

Optical Sensing, Embedded Systems, and Musical Interfaces for Functional Neuroimaging

Avrum Hollinger



Music Technology Area
Schulich School of Music
McGill University
Montreal, Canada

April 13, 2014

A thesis submitted to McGill University in partial fulfillment of the requirements for
the degree of Doctor of Philosophy.

© 2014 Avrum Hollinger

Abstract

Behavioural neuroscience experiments are increasingly making use of functional magnetic resonance imaging (MRI) and require new physical interfaces augmented with sensors and electronic control systems to capture behaviour in synchrony with the MRI scanner, while presenting stimuli and feedback. These interfaces permit neuroscientists to correlate changes in neural activity with changes in behaviour, and thus form a better understanding of the human brain.

Safety and MRI-compatibility are paramount, although repeated testing of new interfaces in the MRI scanner increases development costs. The inherent compatibility of fibre optic sensors ensures the safety and proper functioning of optoelectronic acquisition of behavioural measures within the MRI scanner, and will ensure their compatibility in newer, higher-field MRI scanners.

I have designed and built novel optical sensors, optoelectronic acquisition and control systems, and augmented MRI-compatible musical interfaces for use in musical neuroimaging experiments. The fibre optic sensors were created for the measurement of proximity, position, displacement, flexion, force, vibration, and magnetic field. The embedded systems that have been developed enable optoelectronic signal acquisition, logging, communication, and real-time presentation (including mapping and synthesis) of auditory feedback, with the goal of a self-contained system to perform behavioural experiments. These MRI-compatible musical interfaces include: the Ballagumi, a novel flexible silicone instrument; an optoelectronic piano keyboard; and an optoacoustic cello. The piano and cello have been utilized in neuroimaging experiments, allowing, for the first time, the capture of musician gestures during actual performance on musical instruments within the MRI scanner. Presented herein is the iterative development and testing of these technologies. Additionally, the literature on MRI-compatible musical interfaces and related fields is reviewed, as well as future developments of these technologies and their applications in behavioural neuroscience, neuro-rehabilitation, and music performance.

Résumé

L'imagerie par résonance magnétique fonctionnelle (IRM) est de plus en plus utilisée pour des expériences de neuroscience comportementale. De telles expériences nécessitent la création de nouvelles interfaces augmentées de capteurs et de systèmes de contrôle électroniques pour capter les comportements, présenter des stimuli et produire des retours d'information de manière synchrone avec le scanner. Ainsi, il est possible pour les neuroscientifiques d'établir des corrélations entre les changements de comportement et les changements d'activité cérébrale; ceci leur permet donc de développer une meilleure compréhension du cerveau humain.

Sécurité et compatibilité sont de plus haute importance, bien que des essais répétés de nouvelles interfaces dans le scanner IRM, afin de garantir qu'elle soit sécuritaire, augmente donc le prix de son développement. Les capteurs à fibre optique permettent d'atteindre ces objectifs avec les scanners actuels et les nouvelles générations de scanners qui produisent des champs magnétiques encore plus puissants.

J'ai développé et construit de nouveaux capteurs optiques, des systèmes d'acquisition optoélectroniques, et des interfaces musicales compatibles avec les scanners IRM dans le cadre d'études d'imagerie cérébrale. Les capteurs à fibre optique que j'ai conçus ont été créés pour mesurer la proximité, la position, le déplacement, la flexion, la force et le champ magnétique. Les systèmes embarqués développés durant cette thèse ont pour objectif de fournir des systèmes autonomes permettant de réaliser des expériences comportementales. À ce titre, ils permettent l'acquisition optoélectronique de signaux, leur communication, leur enregistrement et la production de stimuli auditif en temps-réel (incluant le mapping et la synthèse). Voici une liste des interfaces de contrôle musicales mises au point durant ce projet de doctorat: le Ballagumi, un nouvel instrument flexible fait en silicone; des claviers de piano optoélectroniques; et un violoncelle optoacoustique. Les pianos et le violoncelle ont été mis à contribution lors d'expériences ayant permis, pour la première fois, de capter les gestes d'un instrumentiste au sein d'un scanner. Dans ce qui suit, je présente le processus itératif de développement et de test ayant mené à ces nouvelles interfaces. De plus, l'état de l'art sur les interfaces musicales compatibles avec les scanners IRM est passé en revue, ainsi que les futures évolutions et applications de ces technologies aux domaines de la neuroscience comportementale, de la réhabilitation cérébrale et de la performance musicale.

