for the combined system. This model, introduced in the mid-1960s [12], has proved to be extremely useful in many fields, from physics to biology.

In electronics an important subclass of Boolean networks is the feedback shift-registers. While the theory of linear FSRs is matured, the tools developed for linear FSRs were inadequate for nonlinear FSRs. In the last decade a novel algebraic structure has emerged which is capable of capturing not only linear and non-linear FSRs, but more general Boolean networks as well. In this study we have introduced this new structure and demonstrated its properties on examples. We have also built up the description of a particular nonlinear feedback shift-register which produces maximal length output sequence.

Acknowledgement

The authors acknowledge the financial support from MBIE contract UOOX-1208. We thank the reviewer for his/her suggestions.

Appendix
HDL code with test-bench

An implementation of a FSR equipped with f of Example 14 is given together with a test-bench.

Listing 1. A FSR with feedback function from Example 14 in Verilog HDL.

```
module nfsr(clock, reset, init, state);
    input clock, reset;
    input [1:4] init;
    output [1:4] state;

    reg [1:4] x = 4'b0;
    wire f, tmp1, tmp2, tmp3;

    always@(posedge clock or posedge reset)
    begin
        if (reset == 1)
            x = init;
        else
            x = {x[2:4], f};
    end

    // The feedback function.
    and g1(tmp1, x[1], x[2], x[3], x[4]);
    xor g2(tmp2, x[2], tmp1);
    xor g3(tmp3, x[1], tmp2);
    xor g4(f, 1'b1, tmp3);

    assign state = x;
endmodule
```

Listing 2. Test-bench for the `nfsr` module.

```
module testbench();
    reg [1:4] init = 4'b0001;
    wire [1:4] state;
    reg rst = 1'b0;

    // Beating clock.
    reg clk = 1'b0;
    always #2 clk = !clk;

    // Unit under test.
    nfsr uut(clk, rst, init, state);

    // Stimuli.
    initial begin
        #100 rst = 1'b1;
        #003 rst = 1'b0;
        $monitor("time=%t, clk=%b, rst=%b, state=%4b",
                 $time, clk, rst, state);

        #070 $finish;
    end
endmodule
```

APPENDIX
SCEMATICS OF A NONLINEAR FSR

Figure 3 depicts the schematics of a possible realisation of a feedback shit-register circuit equipped with feedback as in Example 14. The shift-register segment contains four D flip-flops whose sole function is propagating Q onto their outputs as the clock ticks. The feedback is $1 \oplus x_1 \oplus x_2 \oplus x_1x_2x_3x_4$ which contains an "$1\oplus$" term, i.e. a XOR gate with an input of logical 1. Examining the truth-table of XOR shows that such setup can be replaced by a NOT gate (cf., Figure 3). Other possible inputs of the D flip-flops are omitted for clarity.

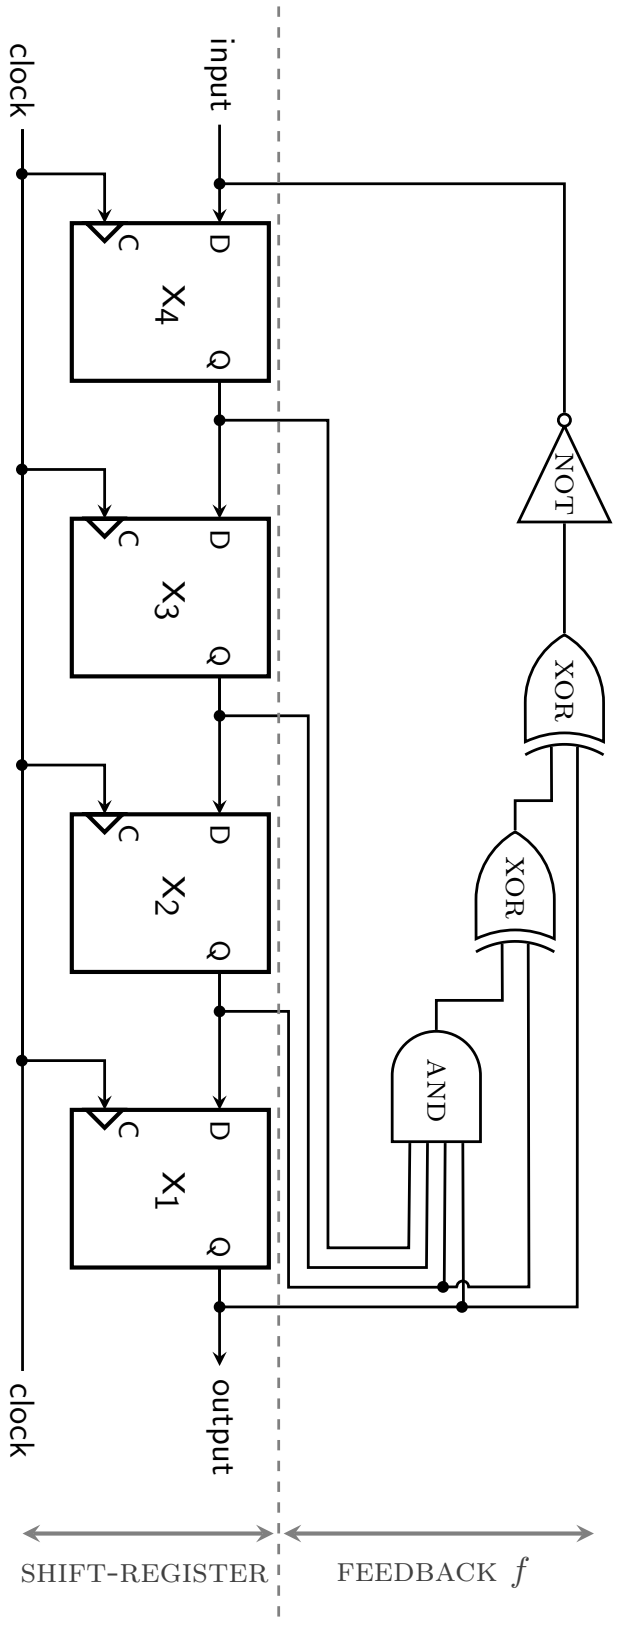

Fig. 3. Schematics of a feedback shift-register circuit equipped with the logic feedback function analysed in Example 14.

REFERENCES

[1] N. Zierler, "Linear recurring sequences," *SIAM J. Appl. Math.*, vol. 7, no. 1, pp. 31–48, 1959.

[2] S. W. Golomb, *Shift Register Sequences*, revised ed. San Francisco, Holden-Day, 1967, Laguna Hills, CA: Aegean Park Press, 1982.

[3] R. W. Ward and T. C. A. Molteno, "Table of Multiple Feedback Shift Registers," Electronics Group at the University of Otago, Tech. Rep. 2009-1, 2009.

[4] R. W. Ward and T. Molteno, "Table of Linear Feedback Shift Registers," University of Otago, Tech. Rep. 2012-1, 2012.

[5] S. W. Golomb and G. Gong, *Signal Design for Good Correlation: For Wireless Communication, Cryptography, and Radar*, 1st ed. Cambridge University Press, 2005.

[6] E. Dubrova, M. Teslenko, and H. Tenhunen, "On Analysis and Synthesis of (n, k)-Non-Linear Feedback Shift Registers," in *2008 Design, Automation and Test in Europe*, March 2008, pp. 1286–1291.

[7] J. Mykkeltveit, M.-K. Siu, and P. Tong, "On the cycle structure of some nonlinear shift register sequences," *Information and Control*, vol. 43, no. 2, pp. 202–215, 1979.

[8] P. Daþrowski, G. Łabuzek, T. Rachwalik, and J. Szmidt, "Searching for nonlinear feedback shift registers with parallel computing," *Information Processing Letters*, vol. 114, no. 5, pp. 268–272, 2014.

[9] C. R. Edwards, "The Logic of Boolean Matrices," *The Computer Journal*, vol. 15, no. 3, pp. 247–253, 1972.

[10] D. Cheng and Y. Dong, "Semi-tensor product of matrices and its some applications to physics," *Methods Appl. Anal.*, vol. 10, no. 4, pp. 565–588, 12 2003.

[11] D. Cheng, H. Qi, and Z. Li, *Analysis and Control of Boolean Networks: A Semi-tensor Product Approach*, ser. Communications and Control Engineering. Springer-Verlag London, 2011.

[12] C. C. Walker and W. R. Ashby, "On temporal characteristics of behavior in certain complex systems," *Kybernetik*, vol. 3, no. 2, pp. 100–108, 1966.

GPS Estimates in the Presence of Correlated Noise

Andrew Soundy, Daniel Schumayer, Timothy Molteno
Department of Physics
University of Otago
tim@elec.ac.nz

***Abstract*—This paper presents an experimental and theoretical investigation of GPS errors. Data from stationary GPS units were gathered and the time dependence of the observed noise is presented. Two quantitative approaches for dealing with time dependent noise are provided.**

I. Introduction

The ubiquity of global navigation satellite systems (GNSSs) has likely not escaped the reader's attention. GPS still has the largest constellation of any GNSS system and all GNSS measurements made in this paper use the GPS constellation. The first GPS satellite was launched in 1978 and the constellation was declared fully operational in 1995 [14]. As of March 9, 2016 the GPS constellation has more satellites in orbit than any other system and the Federal Aviation Administration (FAA) has shown single frequency GPS to have a horizontal accuracy of less than 3.5 m across North America [3], [1].

There are many sources that contribute to GPS errors. The dominant source of error in GPS measurements results from the dispersion of the microwave signal in the ionised upper atmosphere [13]. Dual frequency GPS units are able to largely compensate for the errors from ionospheric effects by calculating the atmospheric dispersion from the different time delays on the L1 (1575.42 MHz) and L2 (1227.60 MHz) signals [8]. However, single frequency GPS units cannot make this correction and dual frequency GPS receivers are considerably more expensive ($\sim$\$2000+) than single frequency receivers ($\sim$\$10+).

Accurate weather data can be used to estimate the amount of moisture in the air, and this information can largely correct for tropospheric errors on GPS measurements [8]. If high-precision measurements are required then the user may be able to use a dual frequency GPS module and real-time weather models to substantially reduce the errors in GPS measurements. However, given the costs involved these methods are outside the means of everyday GPS users.

For the purposes of this paper we take an agnostic position to the source of the noise observed in the GPS signal. We treat the GPS unit as a 'black box', and simply investigate the returned data.

II. Correlated Noise in Single Frequency GPS Measurements

A common method for determining the relative time dependence of a set of measurements is to examine the autocorrelation function of the set of measurements [4], [9]. The autocorrelation function ($\rho(s)$) of a signal indicates how strongly correlated a signal is with itself after a time delay of s seconds. If the autocorrelation function is well defined then it should be bounded by [-1,1], with 1 representing prefect correlation, 0 representing perfect independence, and -1 representing perfect anticorrelation. The autocorrelation function of a set of discrete data $\{X_i\}$ can be estimated by [5],

$$\rho(s) = \frac{\sum_{t=s+1}^{n}(X_t - \mu)(X_{t-s} - \mu)}{\sum_{t=1}^{n}(X_t - \mu)^2}, \tag{1}$$

where μ is the mean of the set of data. Figure 1 shows the autocorrelation function of latitude measurements taken from the Physics Department at the University of Otago. The data was taken with a GlobalSat BU-353S4 GPS unit operating on the L1 frequency (1575.42 MHz) with sampling period of 1 second over an aggregate time of approximately 100 hours, in total more than 355,000 individual measurements.

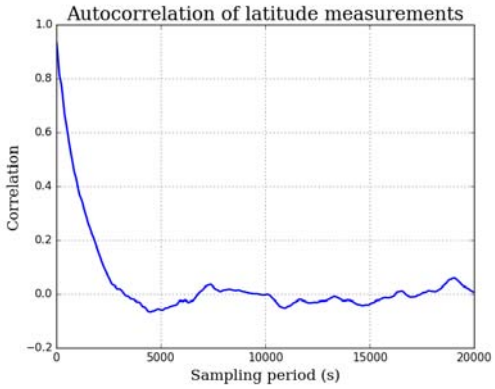

Fig. 1: Autocorrelation of latitude errors taken from a GlobalSat BU-353S4 GPS unit with excellent sky view. The sampling period of the device is 1 second. This data is taken from a data set of $\sim$ 100 hours but the time axis has been cut off at 5.5 hours for readability.

This observation of long-term autocorrelation within GPS measurements agrees with some previous work [12], [11]. However, given the problem of correlated noise the question of how to handle correlated noise remains.

In a previous work we took GPS data from Dunedin, and several locations at different latitudes in Australia and used the Akaike Information Criterion (AIC) to determine the best noise model from the candidate set [15]. The noise models investigated were an independent, identically distributed

Gaussian (i.i.d Gaussian) model, an autoregressive model of order p (AR(p)), a moving average model of order q (MA(q)), a mixed autoregressive moving average model of order p,q (ARMA(p,q)), and the Ornstein-Uhlenbeck model (O-U). According to the AIC the higher order autoregressive, and mixed autoregressive moving average models tended to be the best fit to the data. However, the Ornstein-Uhlenbeck model also performed significantly better than the i.i.d Gaussian model, which was generally the worst fit to the data.

III. The Ornstein-Uhlenbeck Process

An Ornstein-Uhlenbeck (O-U) process is defined by the stochastic differential equation,

$$dx_t = \theta(\mu - x_t)dt + \sigma dW_t, \tag{2}$$

where μ is the long-term average of the estimated variable $\{x_t\}$, θ is a constant parameter of the model, and W_t is the Wiener process. This then implies that σ from Equation (2) is the standard deviation of the normally distributed noise present in the O-U process.