Acknowledgments

Thank you to my supervisor, Marcelo M. Wanderley, for your support and encouragement throughout this process: for giving me the freedom to approach and solve problems in my own way, for being available at all hours, and for pushing me to go beyond my comfortable norms. Thank you Robert J. Zatorre and Virginia B. Penhune for this very fruitful collaboration, for your guidance, and your suggestions (and last-minute corrections). You have provided me with insight and elucidated the path on countless occasions.

Thank you to Melanie Segado for getting the cello project off the ground and into the scanner. Thank you to Rachel Brown, Michael Klein, and other collaborators, lab mates, and research assistants. Thanks to mechanical engineers Anthony Piciacchia and Steven Phillips for your tireless efforts and sharing your CAD modelling skills; this work would not have been possible without you. Other CAD models were drawn by Nabil Shalabi (in association with Larry Lessard and Zen Composites), Bryan Hynes (Hybex Innovations), Antoine Lefebvre, and David Munro. Thanks to cellists Michael Spleit and Erika Donald, luthier Richard Benoit, and bow maker Éric Gagné for your guidance and work on the cello. Thank you, Darryl Cameron, for your help, good cheer, and ensuring I was always surrounded by boxes of parts. Thanks as well to all the Music Technology Area professors. Thanks to Yves Méthot, Harold Kilianski, and Julien Boissinot of CIRMMT for your technical assistance and critical review. Thank you to Michael Ferreira, David Costa, Ron Lopez, Louise Marcotte, and André Cormier for all of your help at the Brain Imaging Centre of the MNI. Thank you to Annie Le Bire, Hélène Drouin, and Jacqueline Bednar for your administrative assistance, I am grateful for all your hard work and easing the process that is graduate school.

I would like to acknowledge the funding and support provided by the Canada Foundation for Innovation (CFI), the Centre for Interdisciplinary Research in Music Media and Technology, and the McGill–Music Travel Award.

Thank you to all of my friends. To Christopher J. Steele, I am grateful for your fellowship and the long discussions, and I hope to get that circuit board in a box for you soon. Thanks to Joseph Thibodeau for your collaboration, patience with my ramblings, and for the well-documented meetings taken in exquisite script. I owe a debt of gratitude to Joseph Malloch, Mark Zadel, and Stephen Sinclair for your gracious

assistance, shared enthusiasm, and making my research life as efficient as possible while also providing me with much needed diversion. Thanks to everyone else at the IDMIL and Music Tech. for your insight and ideas, especially Bertrand Scherrer for helping me practice my french and translating my abstract. Thanks to Sean D. S. for innumerable discussions over lunch and coffee, both technical and in respite. Thanks to Matt F. to solving their months-long issues with your weekend of coding—I think they're still making good use of it. Thanks to Bartek for help with the FHN models. Thanks to Marcin for sharing with us your music and sweet times.