The solution of the Ornstein-Uhlenbeck process is as follows [16].

$$x_t = x_0 e^{-\theta t} + \mu\left(1 - e^{-\theta t}\right) + \sigma \int_0^t e^{\theta(s-t)} dW_s \tag{3}$$

We see that the O-U process is normally distributed with expected value and variance as given by,

$$\mathrm{E}[x_t] = x_0 e^{-\theta t} + \mu\left(1 - e^{-\theta t}\right), \quad \mathrm{var}[x_t] = \frac{\sigma^2}{2\theta}\left(1 - e^{-2\theta t}\right) \tag{4}$$

i.e.

$$x_t \sim \mathcal{N}\left(x_0 e^{-\theta t} + \mu(1 - e^{-\theta t}), \frac{\sigma^2}{2\theta}(1 - e^{-2\theta t})\right). \tag{5}$$

We can then easily move from the continuous case to the discrete case. If we let dt be the sampling period of the discrete case, then we get the following,

$$\mathrm{E}[x_i] = x_{i-1} e^{-\theta dt} + \mu\left(1 - e^{-\theta dt}\right), \quad \mathrm{var}[x_i] = \frac{\sigma^2}{2\theta}\left(1 - e^{-2\theta dt}\right). \tag{6}$$

In the context of using an Ornstein-Uhlenbeck process to model GPS noise, the long-term mean (μ) is the most important property. If we assume that the device is stationary, i.e. μ does not change in time, then we can re-write our measurement in terms of 'true' position, μ, and a noise term n_t which has O-U noise, i.e.

$$x_i = \mu + n_i, \tag{7}$$

where n_i is given below.

$$n_i \sim \mathcal{N}\left(e^{-\theta dt}(x_{i-1} - \mu), \frac{\sigma^2}{2\theta}\left(1 - e^{-2\theta dt}\right)\right) \tag{8}$$

If we substitute $(x_{i-1} - \mu)$ for n_{i-1}, then we get the following pleasant expression,

$$n_i \sim \mathcal{N}\left(n_{i-1} e^{-\theta dt}, \frac{\sigma^2}{2\theta}\left(1 - e^{-2\theta dt}\right)\right), \tag{9}$$

for the distribution of the i^{th} noise variable term. As you can see in Equation (9) the O-U process is clearly time dependent. Indeed, once it is discretised the O-U process simplifies to an autoregressive process of order 1. However, we believe that preferentially using the O-U process over an arbitrary AR(1) process has several benefits. For one, the sampling period of the device is explicit in the formulation of the Ornstein-Uhlenbeck process and so can easily accommodate non-constant sampling periods. Also the other parameters of the O-U process give clear insight into the characteristics of the model. The θ parameter describes the rate at which random measurements will revert towards the mean μ while the σ parameter indicates the amount of noise present between each measurement [17].

IV. Filtering With the Ornstein-Uhlenbeck Process

The Kalman filter is a versatile algorithm that combines sequential measurements with a linear model and provides asymptotically optimal estimates of the state while minimizing the state covariance [10]. One of the convenient aspects of the Kalman filter is that it is a recursive algorithm and so all of the information from the previous measurements is contained in the current state estimate, therefore significantly easing the computational load over each iteration.

Another convenient aspect of the Kalman filter is that if you are using an independent noise model that can be characterised completely by its covariance, e.g. an i.i.d Gaussian model, then you can define your system noise, and measurement noise covariances before the filter begins to iterate thereby further easing the computational load [10]. This, combined with the central limit theorem makes the i.i.d Gaussian noise model a very favourable noise model to use with the Kalman filter. However, given the long-term time dependence of GPS noise, clearly an i.i.d Gaussian model is less than ideal. This brings us to the Ornstein-Uhlenbeck noise model.

Filtering with the O-U model will require expanding the state vector of the Kalman filter. But we need only include a single extra noise variable for each O-U filtered variable, e.g. three extra for latitude, longitude, and altitude. We also have a measurement model and propagation model ready from Equation (9). The state vector ($\hat{\mathbf{x}}(k)$) and state propagation matrix ($F(k)$) to filter an O-U variable are given below, along with the measurement matrix ($H(k)$):

$$\hat{\mathbf{x}}(k) = \begin{bmatrix} \hat{\mu} \\ n(k) \end{bmatrix}, \quad F(k) = \begin{bmatrix} 1 & 0 \\ 0 & \exp(-\theta dt) \end{bmatrix}, \quad H(k) = [1 \quad 1]. \tag{10}$$

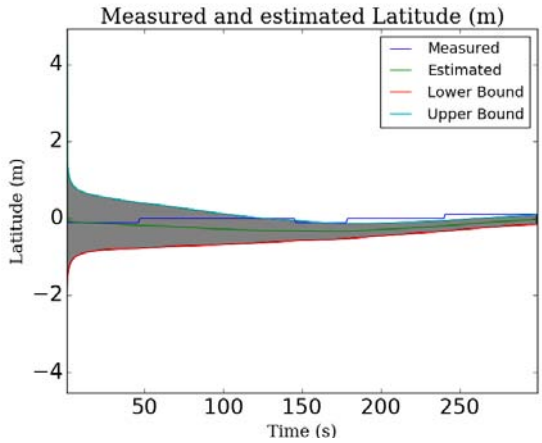

Fig. 2: Above shows the results of filtering latitude measurements with an Ornstein-Uhlenbeck noise model. The estimate is plotted in green and shown within a 95% confidence interval.

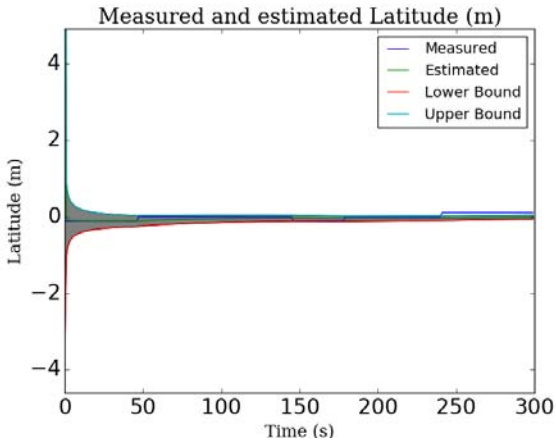

Fig. 3: Above shows the results of filtering latitude measurements with an i.i.d. Gaussian noise model. The system noise is the same as in Figure 2, the only difference is that above we are using an i.i.d. Gaussian noise model rather than an O-U noise model. Note the different behaviour of the uncertainty estimates.

To implement a standard Kalman filter we now only need the covariance matrices for the measurement ($R(k)$) and process noises($Q(k)$). The only slight complication is that for a system where the measurement device is stationary there is virtually no process noise present in the system. However to construct a functioning Kalman filter there exists some computational reasons against having the matrix $Q(k)$ set to exactly zero. Fortunately setting $Q(k)$ to some small value, e.g. 10^{-10}, deals with these satisfactorily. This then gives us the state vector and covariance matrices as follows.

$$R(k) = \left[\frac{\sigma^2}{2\theta}\left(1 - e^{-2\theta dt}\right)\right], \qquad Q(k) \approx [0]\,. \tag{11}$$

(12)

Observe the explicit dependence on the sampling period in the descriptions of the state propagation, and measurement covariance matrices, $F(k)$ and $R(k)$, respectively. This means that the sampling period need not be constant. In our experience some GPS devices have difficulty maintaining a constant sampling period when taking data from a warm start. This would present difficulties for some noise models as the parameters involved can have an implicit time dependence, as is the case for the autoregressive and moving average models [5].

With the sampling period included in the matrices the Ornstein-Uhlenbeck model takes into account the effect of longer or shorter time periods and allows the filter to adjust automatically.

We demonstrate the filtering proposed above in Figure 2. We use real latitude measurements and have plotted a 95% confidence interval around the state estimate. Compare that to Figure 3 where we have used an i.i.d. Gaussian noise model. The Ornstein-Uhlenbeck noise model has a considerably larger uncertainty after the same amount of time which reflects the uncertainty involved when using time correlated measurements.

V. Using the sampling period instead of filtering

We have shown the long-term autocorrelation function of noise observed in single frequency GPS devices. We have also shown how filtering is possible with a time dependent noise model. However, another way to deal with the correlated noise might be to borrow methods that were originally developed to handle the correlated outputs from Markov Chain Monte Carlo (MCMC) systems.

In a MCMC system various outputs $\left\{x^{(n)}\right\}$ are produced. For our situation, let us treat each $x^{(i)}$ as a position measurement. For a MCMC, and for a GPS measurement the current measurement ($x^{(k)}$) is correlated to the previous measurement ($x^{(k-1)}$). If our set of measurements $\left\{x^{(n)}\right\}$ were independent then the variance of our estimate ($\mathrm{var}(x_N)$) would take the form,

$$\mathrm{var}(x_N) = \frac{\bar{\sigma}^2}{N}. \tag{13}$$

Where $\bar{\sigma}^2$ is the sample variance and N is the number of measurements. However, in the case of correlated measurements then the variance of our estimate is given by,

$$\mathrm{var}(x_N) = \frac{\bar{\sigma}^2 \tau_{ff}}{N}, \tag{14}$$

where τ_{ff} is the integrated autocorrelation time. The integrated autocorrelation time for a MCMC is given by [7],

$$\tau_{ff} = 1 + 2\sum_{s=1}^{\infty} \rho(s), \tag{15}$$

where $\rho(s)$ is the autocorrelation function after lag s. In Figure 4 we have fitted a decaying exponential to the empirical

autocorrelation function of GPS data gathered at the University of Otago. The dashed lines in Figure 4 show a 95% confidence bound, as defined by the Box-Jenkins standard error. The Box-Jenkins standard error is given by [6],

$$-\frac{z(1-\frac{\alpha}{2})}{\sqrt{N}}S(h,\rho) \leq \mathrm{E}[X_t] \leq \frac{z(1-\frac{\alpha}{2})}{\sqrt{N}}S(h,\rho), \quad (16)$$

$$S(h,\rho) = \sqrt{1+2\sum_{k=1}^{h-1}\rho^2(k)}, \quad (17)$$

where $z(x)$ denotes the cumulative distribution function of the standard normal probability distribution, α is the desired level of confidence, and h is the highest used index of your autocorrelation function.

In Figure 4 we fit an exponential to the empirical autocorrelation function of a set of latitude measurements. A 95% confidence interval, as defined by the Box-Jenkins standard error, is plotted in dashed lines around zero, i.e. within those bounds the empirical autocorrelation function is predominantly noise.

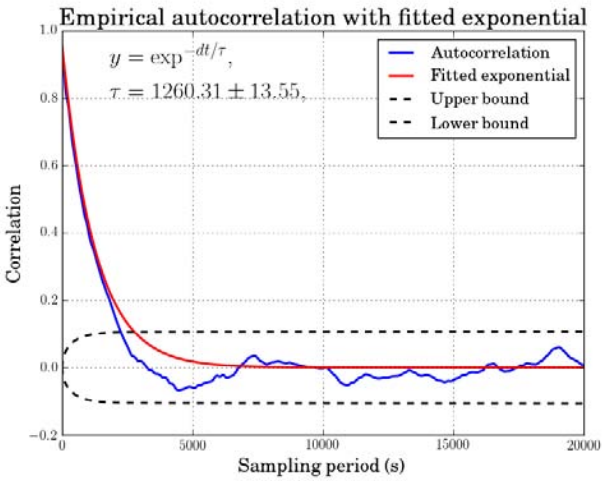

Fig. 4: Autocorrelation of latitude errors with fitted exponential. The dashed lines show the region of 95% confidence of independence, as defined by the Box-Jenkins standard error. The exponential is fitted from the 1800 sampling periods.

The oscillations of the autocorrelation function around zero are within the Box-Jenkins standard error. Therefore we will truncate the sum in Equation (15) so that arbitrary noise is not introduced into our estimate of the integrated autocorrelation time. We propose three time periods of the fitted decay constant (τ) as a 'natural' cut-off point. Therefore after three decay constants we assume that measurements are approximately independent from each other.

A well-known technique for reducing the autocorrelation between measurements in a MCMC is to subsample the Markov chain [7]. That is, instead of using every measurement, you use every n^{th} measurement.

In Figure 5 we used a data set of approximately 550 hours. For every sampling period we generated 20 subsets of data, each with that sampling period and with a random starting point somewhere within the original data set. The mean and variance of these subsets are shown in the figure. Observe how the subsets with longer sampling periods consistently have means that are closer to the true mean. The sample variances of subsets with longer sampling periods also return more accurate estimates of the original variance. If we look at some of the faster sampling periods, even with 300 measurements, some of the subsets have mean altitude estimates that are more than six metres from the true mean.

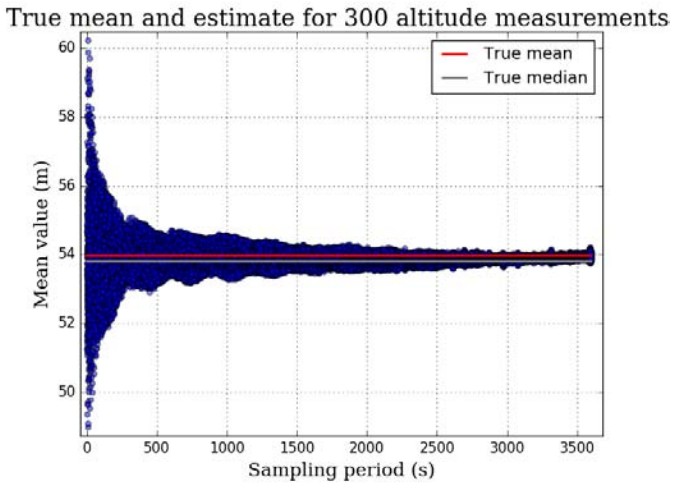

(a)

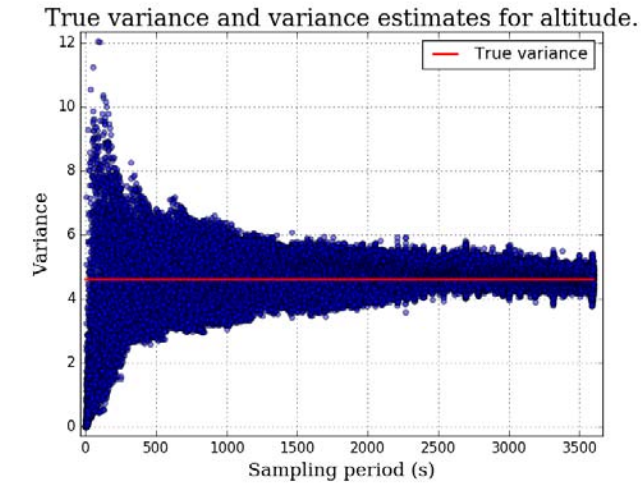

(b)

Fig. 5: For each sampling period we started 20 subsets of data with the same sampling period, but at random times within the full data set. In Figure (a) we see the mean of these subsets, and in Figure (b) we see the sample variances of the subsets. Note that for fast sampling periods $\sim< 50$ the sample variances are artificially low, this is because of the strong correlation between measurements.

However, to be clear you do not gain any information by throwing away measurements. That is, if the cost of sampling is negligible, then any sub-sampling is not only not useful but

will result in worse data analysis [7].

In the case of GPS measurements the cost of sampling is power consumption. Some modern GPS chipsets are able to maintain hot start conditions while using a very small amount of current. As an example, the SiRF STAR IV GSD4e GPS chipset [2] that was used to make the measurements presented here, has a continuous tracking current of 54mW. But it is able to maintain hot start conditions with a current of 50-500μA. From a hot start the Time to First Fix (TTFF) is less than one second, versus up to 35 seconds from a warm or cold start [**?**].

Therefore, if we sought to use a single frequency GPS unit to determine the location of a stationary object, say a marker post for a future building, we could use the cost function of power consumption of the SiRFstarIV chipset to determine the optimal sampling period.

Let's imagine that our GPS device has enough battery life for 300 measurements when continuously sampling at 1Hz. Given a continuous tracking current (I_T) of 54mW and a hot start current (I_H) of 500μA we can calculate the number of measurements possible (N) using the expression.

$$N = \left(\frac{\text{totalcurrent} - I_T}{(dt-1)I_H + I_T} \right) + 1. \tag{18}$$

In Figure 6 we see can see the effect on the number of samples given some sampling period dt.

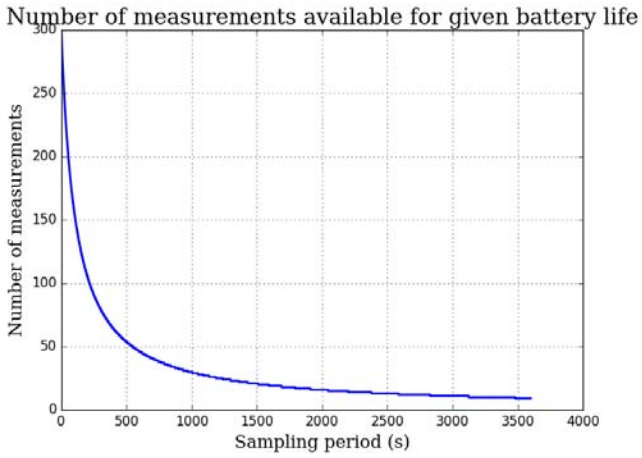

Fig. 6: Given a tracking current of 54mW, a current to maintain hot start conditions of 500μA and enough battery life for 300 measurements at 1Hz we get the above relationship between the number of measurements possible for a given sampling period.

The integrated autocorrelation times from empirical data and the fitted exponential from Figure 4 are both shown in Figure 7. The sum for the integrated autocorrelation time has been truncated at 3 time periods, or around 3700s.

Recalling Equation (14), if we assume that the sample variances for the subsets are reasonably similar then by dividing the integrated autocorrelation time (τ_{ff}) for a given sampling period by the number of measurements possible for

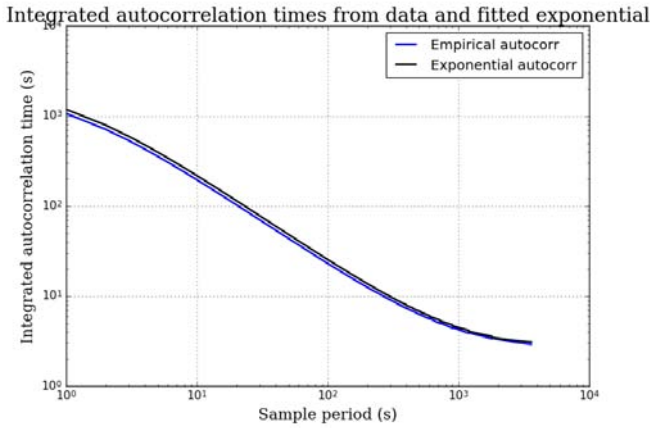

Fig. 7: The integrated autocorrelation times for the empirical data (blue) and the fitted exponential (black) with parameters from Figure 4 plotted on a logarithmic scale for clarity.

that sampling period we can then look for a minimum as an indication of a useful sampling period. This τ_{ff}/N is shown in Figure 8 and achieves a minimum value with a sampling period of 360s.

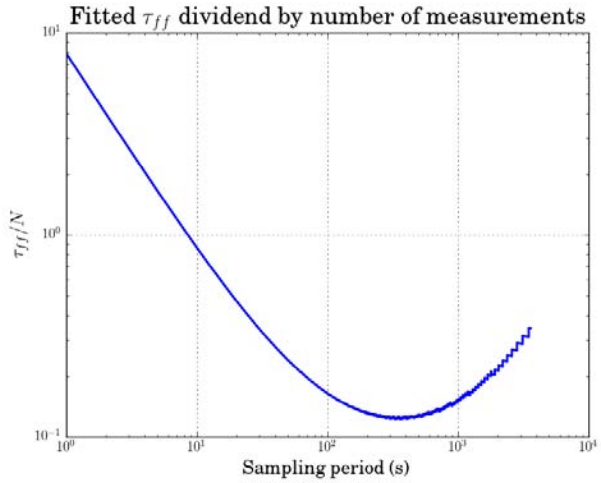

Fig. 8: The integrated autocorrelation time divided by the number of measurements possible for that sampling period. For readability this has been plotted on a logarithmic scale. You can see the minimum τ_{ff}/N is achieved with a sampling period of around 360s, which suggests a natural sampling period to investigate.

To investigate whether this predicted optimal sampling period does perform any better or worse we looked at another data set and, similarly to Figure 5, chose random starting places for subsets of data with a given sampling period. Each sampling period had a thousand subsets of data, and the results are plotted in Figure 9.

Interestingly we see that after our predicted sampling period there is little advantage or disadvantage to having a

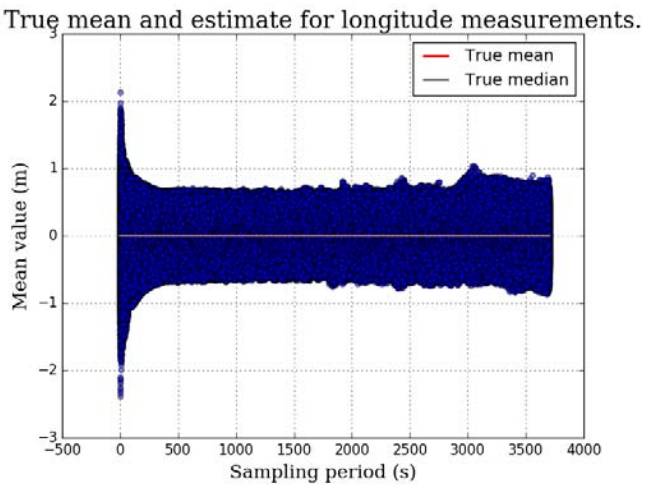

Fig. 9: For each integer sampling period from 1s to 1hr we created a thousand sample sets with random start points within the data set. Above we see the means of all of the sample sets, along with the true mean and median of the entire data set.

longer sampling period. The benefits due to the increased independence of measurements is almost exactly offset by the decrease in the number of possible measurements. However, what is clear is that sampling faster than our predicted sampling period does lead to worse estimates. This is despite the fact that at a sampling period of 1Hz you would be able to make 300 measurements, versus only 70 measurements at 360Hz.

VI. Conclusions

We have presented corroborating evidence of the long-term auto-correlation functions of single frequency GPS measurements. We then proposed two separate mechanisms to accurately estimate stationary positions given this long-term correlation. The first involves using a time-dependent noise model, the Ornstein-Uhlenbeck process, and incorporating that into a Kalman filter. This enables us to maintain fast sampling periods with the extra cost of a larger state estimate for the Kalman filter, and therefore a higher computational cost. We recommend this approach if the computational cost is acceptable. However, in cases where battery life is a significant issue then a low-power alternative is to sample more slowly. This reduces the correlation between measurements and if the rate of power usage is known then an optimal sampling time may be proposed wherein the benefits of greater independence outweigh the reduced number of possible measurements.

References

[1] Anon. Official U.S. government information about the Global Positioning System (GPS) and related topics.

[2] Anon. *Datasheet for USB GPS Receiver : BU-353S4(Datasheet*. GlobalSat Technology Corporation, September 2013.

[3] William J. Hughes Technical Center. Global Positioning System (GPS) Standard Positioning Service (SPS) Performance Analysis Report. Technical Report 86, Federal Aviation Administration, July 2014.

[4] Chris Chatfield. *Time-series forecasting*. CRC Press, 2000.

[5] Jonathan D. Cryer and Kung-Sik Chan. *Time series analysis: with applications in R*. Springer, 2008.

[6] George S. Fishman. *Principles of discrete event simulation*. Wiley, 1978.

[7] Charles J. Geyer. Practical Markov Chain Monte Carlo. *Statistical Science*, 7(4):473–483, 1992.

[8] Paul D. Groves. *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Artech House, 2008.

[9] James Douglas Hamilton. *Time series analysis*, volume 2. Princeton university press, 1994.

[10] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME, Series D: Journal of Basic Engineering*, 82:35–45, 1960.

[11] Alfred Leick, Lev Rapoport, and Dmitry Tatanikov. *GPS Satellite Surveying*. John Wiley & Sons, 2015.

[12] Ailin Mao, Christopher G. A. Harrison, and Timothy H. Dixon. Noise in GPS coordinate time series. *Journal of Geophysical Research*, 104(B2):2797–2816, February 1999.

[13] Pratap Misra and Per Enge. *Global Posiitioning System Signal, Measurements, and Performance*. Ganga-Jamuna Press, 2 edition, 2006.

[14] Aboelmagd Noureldin, T. Ben Karmat, and Jacques Georgy. *Fundamentals of Inertial Navigation, Satellite-based Positioning and their Integration*. Springer, 2013.

[15] Andrew WR Soundy, Bradley J. Panckhurst, and Tim CA Molteno. Enhanced noise models for GPS positioning. *Automation, Robotics and Applications (ICARA), 2015 6th International Conference on*, 2015.

[16] George E. Uhlenbeck and Leonard S. Ornstein. On the theory of the Brownian motion. *Physical review*, 36(5):823, 1930.

[17] Oldrich Vasicek. An equilibrium characterization of the term structure. *Journal of financial economics*, 5(2):177–188, 1977.

Investigation of Light Transmission in Healthy and Rotten Onions

Jason Sun, Rainer Künnemeyer
School of Engineering
University of Waikato
Hamilton, New Zealand

Andrew McGlone, Nathan Tomer
Bioengineering
Plant and Food Research
Hamilton, New Zealand

***Abstract*— Light transmission through healthy and rotten onions was investigated in support our development of optical sensors for internal defects. Two types of systems were used: the first consisted of a broadband light source and a CCD spectrometer. Here the spectra had two peaks in the wavelength range from 670 to 850 nm. The second system employed a 728 nm laser diode and a rotating table, which enabled us to examine the light transmission spatially. Healthy, botrytis and pseudomonas onions were measured. The laser system could detect a signal attenuated by a factor of $3.25 \cdot 10^{-8}$ to $1.81 \cdot 10^{-07}$ due to the onions. The light level decreased exponentially with distance between source and detector. The effective attenuation was estimated for each onion and increased proportional to the degree of rottenness. The results of this study reveal the light transportation in onions and the effect from internal rots, which provide requirements for signal processing and optical systems.**

Keywords—onion; internal defects; sensor; optics.

I. Introduction

Onion is one of the largest fresh vegetable industries in the world. To meet the quality demand from consumers, non-destructive methods for quality assessments must be developed. NIR spectroscopy [1-3] is one optical technique routinely used for detecting internal defects if the defect is large and/or generalised in the tissue. However, it becomes less sensitive and more problematic as the defect becomes smaller and particularly if the defect is localised to certain regions. In onions, rotten tissue is often localised to the stem end or lies discretely along individual layers. The standard NIR technology cannot reliably determine the presence of such rots, which are a significant economic problem for the global onion industry.

In the VIS-NIR range of the electromagnetic spectrum, agricultural products can be considered as semi-transparent or turbid. There are two optical phenomena that describe how light interacts with biological tissues: absorption and scattering. Absorption is primarily due to the chemical composition of the tissue (pigments, chlorophylls, water etc.). Scattering depends on microscopic changes in refractive index caused by membranes, air vacuoles or organelles.

Understanding the light transport properties of onions is an important part of the design of optical methods for internal quality assessments [4]. In order to achieve reliable detection with a reasonable signal to noise ratio (SNR), three aspects must be considered in a design process: light source power, light transmission in onions, and detector sensitivity. The requirements for light source and detector depend on light transport in onions and the optical geometry. Therefore, we conducted this study to investigate the light transmission, which will aid our future design of a reliable sensor that can assess the internal quality of onions. We presented here two optical techniques for investigations: a traditional NIR spectroscopy technique and a laser based tomography technique [5].

II. Material and Methods

A. Onion Inoculation

In order to provide rotten onions for testing two common types of infections were grown. Botrytis fungus, which causes neck rot, and pseudomonas bacteria, which causes slippery skin rot. A wound was created by cutting the dry skin on the stem end off, a soldering iron was then used to produce a small wound on the exposed flesh. A solution containing botrytis or pseudomonas was injected into the wound on different onions. Three types of onions were then stored in an incubator at 23°: a healthy control, botrytis infected and pseudomonas infected. After two weeks the temperature of the incubator was reduced to 2° in order to slow down rot growth. All the measurements began after the onions were kept at room temperature for five hours.