Thank you to my family, close and extended. To my parents for your love, support, shelter, and opportunity. For being there when I needed it and pushing me toward wonderment and inquiry. To my uncle Allan for your technical knowledge and taking us out to dinner. To my uncle Sandy for inducing me to focus on a clear goal and on finishing. To my aunt Bevy for your interest, encouragement, help with typesetting, and proofreading. Thank you to my brother and sister-in-law, and most of all, my nephews, for asking discerning questions and forcing me to answer in plain terms. My father-in-law and mother-in-law, I thank for your kind words, shared meals, and wonderful daughter. To my brother-in-law, thank you for your camaraderie and intrigue, and sorry for keeping your sister away from home—you're welcome to visit any time. And thank you to my lovely wife, for the joy and elevation that you bring to my life; for taking me to places close and afar; and for just getting me out of the house. I'll get back to doing the dishes soon. Thank you as well for your help moulding the Ballagumi, laying up fiberglass for the cello, and drawing figure 6.1. Mwwaaa!

Thank you.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Previous Work | 2 |
| 1.2 | Design and Implementation | 4 |
| 1.3 | Contributions | 5 |
| 1.4 | Overview | 5 |
| 2 | Background | 7 |
| 2.1 | Neuroimaging, Interfaces, and Music | 9 |
| 2.1.1 | Neuroimaging Techniques | 9 |
| 2.1.2 | Feedback Manipulation in Auditory-Motor Integration Tasks | 11 |
| 2.1.3 | Probing the Auditory-Motor System with Electronic Musical Instruments | 14 |
| 2.1.4 | Musical Tasks | 16 |
| 2.1.5 | Non-Musical MRI-Compatible Interfaces | 21 |
| 2.1.6 | Discussion | 30 |
| 2.2 | Optical Sensing | 32 |
| 2.2.1 | An Introduction to Fibre Optic Sensors | 32 |
| 2.2.2 | Interferometric Sensors | 34 |
| 2.2.3 | Polarization Sensors | 37 |
| 2.2.4 | Grating-Based Sensors | 38 |
| 2.2.5 | Scattering and Luminescence-Based Sensors | 41 |
| 2.2.6 | Discussion | 43 |
| 3 | Embedded Systems Design | 45 |
| 3.1 | Preliminary Efforts | 48 |

| | | |
|----------|--|-----------|
| 3.2 | Hub and Node Hardware Platform | 49 |
| 3.2.1 | Analog Front-End | 50 |
| 3.2.2 | Mapping | 51 |
| 3.2.3 | Synthesis | 52 |
| 3.2.4 | Discussion | 52 |
| 3.3 | Fungible-1 | 53 |
| 3.3.1 | Optoelectronic Acquisition Board | 53 |
| 3.3.2 | Hub Board | 59 |
| 3.3.3 | Discussion and Applications | 60 |
| 3.4 | Fungible-2 | 62 |
| 3.4.1 | Optoelectronic Acquisition Board Overview | 63 |
| 3.4.2 | Power Regulation and Distribution | 63 |
| 3.4.3 | Optical Subsystem | 66 |
| 3.4.4 | Internal PSoC Hardware and Firmware | 67 |
| 3.4.5 | SPI Communication and MR-COSBBE | 74 |
| 3.4.6 | Discussion | 79 |
| 3.5 | Physical Modelling Synthesis | 80 |
| 3.5.1 | FitzHugh-Nagumo Excitable Cell | 81 |
| 3.5.2 | Ishizaka-Flanagan Vocal Fold Model | 83 |
| 3.5.3 | Real-Time Synthesis | 84 |
| 3.5.4 | Discussion | 87 |
| 3.6 | Conclusion | 87 |
| 4 | Ballagumi | 89 |
| 4.1 | Interface Design | 90 |
| 4.1.1 | Design Concept | 90 |
| 4.1.2 | Optical sensor design | 90 |
| 4.1.3 | Initial Prototype | 91 |
| 4.1.4 | Sculpture, Mould-Making, and Casting | 93 |
| 4.2 | Signal Acquisition, Mapping, and Synthesis | 97 |
| 4.2.1 | Qualitative Study | 97 |
| 4.2.2 | Gesture Analysis | 98 |
| 4.3 | Discussion | 105 |