In the cross section images (Fig. 1), we can see that pseudomonas caused rotten tissues on certain layers, while botrytis infected more layers and occupies a whole area. Both types of rots are more localised to the stem end where the inoculation took place.

Fig. 1. Cross sections of pseudomonas infected (left) and botrytis infected (right) onions.

B. Experimental Setup

Two types of light sources were used in the investigation: broadband QTH lightbulb and laser diode. The broadband light source with the CCD spectrometer allowed us to have a general picture of light transmission through the onion in the range of 650 to 950 nm. The laser system enabled us to study the light transmission spatially.

1) BroadBand Light Source Measurments

As the rots were more localised to stem end (Fig. 1), the light source and detector were arranged in a 'V' shape (Fig. 2) in which more light would go through the rotten tissues. It was expected that, that way, the spectra collected from CCD spectrometer were more affected by rotten tissue. The light source was a 250 W tungsten halogen lamp (6334NS, Newport, Irvine, CA, USA) driven by a digital radiometric power supply (69931, Newport, Irvine, CA, USA). The source light was coupled to the onion using a 14 mm diameter light pipe. The transmitted light was detected by a CCD spectrometer (CD024321, Control Development, South Bend, IN, USA) through a multimode fibre (M35L01, Thorlabs, Newton, NJ, USA). The integration time was 0.1 s. Both light pipe and fibre had an in-house designed adaptor for the attachment to the onion, which had a rubber seal to prevent unwanted stray light.

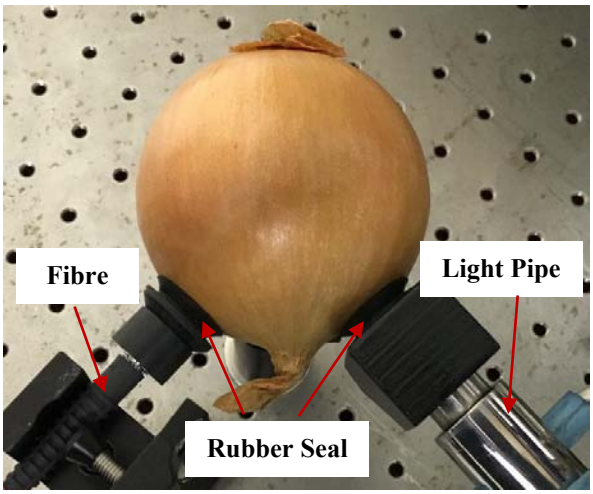

Fig. 2. Experimental setup of the broadband light source measurements.

2) Laser Light Source Measurements

The first transmission peak in the spectra collected in the broadband light source measurements varied around 710 nm, 728 nm was chosen due to the availability of the lasers. This wavelength was also away from the chlorophyll absorption, so that it will not be affected by the green tissues in onions. More importantly, this wavelength showed higher transmission through onions especially the healthy ones, from the broadband light source measurements. The laser diode (HL7302MG, Opnext, Fremont, CA, USA) was mounted in a laser diode mount (LDM21, Thorlabs, Newton, NJ, USA) driven by a laser diode driver and controller (ITC102, Thorlabs, Newton, NJ, USA). Two lenses were used to focus the light to a 5 mm radius spot. An adjustable pinhole was used to remove the elliptical shape of the beam and any stray light. The laser was operating to output 16 mw light power to the onion.

The detector consisted of a silicon photodiode (FDS100, Thorlabs, Newton, NJ, USA) and a trans-impedance amplifier (DLPCA-200, Femto, Berlin, Germany) which was mounted on the turn-table (Fig. 3). The cable for connecting them was made as short as possible to reduce noise. The stepper motor was used to rotate the turn table, which was driven by an Arduino based controller with a microstep driver (JK1545, Jingkong Motor & Electric Appliance, Changzhou, China). A Matlab script was written to control the turn table and record the subsequent data.

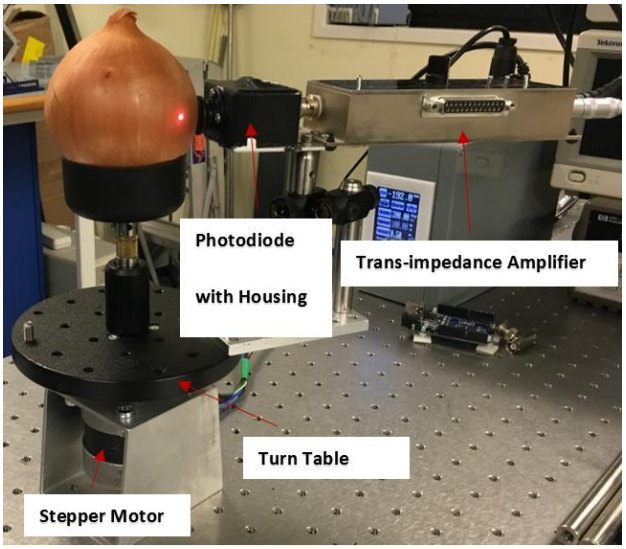

Fig. 3. Experimental setup of the laser light source measurements.

The onion was put on a cup so that it would stay stationary while the table was rotating. As the detector was mounted on the table, it is stationary in the reference of the onion. In each scan, the detector position was fixed and source (laser beam) location rotated around the onion as illustrated in Fig. 4. The table rotated 11.25° each time, which formed 21 source locations from 67.5° to 292.5°. For examining the whole onion, measurements were carried out on three positions: top, middle and bottom. In total, each onion was measured with 63 source locations and 3 detector positions.

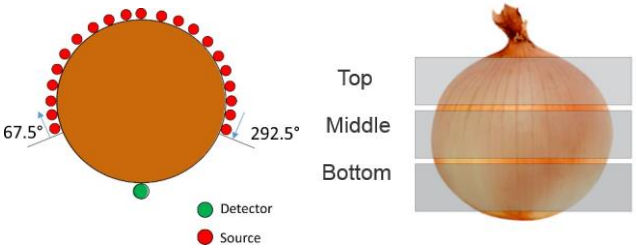

Fig. 4. Source and detect locations of laser light system

III. Results and Discussion

A. Broadband Light Source Measurements

There were in total 33 onions measured, 11 of each control, botrytis infected and pseudomonas infected. The spectra were normalised by a spectrum with reduced integration time taken when the onion was not on the light path. The normalisation allowed us to look at how the light was transmitted without the

effect from light source and detector. From the spectra, it was found that the light was attenuated less from about 670 to 850 nm which yielded low noise spectra in that range (Fig. 5).

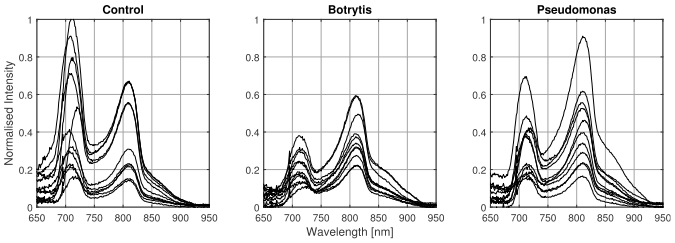

Fig. 5. Spectra of control, botrytis and pseudomonas from broadband light source measurements.

There were two peaks, the second peak was at about 810 nm which was more stable than the first one (Table I). Chlorophyll causes the instability of the first peak as it absorbs the light up to about 710 nm. The control onions had more green tissue so their 710 nm peak varied the most. Therefore, in order to avoid the interference from green tissue in onions, wavelength higher than 710 nm should be used.

TABLE I. WAVELENGTH OF THE FIRST AND SECOND PEAKS IN THE SPECTRA.

Onion Type	First Peak (nm)	Second Peak (nm)
Control	711 ± 4.2	809 ± 1.4
Botrytis	712 ± 1.3	812 ± 1.8
Pseudomonas	714 ± 3.3	811 ± 1.1

B. Laser Light Measurements

Three onions were used in the laser light measurements. Botrytis and pseudomonas infected onions had a large rotten area (Fig. 6) so that we could investigate the light transmission which was strongly affected by rotten tissue.

Fig. 6. Cross sections of the three onions measured: control (left), botrytis (middle) and pseudomonas (right).

The laser system allowed us to calculate the light transmittance through onion with the current optical system using the following equation:

$$Transmittance = \frac{V}{\mathrm{RG}P_{in}} \qquad (1)$$

Where V is voltage output of trans-impedance amplifier, P_{in} is the laser power, R is the responsivity of the photodiode, G is the trans-impedance gain of the transimpedance amplifier.

Three replicates were carried out for each scan. Error bar indicates the standard deviation (Fig. 7). The profile was not perfectly symmetric due to the shape and internal structure of the onions. It was found that the transmittance reduced exponentially when the incident angle was approaching 180°. This indicates that the light transmission reduced exponentially with the separation between source and detector (optode distance). It was lowest at around 180° where there was the largest optode distance.

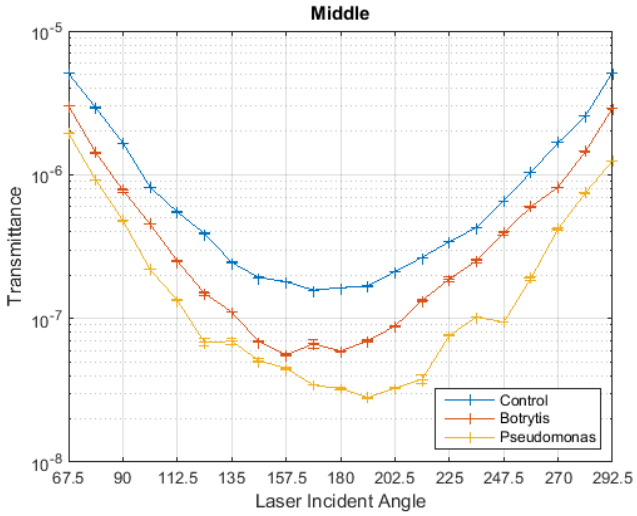

Fig. 7. Light transmittance at laser incident angles from 67.5° to 292.5° for the three onions using measurements at middle position.

For comparing the light transmission in different onions, table II tabulates the transmittance at light incident angle 180° (largest optode distance) for the three onions. The highest 180° transmittance was measured on the healthy onion at the bottom position, which was $1.81 \cdot 10^{-07}$. The rotten onions had lower transmittance than those of the healthy one, pseudomonas onion was lowest for middle and bottom positions ($3.25 \cdot 10^{-08}$ and $3.46 \cdot 10^{-08}$ respectively), and botrytis had transmittance of $3.26 \cdot 10^{-08}$ for top measurements, which was the lowest.

TABLE II. TRANSMITTANCE AT LIGHT INCIDENT ANGLE 180° FOR CONTROL, BOTRYTIS AND PSEUDOMONAS ONIONS UNDER THREE DIFFERENT MEASUREMENT POSITIONS.

Position	Control	Botrytis	Pseudomonas
Top	$1.54 \cdot 10^{-07}$	$3.26 \cdot 10^{-08}$	$4.52 \cdot 10^{-8}$
Middle	$1.64 \cdot 10^{-07}$	$5.93 \cdot 10^{-08}$	$3.25 \cdot 10^{-08}$
Bottom	$1.81 \cdot 10^{-07}$	$6.49 \cdot 10^{-08}$	$3.46 \cdot 10^{-08}$

Because the transmittance reduces exponentially with the optode distance, optode distance for each source and detector arrangement must be considered. The diameter of the onion at each sensor positon was measured with an assumption that an onion has a circular shape. The optode distance was then calculated for each light incident angle. For the control onion (similarly for other onions), from 67.5° to 180° and 180° to 292.5°, the transmittance changed roughly linearly with optode distance in a log scale (Fig. 8). A linear fit was applied to all data points at each sensor position, the slope was the attenuation coefficient (Figure 8 and Table III).

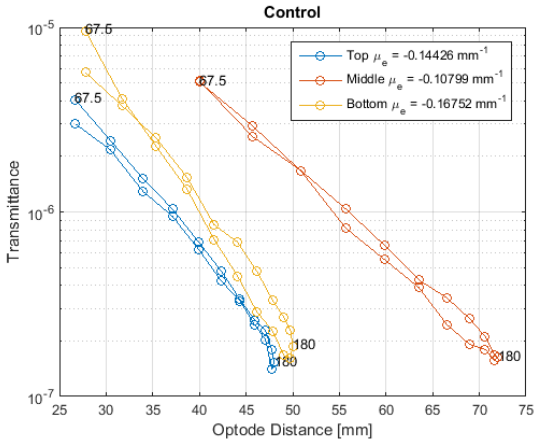

Fig. 8. Light transmission versus optode distance for the control onion at 728 nm.

The effective attenuation indicates how much of the light was attenuated per mm. The middle position had the lowest effective attenuation among all three onions, while the bottom position was the highest. Using control onion as a reference, the effective attenuation was directly linked to the degree of rottenness in onions. For example, botrytis onion had almost no rot at bottom position, its effective attenuation was close to that of healthy one. The effective attenuations of pseudomonas were all higher than those of control one as rot was fully developed.

TABLE III. EFFECTIVE ATTENUATION AT 728 NM FOR CONTROL, BOTRYTIS AND PSEUDOMONAS ONIONS.

Position	Control (mm^{-1})	Botrytis (mm^{-1})	Pseudomonas (mm^{-1})
Top	-0.144	-0.155	-0.163
Middle	-0.108	-0.123	-0.128
Bottom	-0.168	-0.173	-0.182

IV. CONCLUSION

We have studied the light transmission in healthy and rotten onions using two different systems. The broadband light source measurements showed a spectrum with double peaks (710 & 810 nm) in the detectable range of 670 to 850 nm. The laser system using 728 nm achieved highest transmittance of $1.81 \cdot 10^{-07}$ at the bottom positon. The transmittance was lower for the rotten onions (from $3.25 \cdot 10^{-08}$ to $6.49 \cdot 10^{-08}$). The effective attenuation was estimated, which was related to the degree of rottenness in onions. This work allowed us to have an understanding of signal attenuation in onions and the effect from the rots, so that we will be able to design our optical source and detector accordingly.

ACKNOWLEDGMENT

We would like to acknowledge Keith Sharrock for inoculating the onions. Jason Sun acknowledges the financial support of PhD scholarships from Plant and Food Research and University of Waikato.

V. REFERENCES

[1] B. M. Nicolaï, K. Beullens, E. Bobelyn, A. Peirs, W. Saeys, K. I. Theron, *et al.*, "Nondestructive measurement of fruit and vegetable quality by means of NIR spectroscopy: A review," *Postharvest Biology and Technology,* vol. 46, pp. 99-118, 2007.

[2] V. A. McGlone, P. J. Martinsen, C. J. Clark, and R. B. Jordan, "On-line detection of Brownheart in Braeburn apples using near infrared transmission measurements," *Postharvest Biology and Technology,* vol. 37, pp. 142-151, 8// 2005.

[3] B. P. Khatiwada, K. B. Walsh, and P. P. Subedi, "Internal defect detection in fruit by using NIR spectroscopy," 2016, pp. 337-342.

[4] W. Wang and C. Li, "Measurement of the light absorption and scattering properties of onion skin and flesh at 633 nm," *Postharvest Biology and Technology,* vol. 86, pp. 494-501, 12// 2013.

[5] E. K. Kemsley, H. S. Tapp, R. Binns, R. O. Mackin, and A. J. Peyton, "Feasibility study of NIR diffuse optical tomography on agricultural produce," *Postharvest Biology and Technology,* vol. 48, pp. 223-230, 5// 2008.

Enhancing Students' Conceptual Understanding Using Visualisation Software and Collaboration

Elliot Varoy*, Nasser Giacaman and Gerard B Rowe
Department of Electrical
and Computer Engineering
The University of Auckland
Auckland, New Zealand
*evar872@aucklanduni.ac.nz

***Abstract*—Students entering into Engineering education at the tertiary level aren't prepared. They lack fundamental understanding of the requisite subjects, namely physics and calculus, resulting in a snowball effect of misunderstanding in the tertiary sector. In order to remedy this issue we need to focus on improving students' conceptual understanding instead of enforcing rigourous memory-based learning. This paper recognises two viable teaching methodologies, visualisation and collaboration, by proposing, developing and evaluating a mobile and tablet application that utilises these ideas to encourage deeper conceptual understanding in students.**

I. Introduction

Core weaknesses in students' conceptual understanding of pre-requisite Engineering knowledge have been exposed. Many tertiary institutes include diagnostic testing at the entry level of tertiary education . Most of this early diagnostic testing is based upon the secondary level curriculum, which includes exams students must pass to gain entrance to tertiary education. Despite this requirement, the diagnostic examinations have shown that students carry over a weak grasp of fundamental concepts from their secondary level learning [1]. These diagnostic tests often introduce concepts to students in a different perspective when compared with secondary level exams, exposing students' reliance on memorisation techniques in place of building conceptual understanding. This style of learning also results in a lower rate of information retention [2]. This issue has recently been brought to light in New Zealand education by the Tertiary Education Commission. There is doubt being cast upon New Zealand's national qualification for secondary school students, NCEA (National Certificates of Educational Achievement), and its ability to prepare students for higher education. This is especially prevalent in STEM (Science, Technology, Engineering and Mathematics) subjects and is heavily impacting the preparedness of students entering into engineering education.

Many of the current tools and teaching methodologies being used to support the NCEA curriculum are also enforcing learning based around memorisation. Students are primarily relying on textbooks and teacher notes to learn core concepts. While this is good for establishing initial familiarity with the subject, students are struggling to develop a deeper understanding and finding it difficult to apply this knowledge. Kinetic learning using practical work experiements is useful for developing understanding, but has underlying issues with organisation and structure within a classroom setting [3]. This can also be difficult to apply to more abstract concepts often found in STEM subjects.

This paper investigates methods to improve students' ability to retain knowledge and apply it in unfamiliar contexts to remedy the above situation. To fulfill this goal it is important to discover methods for encouraging conceptual understanding over rote memorisation. The educational methodologies applied in this paper have been underpinned by previous research into conceptual understanding. We will discuss two of these core concepts, collaboration and visualisation, and their relevance in secondary school education. With software rapidly becoming commonplace within the education system, there is space for novel applications of these concepts. This is explored through the development and testing of an educational application encouraging both visualisation and collaboration for students studying electromagnetism at the secondary level. This application has been evaluated with students at a New Zealand Secondary School using mobile and tablet devices.

II. Background

This section will discuss the existing literature exploring conceptual understanding and how to encourage critical thinking in students. As electromagnetism is being explored as a use case for the developed theories, there will also be a discussion on some common misconceptions students encounter which can be used as a focus point for the application development. We have evaluated existing software which utilises mobile and tablet technology to explore any potential areas for improvement.

A. Conceptual Understanding

In order to improve the ability for students to retain their secondary school knowledge, and apply it in new contexts, its important to establish strong conceptual understanding. This can be achieved by utilising a technique known as visualisation. Visualisation can be defined by 5 core categories, concrete, verbal, symbolic, visual and gestural [4] It can be beneficial to combine these canonical models to encapsulate multiple learning styles of students [5]. Verbal visualisation can be further defined to include the use of speech or text used to describe a model. This also builds upon the use of metaphors and analogies to give context to abstract concepts. The visual category involves the use of graphs, diagrams and animation, including both two dimensional and three dimensional models. People naturally develop mental models while learning, a mechanism which allows us to describe the purpose of a system, explain how it functions and predict future outcomes [6]. Students will develop unique mental models to aid their understanding of concepts, and it can be difficult to produce an external model that remains consistent between peers. Despite this limitation, when utilised correctly, visual models will help students simplify and conceptualise difficult real-world concepts [4]. STEM subjects are inherently tied with models, as they are required to represent and prove abstract and physical theories among researchers. Thus it is especially important for students interested in Science and Engineering to gain early familiarity with model based learning. This will encourage them to consistently use models and enhance their skills for developing and applying them. Current methodologies for applying visualisation are predominantly two dimensional diagrams and models. Teachers are faced with the difficulty of explaining complex three dimensional concepts through two dimensional mediums. One current solution is to incorporate concrete visualisation through the use of physical experiments. While this theoretically has potential for significant improvements in conceptual understanding, there are various underlying issues, with evidence suggesting they are actually unproductive [7]. Without technical lab assistants available, the instructors priority will shift from encouraging content acquisition to focusing on management and organisation [8] resulting in poor student engagement and learning.

Collaboration is also a powerful methodology for encouraging conceptual understanding. This type of social interaction motivates students to verbalise their current understanding of concepts. This requires them to elaborate on what they perceive, to accurately explain how, or why, phenomena occur. These explanations help to consolidate their own learning, while simultaneously offering a new perspective for their peers [9], [10]. Collaborative situations also encourage students to negotiate their own interpretations of the content [11]. This is helpful for eliminating common misconceptions among students who have misinterpreted the literature, as the arguments for their logic will be faulted. Classical classroom teaching isolates students and suppresses their natural inclination to socialise [10]. Teachers will attempt to remedy this by engaging students with global questions. However, this methodology will often encourage specific students to engage with the teacher while others feel no obligation, or are too nervous, to participate [12]. Sectioning students into smaller groups can remedy this situation and eliminate the issue of students not participating but still retaining diversity in thinking styles [10]. Gokhale's research into collaborative learning [9] concluded, with statistical significance, that students learning through collaborative work performed better at tasks requiring critical thinking than those who studied individually. While collaboration encourages these aspects of learning, it must be introduced carefully. If groups are too small, there will be a lack in diversity which limits critical discussion [9], with larger groups reintroducing the issue of non-participation. Arguments between peers are a powerful tool for discussion but can be destructive if there are consistent disagreements or clashing personalities and must be accounted for.

B. Common Misconceptions in Electromagnetism

To implement the techniques for conceptual understanding discovered in the literature review, and evaluate the application of these techniques, we have used electromagnetism as a use case subject. We have defined some core misconceptions that are common in secondary school students studying electromagnetism, and aim to remedy this by utilising the methodologies above. A study [13] performed on students comprehension of electricity and magnetism discovered students had difficulty recognising that charges without a velocity perpendicular to a magnetic field wouldn't experience a force. Two dimensional representations of this require students to attempt to visualise an imaginary field in or out of the paper to recognise a perpendicular action, which could be a direct factor regarding this misconception. The study also discovered the concept of induction also faced similar difficulties. Student can recognise that movement is required for induction to occur, but fail to understand that the resulting change in flux is what is required for induction to occur. This misconception was also identified by Secondary School teachers involved in the evaluation introduced in this paper. The two Secondary School teachers involved also highlighted the difficulties surrounding the DC motor and Generator, specifically the impact of a split-ring commutator. The motor or generator is inherently a three dimensional and dynamic construct, making it extremely difficult to model using the traditional pen and paper approach. Both teachers outlined that students struggled to recognise how the commutator was utilised in the circuit, defaulting to simply memorising how it impacts the system. This was identified as a key aspect to investigate when implementing the visualisation model with our software development.

III. Implementing Visualisation and Collaboration Software

The review of literature surrounding conceptual thinking and misconceptions in electromagnetism education has helped to define a scope for developing a novel solution to remedy the limited understanding of students entering tertiary education.

The goal is to determine if implementing a software application, which incorporates both visualisation and collaboration methodologies, can effectively improve students' conceptual understanding of STEM based subjects. A mobile and tablet based application has been developed, using physics, specifically electromagnetism, as a focus subject. The underlying mechanics of the application will be applicable to other STEM subjects.

The developed application combines two categories of visualisation, visual and verbal, to present educational concepts. The visual aspect is focused on remedying the current issues with two dimensional representations that is rampant in Secondary School education. We have developed a framework that allows activities to be created which display different electrical circuits and components. Students are able to interact with the activity using multiple touch gestures to translate, rotate and zoom to change their perspective of the circuits involved. This allows students the freedom to explore concepts in three dimensions, from any position they wish to move the camera, rather than being restricted to a single two dimensional perspective in one direction. This is designed with the purpose of letting students conduct self-learning through exploration, altering their view to focus on the specific concepts they are struggling with.

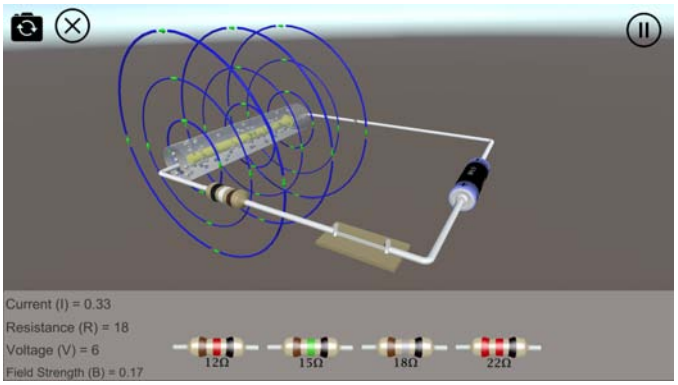

Fig. 1. Screenshot of activity visualising magnetic field generated by an electrical current

Each activity also has a user interface located at the bottom of the screen as shown in Figure1. Students have the ability to touch components within the activity to enable configurable options in the user interface. As an example, a circuit resistor may be touched to display various resistances to apply. Students will be able to alter components of the circuit, and observe in real-time how the circuit is affected by their changes. This adds an extra element to the visualisation when compared with the use of static diagrams to express state changes. With a multitude of different variables to change, students are capable of creating various different states to establish consistent patterns of logic. This pattern recognition will help students link multiple concepts and encourage them to recognise how and why different components impact each other, rather than memorising specific interactions. We have also included a list of relevant circuit variables that will vary as components are interacted with, to aid in recognising relationships between components.

The verbal visualisation is introduced through use of text, also encouraging the collaborative aspect of the application. As students interact with components in an activity, they may be prompted with a question regarding their changes, as shown in Figure 2. This style of questioning encourages students to formulate their own theories by building upon their established knowledge, to predict what will occur. Following their submission, students are prompted to discuss their answer with their peers. As students have already developed their own perspective of what will occur, encouraging collaboration at this will ideally establish constructive arguments. If students have conflicting answers, they will need to explain their theories with logic. This will encourage deeper levels of thinking by both the active (speaking) and passive (listening) student further enhancing their conceptual understanding [14]. Students also maintain the ability to manipulate the camera perspective to utilise the visual aspects of the application during the collaboration phase. After a discussion has occurred, students are offered the option to alter their original answer, and following this are presented with the correct answer. The application will then enforce the change made by the student, allowing them to actively see the visual representation of the answer.

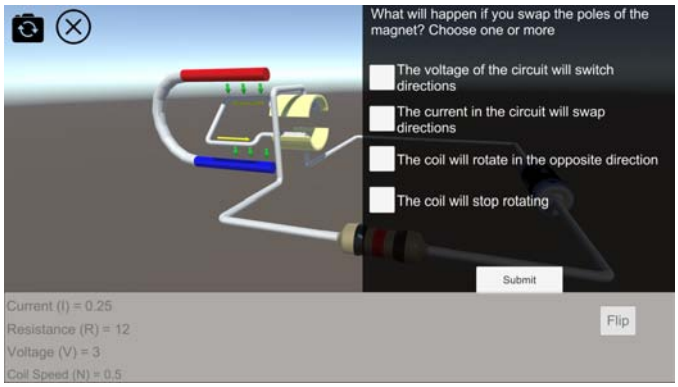

Fig. 2. Screenshot of question prompt when student attempts to flip the magnet poles

The completed prototype application included the following activities, each targeting a specific electromagnetic concept:

- Magnetic fields generated by a current
- Effects of a solenoid in a direct current circuit
- Induction effect experienced when moving a magnet through a solenoid
- Direct current motor
- Electricity generator

These activities were organised based on the principle of scaffolding and progression [15], where easier concepts are introduced prior to more complex concepts. This was also effective as later activities required pre-requisite knowledge

which was encountered in the earlier activities. Later activities will remain unavailable to students until they have completed the requisite activity. This is done by interacting with all components in the activity until all potential questions have been encountered at which point the user is notified that new content is available. This allows us to manage the concepts being accessed by students, preventing the risk of cognitive overload.

IV. Evaluation Methodology

To determine the effectiveness of this application on improving students conceptual understanding, an evaluation was conducted with a New Zealand Secondary School. Two Year 12 classes studying NCEA Level 2 Physics were used for this evaluation. The students had recently completed their electromagnetic education and hence it was assumed they had knowledge of the basic concepts of electromagnetism. A total of 40 students were involved in this study, with the cohorts being split into two classes of 23 and 17. The researcher conducted two 1-hour sessions with each of the classes with the following method. At the beginning of the first session, students were given a preliminary test consisting of 10 multiple choice questions to evaluate their understanding of electromagnetic concepts. These questions were based off previous research by the physics-education community ([1], [13]). An example of these pre-test questions can be seen in Figure 3.

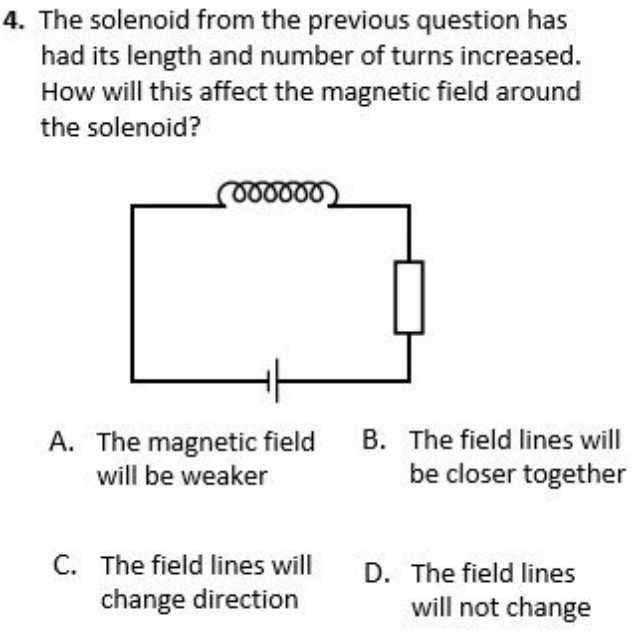

Fig. 3. Example of a question used in the preliminary test

After completing the test, android devices preloaded with the application were distributed among the class. Students were required to share each device between two to three members which aided with collaboration between peers. The mechanics of the application were explained, though they are aligned with common gesture functionality and were picked up quickly. The students were then free to complete the first three activities in the application. This process remained fairly unstructured, allowing students to self-direct their learning and explore the simulations from their own perspectives. The students' learning was also guided by the embedded questions which ensured students were exposed to the common misconceptions discovered in the literature. The second session began with students completing the final two activities in the application involving direct current motors and electrical generators. Following this, the students were given a post-test based on similar content to the previous test. An example of the comparison questions can be seen in Figure 4 and Figure 5. Formatting the question content in this fashion helps the validity of comparing knowledge of concepts prior to and following the use of the application. The study was concluded with a voluntary questionnaire which included 19 questions using a likert scale to assess the students learning experience and the ease of use of the application. Also included were four open-ended questions which allowed students to expand on their experience with the application.

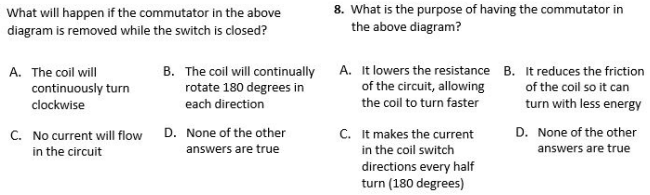

Fig. 4. Preliminary test question

Fig. 5. Post test question

V. Results

Our research question was to define whether there is a significant difference in marks following the use of the developed application. The overall results for the two classes of students can be found in Table I. The mean for all students on the pre-test was 4.48, with the post-test rising to a mean of 5.23. While 13 students saw reduced marks in the post-test, the overall mean difference was an increase of 0.75. In order to determine whether this increase is of statistical significance, a paired t-test has been performed.

We can define the null hypothesis as:

- H_0: No difference in mean between pre-test and post-test marks.

And the alternative hypothesis as:

- H_1: A difference in mean between pre-test and post-test exists.

The graph displayed in Figure 6 depicts the difference in scores per question, between the pre-test and post-test. We can see a maximum improvement on question 8, with the number of marks gained here increasing by 24. There is also a large negative difference on question 5 totaling -22. Question 2 and

TABLE I
STUDENT SCORES IN PRE-TEST AND POST-TEST OUT OF 10

Class	Pre-Test	Post-Test	Difference
1	2	3	1
1	3	6	3
1	4	5	1
1	5	6	1
1	4	5	1
1	5	4	-1
1	2	2	0
1	3	6	3
1	5	7	2
1	3	2	-1
1	2	4	2
1	6	4	-2
1	5	5	0
1	4	4	0
1	6	8	2
1	5	4	-1
1	5	3	-2
2	4	4	0
2	5	8	3
2	4	3	-1
2	7	10	3
2	5	5	0
2	5	8	3
2	6	9	3
2	6	7	1
2	4	1	-3
2	8	9	1
2	6	5	-1
2	5	9	4
2	7	9	2
2	6	9	3
2	4	2	-2
2	3	7	4
2	3	2	-1
2	1	4	3
2	5	2	-3
2	2	1	-1
2	4	1	-3
2	4	6	2
2	6	10	4
Mean	4.48	5.23	0.75

question 9 also have significant differences between tests and can be explored further. The remaining questions have little variance and this is expected between tests if no impact has been made on that particular concept.

VI. DISCUSSION

From the results in Table II we can recognise the p-value founds is less than 0.05 and thus can reject the null hypothesis of having no difference in mean score between pre and post test marks. Considering the lower value of 0.0950 and upper value of 1.405 for our confidence interval, we can say with 95% certainty that the mean difference in score lies between these values, and hence it is likely it will fall in the positive region. This is significant evidence for the application of collaboration and visualisation with a mobile or tablet and it's potential to positively impact students' conceptual understanding. However, as all students were required to participate in the study, we were unable to form a control group. This lowers the validity of our claim and limits our potential to extrapolate further.

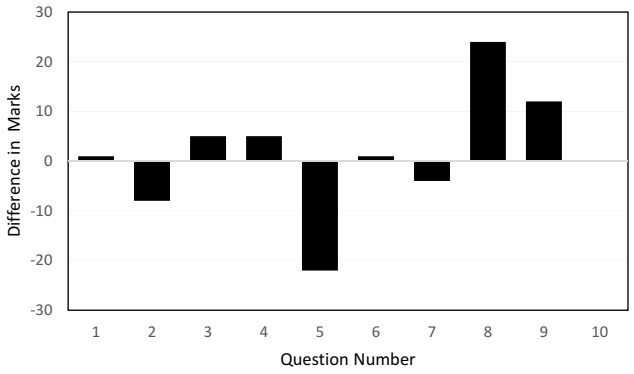

Fig. 6. Column Graph Representing Difference in Marks of Pre and Post Test by Question

The overall breakdown of marks was a net positive result, however the question by question break down was more interesting. Question 8 received the highest improvement after students used the application. This question focused on the use of a commutator within both generators and direct current motors. Both secondary school teachers highlighted this as a core issue early on in the development process, so specific attention was paid to ensure we could assess this misconception in particular. The extremely large net gain in marks is a strong indicator that this type of software implementation has a powerful impact on concepts that are commonly difficult for teachers to explain through traditional two dimensional means. We are unlikely to see as impressive results for other concepts in the test, as while they are complex in nature, they don't face the same dynamic, three dimensional challenges that commutator explanations do.

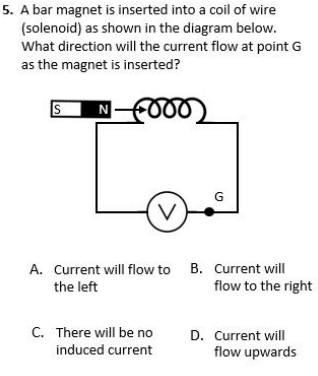

Fig. 7. Preliminary test question 5

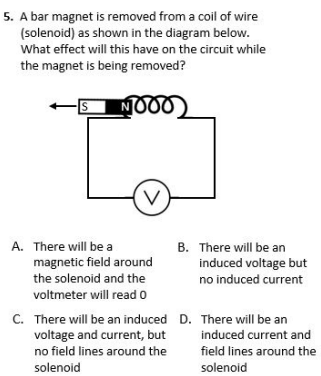

Fig. 8. Post test question 5

On the opposite end of the scale, question 5 had a very large decrease in correct answers. There is potential cause for concern here, as there is a chance that this application is negatively impacting student understanding for specific types of material. Upon revisiting the content of question 5, as shown in Figure 7 and Figure 8, it appears that while they are built upon the same concept of a magnet cutting a solenoid

TABLE II
PAIRED T-TEST (2-TAILED) RESULTS

				95% Confidence Interval				
	Mean	Std. Deviation	Std. Error Mean	Lower	Upper	t-value	df	p-value
Difference	0.750	2.170	0.340	0.0950	1.405	0.228	39	0.030

there are some distinct differences in the answers. To be able to accurately answer this question in the post-test, a much deeper understanding of multiple linking concepts is required. The pre-test version required only the knowledge of a rule governing the direction of induced current in this scenario. This makes it difficult to accurately assess whether the application has hindered students understanding, or if it is simply because of the difference in understanding required to answer each question.

VII. Conclusion

We have developed a mobile and tablet based application that incorporates the learning methodologies of visualisation and collaboration in an attempt to improve secondary school students' conceptual understanding of electromagnetism. The results of an evaluation involving 40 secondary school students have shown that there was a statistically significant improvement in student marks following the use of our software application. As the complex and model based nature of electromagnetism represents the wider content of STEM subjects, we believe that these results would be replicated in similar areas that lead to Engineering education. Improving the conceptual understanding of students at this level will better prepare them for tertiary study and help to alleviate the issues with current early undergraduates.

Acknowledgment

The authors would like to extend our thanks to the Secondary School faculty and students who partook in this study.

References

[1] C. R. Smaill, G. B. Rowe, E. Godfrey, and R. O. Paton, "An investigation into the understanding and skills of first-year electrical engineering students," *IEEE Transactions on Education*, vol. 55, no. 1, pp. 29–35, 2012.

[2] R. M. Felder and R. Brent, "Understanding student differences," *Journal of Engineering Education*, vol. 94, no. 1, pp. 57–72, 2005. [Online]. Available: http://dx.doi.org/10.1002/j.2168-9830.2005.tb00829.x

[3] S. Klainin, "Practical work and science education," *Development and dilemmas in science education*, pp. 169–188, 1988.

[4] J. K. Gilbert, *Visualization: A Metacognitive Skill in Science and Science Education*. Dordrecht: Springer Netherlands, 2005, pp. 9–27. [Online]. Available: http://dx.doi.org/10.1007/1-4020-3613-2_2

[5] B. Buckley, C. Boulter, and J. Gilbert, "Towards a typology of models for science education," *Exploring models and modelling in science and technology education*, pp. 90–105, 1997.

[6] W. B. Rouse and N. M. Morris, "On looking into the black box: Prospects and limits in the search for mental models." *Psychological bulletin*, vol. 100, no. 3, p. 349, 1986.

[7] D. Hodson, "Experiments in science and science teaching," *Educational Philosophy and Theory*, vol. 20, no. 2, pp. 53–66, 1988. [Online]. Available: http://dx.doi.org/10.1111/j.1469-5812.1988.tb00144.x

[8] ——, "Philosophic stance of secondary school science teachers, curriculum experiences, and children's understanding of science: Some preliminary findings," *Interchange*, vol. 24, no. 1, pp. 41–52, 1993. [Online]. Available: http://dx.doi.org/10.1007/BF01447339

[9] A. A. Gokhale, "Collaborative learning enhances critical thinking," 1995.

[10] B. S. H. William Rau, "Humanizing the college classroom: Collaborative learning and social organization among students," *Teaching Sociology*, vol. 18, no. 2, pp. 141–155, 1990. [Online]. Available: http://www.jstor.org/stable/1318484

[11] C. van Boxtel, J. van der Linden, and G. Kanselaar, "Collaborative learning tasks and the elaboration of conceptual knowledge," *Learning and Instruction*, vol. 10, no. 4, pp. 311 – 330, 2000. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0959475200000025

[12] D. A. Karp and W. C. Yoels, "The college classroom: Some observations on the meanings of student participation." *Sociology & Social Research*, 1976.

[13] D. P. Maloney, T. L. OâĂŹKuma, C. J. Hieggelke, and A. Van Heuvelen, "Surveying studentsâĂŹ conceptual knowledge of electricity and magnetism," *American Journal of Physics*, vol. 69, no. S1, pp. S12–S23, 2001.

[14] N. M. Webb, "Student interaction and learning in small groups," in *Learning to cooperate, cooperating to learn*. Springer, 1985, pp. 147–172.

[15] D. Fisher and N. Frey, "Guided instruction: How to develop confident and successful learners." *ASCD*, 2010.

A system for non-destructively determining variation in acoustic velocity within a harvested tree stem.

Michael Frampton
School of Electrical and
Computer Engineering
University of Canterbury
New Zealand, 8025
Email: michael.frampton@pg.canterbury.ac.nz

Michael Hayes
School of Electrical and
Computer Engineering
University of Canterbury
New Zealand, 8025
Email: michael.hayes@canterbury.ac.nz

***Abstract*—Acoustic methods are increasingly being employed by the forestry industry to determine wood quality. Commonly, acoustic velocity is used to determine the stiffness of a harvested tree stem, which is useful for grading and sorting purposes. Existing systems provide a single value for the velocity of a stem. However, a stem does not have a uniform velocity; it varies considerably over its length. This paper investigates two different approaches for non-destructively finding the variation in velocity in a stem. An experiment was performed to capture acoustic signals along the length of a complete stem, which was then sawn into segments and further measurements taken. An analysis was performed on the complete-stem data to find the velocity variation, however the proposed algorithms were unsuccessful. Further development of these algorithms is required.**

I. Introduction

Foresters and wood-processors are interested in maximising the value of their stock. Typically, a stem is sawn up into approximately three logs of length 5 to 8 m. From these logs a number of finished products can be produced including veneer, timber, pulp, packaging, or roundwood posts. Each of which has different wood requirements, and each has a different return value for the seller. The return on structural wood can be double that of lower value products such as pulp/paper and packaging products [1]. Unfortunately, the amount of structural wood being produced is as low as 10% of the total yield [2]. Thus, it is sensible to explore ways of maximising the proportion of structural wood produced from a stand of trees. Traditionally, the industry has relied on human inspection of a harvested stem to determine its value [3]. Clearly this is not ideal since it requires experience, is prone to errors, and provides little accountability to recipients of the product. More recently, acoustic tools have been employed in production lines to measure acoustic velocity [4].

Significant variability is observed when assessing the wood quality of any harvested stand of trees [3]. This is true between families of trees, within families, and even between clones grown in similar conditions [5]. Additionally, there is significant variation within a harvested stem. Several authors have documented the variation in wood stiffness within harvested stems, measured destructively [6], [7]. Carter et al. measured the resonant velocity of 6 stems. They then sawed the stems into boards and measured the resonant velocity of each board. They found that the velocity of the total stem is equal to the weighted-average of the velocity of each individual board (the law of mixtures) [8]. They found that the green Modulus of Elasticity (MoE) is a good predictor of the dry MoE ($\mathrm{E}_{\text{air-dry}} = 1.06\ \mathrm{E}_{\text{green}}$, $R^2 = 0.86$).

Walker and Xu [9] sawed 62 *pinus radiata* trees into boards which were tested using machine stress grading (static bending test). This permitted a stiffness distribution map to be created for this species. They found that the least stiff wood occurs in the pith of the stem and the most stiff at the outerwood. In addition, they found that the lowest 2.5 m of a tree contains a higher proportion of low stiffness wood, probably due to the large amount of heartwood in that region.

As far as the authors have found, there has not yet been any effort to determine the longitudinal variation in stiffness of a log non-destructively. If this were possible, the technology could prove very attractive. Wood processors would more effectively be able to assign stock to the appropriate products and thus maximise the value.

This paper explores methods of determining the acoustic velocity variation within a harvested stem. A system of capturing acoustic signals from a harvested stem is described in Section II. Two different approaches of processing the captured data are assessed in Section III and Section IV. An experiment was performed to capture acoustic signals from a harvested stem, described in Section V. Results are provided in Section VI. Finally, conclusions are drawn in Section VII and suggestions provided for further work.

II. Proposed system

The proposed system will use an array of transducers to measure the variation in acoustic velocity along the length of a stem. The system will place a number of acoustic transducers in contact with the stem, along its length, each separated by a known distance. The stem is then struck at the base to

induce a stress wave. A sampling device is used to capture the signal at each transducer. By measuring the time taken for the wave to propagate between the transducers of each segment, the incremental velocity can be determined. This would be a significant improvement over existing methods, which only provide a single velocity value for the entire log.

This method is an extension of existing acoustic systems, which come in two varieties, *time-of-flight* (ToF) and *resonance* systems [10].

The ToF and resonance methods estimate the stiffness by measuring the acoustic velocity within the wood. The relationship between stiffness and speed for a plane-wave propagating in an isotropic lossless-medium is given by

$$E = c^2 \rho, \tag{1}$$

where E is the MoE (also known as the *stiffness*), c is the wave speed and ρ is the wood's density.

The ToF and resonance methods, and their application to this problem are further discussed in Section III and Section IV.

III. The Time-of-Flight method

ToF systems measure the time taken for an acoustic wave to travel between two transducers separated by a known distance z. After placing the transducers in the wood an stress wave is excited either using an electrically driven transducer, or by striking the wood with a hammer. The propagation time between the transducers τ is then measured, and the speed $c = z/\tau$ is calculated.

The difficulty of this method is locating exactly when the waveform starts. The simplest method is to define a threshold level. When the amplitude exceeds that level, the start of the signal is tagged. Alternatively, a peak finding algorithm can be used to find the location of the first peak in the signal. The propagation time is difficult to determine because wood tends to attenuate higher frequencies [11]. This causes the signal's rise-time to be extended, which causes the measured delay to increase. As a result of wood's attenuation, it is also dispersive, i.e. different frequencies travel at different speeds. When making a measurement it is not clear exactly which frequency is being measured. Figure 1 shows an example of two signals measured at different distances along the stem. Attenuation of high-frequency components has modified the wave's shape and slowed the rise-time.

IV. The resonance method

The resonance method uses a single transducer placed against the end of a harvested stem. The operator strikes the end of the stem with a hammer which causes a stress wave to propagate. The signal reflects back and forth along the stem, establishing a standing wave, the frequencies of which are given by

$$\omega = \frac{n\pi c}{L}, \tag{2}$$

(assuming a lossless medium), where $n = 1, 2, \ldots$. These frequencies are known as the *resonant harmonics* of the system, or simply the *harmonics*.

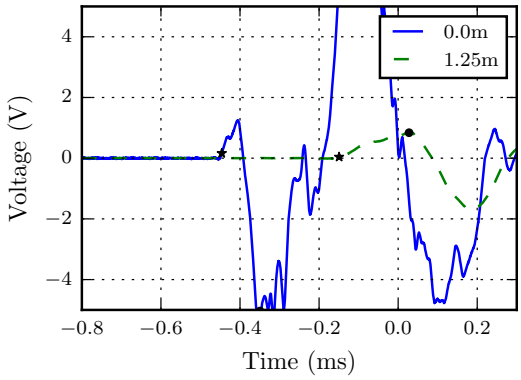

Fig. 1. This plot illustrates the difficulty of the ToF method. Two signals were measured at the indicated distances from the base of the stem. As the signal propagates through the medium it is attenuated, particularly at higher frequencies. The ⋆ indicates the detected threshold time. The • indicates the detected peak time.

In the following section, a derivation is provided for the transfer function of a signal propagating using the resonant method. The transfer function is generalised to allow the signal to be evaluated at any location along the length of the stem. Next, a method is proposed which may allow the resonance method to be applied to the problem at hand.

A. Transfer function derivation

In the following section, a transfer function is derived for the amplitude and phase of a stress-wave excited in the stem. Consider a signal $x(t)$ generated at one end of the stem. The signal will propagate to the other end where it will be reflected back. In the time domain, the propagating signal is of the form

$$\begin{aligned} y(z,t) = {} & h(z,t) * x(t) \\ & + (h(L,t) * x(t))\Gamma_{\mathrm{e}} * h(L-z,t) \\ & + [(h(L,t) * x(t))\Gamma_{\mathrm{e}} * h(L,t)]\Gamma_{\mathrm{s}} * h(z,t) \\ & + ([(h(L,t) * x(t))\Gamma_{\mathrm{e}} * h(L,t)]\Gamma_{\mathrm{s}} * h(L,t)) \\ & \qquad \Gamma_{\mathrm{s}} * h(L-z,t) \\ & + \ldots . \end{aligned} \tag{3}$$

Where $h(z,t)$ is the impulse response of the wood, $x(t)$ is the applied input signal, L is the length of the stem, Γ_{s} and Γ_{e} are the start and end reflection coefficients, respectively, and $*$ denotes the convolution operator. The first term is the first *incident* wave (forward travelling), the next is the *reflected* wave (backward travelling), followed by incident, reflected, *ad infinitum*.

Fourier transforms are then taken on both sides.

$$\begin{aligned} Y(z,f) = {} & H(z,f)X(f) && \text{(incident)} \\ & + H(L-z,f)H(L,f)\Gamma_{\mathrm{e}}X(f) && \text{(reflected)} \\ & + H(z,f)H^2(L,f)\Gamma_{\mathrm{e}}\Gamma_{\mathrm{s}}X(f) && \text{(incident)} \\ & + H(L-z,f)H^3(L,f)\Gamma_{\mathrm{e}}\Gamma_{\mathrm{s}}^2X(f) && \text{(reflected)} \\ & + \ldots \end{aligned} \tag{4}$$

This can be divided by $X(f)$ to obtain a transfer function for the system $T(z, f)$. This can be expressed as a sum

$$T(z,f) = H(z,f)\sum_{n=1}^{\infty} H^{2(n-1)}(L,f)\,(\Gamma_{\mathrm{e}}\Gamma_{\mathrm{s}})^{n-1} \quad \text{(incident)}$$
$$+ H(L-z,f)\sum_{n=1}^{\infty} H^{2n-1}(L,f)\Gamma_{\mathrm{e}}^{n}\Gamma_{\mathrm{s}}^{n-1}. \quad \text{(reflected)} \tag{5}$$

The reflection coefficients will closely approximate unity $\Gamma_{\mathrm{e}}, \Gamma_{\mathrm{s}} \approx 1$ because the wood's acoustic impedance is much greater than that of air. (5) is a geometric series, the sum of which is given by

$$T(z,f) = \frac{H(z,f)H(L,f)}{1-H^2(L,f)} + \frac{H(L-z,f)H(L,f)}{1-H^2(L,f)}. \tag{6}$$

The wood's one-way transfer function is assumed to be of the form

$$H(z,f) = e^{-\gamma z}, \tag{7}$$

where $\gamma(f)$ is the *propagation coefficient* of the wood, and is a complex, frequency dependent function. Making these substitutions, the transfer function becomes

$$T(z,f) = \frac{e^{\gamma(L-z)} + e^{-\gamma(L-z)}}{e^{\gamma L} - e^{-\gamma L}} \tag{8}$$
$$= \frac{\cosh\gamma(L-z)}{\sinh\gamma L}. \tag{9}$$

The propagation coefficient can be expanded into real and imaginary components $\gamma = \alpha + jk$. α is known as the *attenuation coefficient*. $k = \omega/c$ is the *angular wave number*, with units of radians per metre. For this analysis it is assumed that the attenuation in the stem is small, or $\alpha \approx 0$. The transfer function then reduces to

$$T(z,f) = \frac{\cos\left(\frac{\omega}{c}(L-z)\right)}{j\sin\left(\frac{\omega}{c}L\right)}. \tag{10}$$

The poles of this transfer function occur at the frequencies given by (2). This is a generalised transfer function for the signal response at any point along the length of the stem. The transfer function's magnitude and phase for the first three harmonics are plotted in Figure 2. This plot shows the transfer function in the lossy case (9).

B. Reflected wave cancellation

The stem transfer function (10) does not have linear phase with respect to distance, as demonstrated by Figure 2(b). The fundamental harmonic has a $-\pi$ phase shift at $z = L/2$. If the reflected terms of (4) are removed, then the sum of the series reduces to,

$$T_{\mathrm{RC}}(z,f) = \frac{H(z,f)H(L,f)}{1-H^2(L,f)}. \tag{11}$$

As in the previous derivation, the one-way wood transfer function (7) can then be substituted. The transfer function becomes

$$T_{\mathrm{RC}}(z,f) = \frac{e^{-\gamma z}}{\sinh\gamma L}, \tag{12}$$

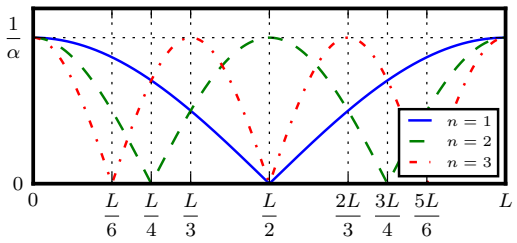

(a)

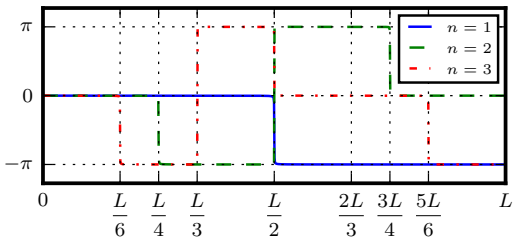

(b)

Fig. 2. The magnitude (a) and phase (b) of the stem's transfer function $T(z, f)$ are plotted against distance along the stem, evaluated at the frequencies of the first 3 harmonics.

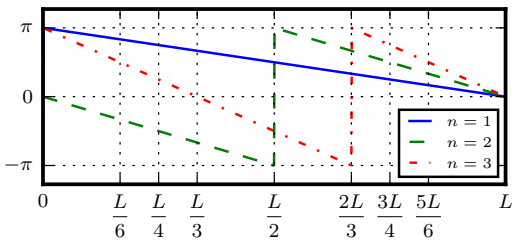

Fig. 3. The phase of the reflected wave cancelled model (12) for the first three harmonics.

which has a linear phase with respect to distance. This *reflected wave cancellation* has linear phase and constant magnitude. The phase for the first three harmonics is plotted in Figure 3.

Of course, this method is dependent on being able to cancel the reflected wave components.

V. Experimental Method

An experiment was performed to capture a series of waveforms along the length of an 8 metre *pinus radiata* stem. The experiment was performed in two separate stages. Firstly, a series of waveforms were captured from the complete stem. Secondly, the stem was sawn into 1m segments, on which resonance measurements were performed.

For the first stage, referred to as the *stem* experiment, four Fakopp SD-02 acoustic transducers were employed. A diagram of the experimental setup is shown in Figure 5. For clarity, the labels *start*, *forward*, *reverse* and *end* were assigned to

Fig. 4. The 8m *pinus radiata* stem.

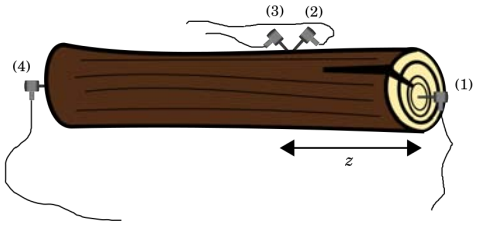

Fig. 5. The experimental setup. The transducers are labelled from the right: *start* (1), *forward* (2), *reverse* (3), *end* (4).

the transducers. The forward and reverse transducers were placed at an angle of 45° to the stem. This allows the *inplane* (longitudinal) and *outplane* (radial) signals to be deduced. The inplane signal is given by

$$y_{\text{inplane}} = \frac{y_{\text{forward}} - y_{\text{reverse}}}{2}. \tag{13}$$

The out of plane signal is given by

$$y_{\text{outplane}} = \frac{y_{\text{forward}} + y_{\text{reverse}}}{2}. \tag{14}$$

The four transducers were attached to the inputs of a four channel Tektronix DPO3014 oscilloscope. The base of the log (near the start transducer) was struck with a hammer to induce a stress wave in the stem.

The forward and reverse transducers were moved up the stem in 250 mm increments. At each increment the stem was struck 5 times and a signal captured for each strike.

Preliminary testing revealed that repeated striking of the base of the stem caused permanent damage to the wood. This creates a significant delay and damping on the signal in the damaged region. In response to this, a metal plate was fixed to the end of the stem. This has the effect of spreading the hammer blow out over a larger area in an attempt to reduce the amount of damage.

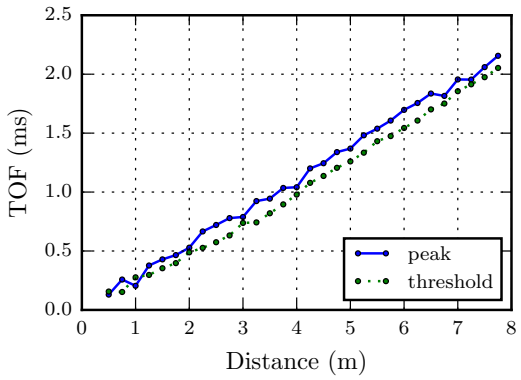

Fig. 6. ToF from the start transducer to the inplane signal, measured along the length of the stem.

In the second part of the experiment, referred to as the *segment* experiment, the stem was sawn up into 1 m segments. In this experiment the start and end transducers were inserted into each face of the sawn segment. The stem was then struck at the end closest to the base. The signal for each transducer was captured and recorded on the oscilloscope.

VI. Analysis and Results

As discussed in the preceding sections, two different methods were assessed to determine the speed variation along the stem, the time-of-flight method and the resonance method with reflection cancellation. In addition, the ToF and resonance methods were used to measured the acoustic speed of the cut segments as a comparison for the other methods.

A. Time-of-flight results

Figure 6 shows the time-of-flight from the start transducer to the inplane signal using both threshold and peak methods. As the transducers were moved along the stem, the peak method measured a longer ToF. As expected, both methods show an approximately linear trend with distance.

As described in Section V, the inplane signal was captured in incremental steps, not simultaneously. In a production system the signals would be captured at a number of positions along the stem simultaneously. This was not possible in this experiment due to the limited number of transducers and capturing channels on the available oscilloscope. The *cumulative* speed is the speed from the base to some distance along the stem

$$c_{\text{cumulative}}[i] = \frac{z[i]}{\tau[i]}, \tag{15}$$

where i is measurement number, $z[i]$ represents a vector of all measured distances, and $\tau[i]$ is a vector containing the ToF from the base to that distance.

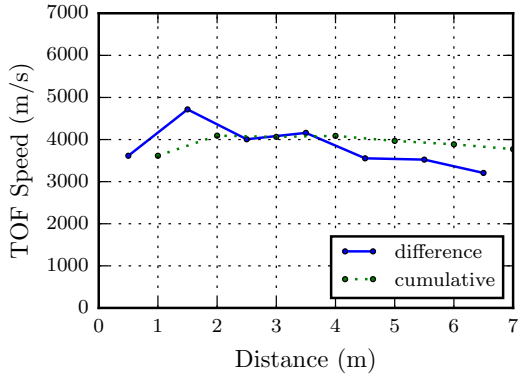

Fig. 7. Speed measured along each segment using the threshold method. Both cumulative and differential speeds are plotted. The measurements were analysed at 1 m increments. The differential ToF is plotted at the midpoint of measurement positions.

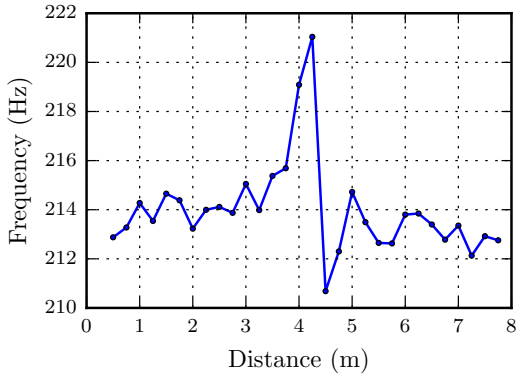

Fig. 8. The measured resonant frequency of the fundamental harmonic without reflection cancellation. The node in the centre of the stem results in a reduced signal magnitude which increases the uncertainty in the measured frequency.

To determine the variation in speed along the length of the stem, the difference between adjacent ToF measurements was taken

$$c_{\text{diff}} = \frac{z\,[i] - z\,[i-1]}{\tau\,[i] - \tau\,[i-1]}. \tag{16}$$

The cumulative and differential speeds are plotted in Figure 7. The signals were captured at intervals of 0.25 m, but the stem was later sawn into 1 m segments, thus the differential and cumulative ToF was evaluated at 1 m increments to enable comparison with the segment results.

B. Resonance results

As discussed in Section IV, the resonance method can be used to determine the speed variation within the stem if the reflected wave components are cancelled. The following section describes the method used to cancel the reflected components, and the results obtained.

The method used to cancel the reflected wave components is to nullify (set to zero) the parts of the waveform which contain a backward travelling wave and do not overlap part of a forward travelling wave. Consider a wave measured at a location $z = z_i$ along the stem. If the stem's wave speed is assumed constant then the first incident wave arrives at $\tau_{I0} = z_i/c$. The first reflected wave arrives at $\tau_{R0} = (2L - z_i)/c$ and so on. These arrival times can be generalised as

$$\tau_{In} = \frac{2Ln + z_i}{c} \quad n = 0, 1\ldots \tag{17}$$

$$\tau_{Rn} = \frac{2L\,(n+1) - z_i}{c} \quad n = 0, 1\ldots \tag{18}$$

The cancellation method zeros the parts of the waveform from $\tau_{R(n)}$ to $\tau_{I(n+1)}$. Thus, only a portion of the reflected wave is cancelled; the portion where the incident and reflected overlap. The speed used to calculate the arrival times was determined using the ToF measured at that particular distance. Local variations in speed mean that the arrival times of (17) and (18) are only an approximation of the actual arrival times. The resonant frequency without reflection cancellation is shown in Figure 8. This frequency should be constant with distance. The plot demonstrates that largest variation in the resonant frequency occurs near the node in the center of the stem, where the signal level is small. This is consistent with the derived transfer function of Figure 2.

To determine the phase variation along the length it is necessary to select a frequency at which the phase is evaluated. In this case, the average of the fundamental resonance shown in Figure 8 was used, approximately 213 Hz, denoted f_0. The phase was then evaluated at 1 m increments, this is shown in Figure 9.

The phase is referenced to the *start* transducer, placed at the start of the stem. If the reflection cancellation method worked successfully, then phase should decrease approximately linearly according to (12). The delay in seconds can then be derived from the phase angle's deviation from the reference.

$$\Delta\angle Y(z, f_0) = \angle Y(z, f_0) - \angle Y(z_0, f_0) \tag{19}$$

$$= \gamma z, \tag{20}$$

where γ is the aforementioned propagation coefficient. It is assumed that the loss in the medium is small, so the real part of γ is ignored,

$$\Delta\angle Y(z, f_0) = -kz \bmod 2\pi, \tag{21}$$

$$= -\frac{2\pi f_0}{c} z \bmod 2\pi. \tag{22}$$

A negative phase difference indicates that the signal is delayed relative to the start signal. Thus, due to causality, the signal should only become more negative as distance increases along the stem. Unfortunately this is not the case, the phase plotted in Figure 9 leads the reference signal for the first section from 1-3 m. Ideally, the phase would be used to calculate the speed using (22). This speed would then be compared against the speed measured destructively.

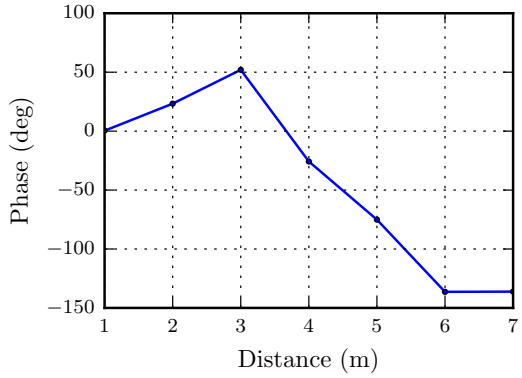

Fig. 9. The resonant phase of the inplane signal with reflected wave cancellation, measured at 1 m increments. The positive phase from 1-3 m indicates that the signal is leading the reference signal, which is not causal.

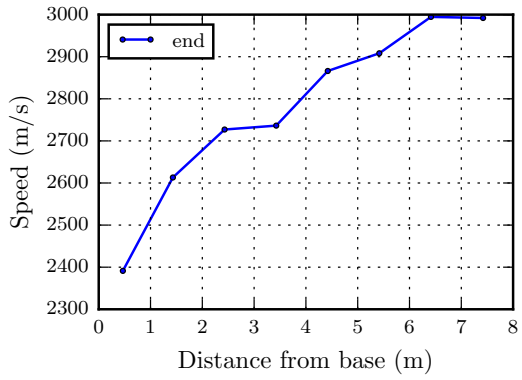

Fig. 10. The speed measured using the resonant frequency of the cut segments. Each data point represents the speed of a single segment, distances are plotted at the midpoint of the segment.

C. Segment results

The acoustic signals captured from the cut segments were used to determine the variation in speed along the length of the stem. The frequency of the fundamental harmonic was used to determine the acoustic speed. The speed is plotted in Figure 10. This plot shows that there is significant variation along the length of the stem.

The segment speed is significantly lower at the base of the stem than at 8 m. This is consistent with Walker and Xu's findings regarding the lowest harvested log, which has a higher proportion of heartwood and a higher microfibril angle [9].

VII. Conclusion

This paper explored some potential methods of non-destructively measuring the variation of acoustic velocity within a harvested tree stem. Two methods were considered: the time-of-flight method and the resonance method. The stem was then cut into 1m lengths so that the speed along each segment could be measured.

The ToF method measured speeds between $3200\,\mathrm{ms}^{-1}$ and $4800\,\mathrm{ms}^{-1}$. The segment resonance measured speeds of between $2400\,\mathrm{ms}^{-1}$ and $3000\,\mathrm{ms}^{-1}$, significantly lower. It is well established that the ToF method measures a considerably higher speed in wood [10]. The underlying mechanism of the difference between the two methods is not well understood, though several possibilities have been proposed. For further discussion see Legg and Bradley's review paper [10].

It was found that the ToF speed measured in the stem does not follow the same trend as the segment speed. This is probably due to the high uncertainty associated with the ToF method due to the difficulty of identifying the waveform start. For this reason, future work will focus on resolving the aforementioned issues with the resonance method.

A reflection-cancellation algorithm was applied to the signals in an attempt to use the spectrum-phase for calculating the wave speed. Unfortunately, the subsequent signals display a positive-going phase as the sensors were moved along the stem, a non-causal phenomenon. This may indicate a systematic error in the method used. Future work will attempt to resolve this issue.

Acknowledgment

The authors would like to thank the NZ School of Forestry for funding this research. Technical assistance provided by Lachlan Kirk and Nigel Pink has been invaluable in the execution of this project.

References

[1] L. A. Apiolaza, "Very early selection for solid wood quality: screening for early winners," *Annals of Forest Science*, vol. 66, no. 6, p. 601, 2009.
[2] D. Cown and L. van Wyk, "Profitable wood processing-what does it require? good wood!" *New Zealand Journal of Forestry*, vol. 49, pp. 10–15, 2004.
[3] X. Wang, P. Carter, R. J. Ross, and B. K. Brashaw, "Acoustic assessment of wood quality of raw materials: a path to increased of profitability," *Forest Products Journal*, vol. 57, no. 5, pp. 6–15, 2007.
[4] F. Divos, "Acoustic tools for seedling, tree and log selection," *The Future of Quality Control for Wood & Wood Products*, p. 5, 2010.
[5] B. Dyck, "Precision forestrythe path to increased profitability," in *Proceedings of the 2nd International Precision Forestry Symposium, Seattle, Washington; University of Washington, Seattle*, 2003, pp. 3–8.
[6] G. J. Searles, "Acoustic segregation and structural timber production," Ph.D. dissertation, Edinburgh Napier University, Edinburgh, Scotland, 2012.
[7] D. Cown, D. McConchie *et al.*, "Wood property variations in an old-crop stand of radiata pine," *New Zealand Journal of forestry science*, vol. 10, no. 3, pp. 508–520, 1980.
[8] P. Carter, S. Chauhan, and J. Walker, "Sorting logs and lumber for stiffness using director hm200," *Wood and Fiber Science*, vol. 38, no. 1, pp. 49–54, 2007.
[9] P. Xu and J. C. F. Walker, "Stiffness gradients in radiata pine trees," *Wood Science and Technology*, vol. 38, no. 1, pp. 1–9, 2004.
[10] M. Legg and S. Bradley, "Measurement of stiffness of standing trees and felled logs using acoustics: A review," *The Journal of the Acoustical Society of America*, vol. 139, no. 2, pp. 588–604, 2016.
[11] V. Bucur, *Acoustics of Wood*. Berlin, Heidelberg: Springer-Verlag, 2006.