| | | |
|----------|---|------------|
| 5 | fMRI-Compatible Keyboard | 109 |
| 5.1 | First Generation Keyboard and Limitations | 109 |
| 5.2 | First Generation Keyboard with Upgraded Electronics | 112 |
| 5.2.1 | Calibration and Testing | 114 |
| 5.3 | Second Generation Keyboard | 115 |
| 5.3.1 | Mechanical Keyboard | 116 |
| 5.4 | Sensing Methods | 120 |
| 5.4.1 | Testing of Sensing Methods | 123 |
| 5.5 | MRI Testing | 133 |
| 5.5.1 | First Generation Keyboard | 133 |
| 5.5.2 | Second Generation Keyboard | 134 |
| 5.6 | Discussion and Future Improvements | 142 |
| 5.7 | Conclusion | 143 |
| 6 | Optoacoustic Cello | 145 |
| 6.1 | Why a Cello? | 145 |
| 6.2 | Design Goals | 146 |
| 6.3 | Implementation | 148 |
| 6.3.1 | Parts of a Cello | 148 |
| 6.3.2 | Modification or Fabrication? | 148 |
| 6.3.3 | Cello Body | 150 |
| 6.3.4 | Fingerboard | 157 |
| 6.3.5 | Bridge | 170 |
| 6.3.6 | Bow | 176 |
| 6.3.7 | MRI-Compatible Cello Support Fixture | 189 |
| 6.4 | Testing and Results | 192 |
| 6.4.1 | Fingerboard Sensors | 192 |
| 6.4.2 | Bridge Pickup | 193 |
| 6.4.3 | Bow Sensors | 196 |
| 6.4.4 | MRI-Compatibility | 207 |
| 6.5 | Discussion | 211 |
| 6.5.1 | Cello Body | 213 |
| 6.5.2 | Fingerboard | 214 |

| | | |
|----------|--|------------|
| 6.5.3 | Bridge | 215 |
| 6.5.4 | Bow | 215 |
| 6.6 | Conclusion | 219 |
| 7 | Closing Remarks | 221 |
| 7.1 | Limitations and Future Work | 222 |
| 7.2 | Musical Applications | 224 |
| 7.3 | Applications in Neuroscience | 225 |
| 7.4 | Applications in Neuro-Rehabilitation | 226 |
| | Appendices | 231 |
| A | Positioning of optical fibres within urethane cello fingerboard | 231 |

List of Figures

| | | |
|------|---|----|
| 3.1 | Block diagram of optoelectronic acquisition system. | 46 |
| 3.2 | Acquisition electronics for first-generation MRI-compatible keyboard. . . | 49 |
| 3.3 | Hub and node hardware platform block diagram for optically-sensed interfaces [117]. | 51 |
| 3.4 | Analog front end, showing an example configuration utilizing four optical channels: three differential and four single-ended [114]. | 55 |
| 3.5 | PCB layout of <i>fungible-1</i> optoelectronic acquisition board. | 57 |
| 3.6 | Optoelectronic acquisition and control board with fibre optic connectors. . . | 58 |
| 3.7 | Optoelectronic acquisition board connected to fibre optic interface, hub board, and PC [114]. | 60 |
| 3.8 | <i>Fungible-2</i> optoelectronic acquisition board PCB layout. | 65 |
| 3.9 | <i>Fungible-2</i> optoelectronic acquisition board. | 66 |
| 3.10 | Normalized ADC and photodiode signals with LED voltage control. . . . | 68 |
| 3.11 | Sample layout of <i>fungible-2</i> on-chip analog hardware configuration. | 70 |
| 3.12 | <i>Fungible-2</i> sample embedded digital hardware. | 71 |
| 3.13 | MR-COSBBE internals with several <i>fungible-2</i> optoelectronic acquisition boards mounted together. | 75 |
| 3.14 | MR-COSBBE and other equipment set up in the MRI control room. | 75 |
| 3.15 | Waveform of FHN model in astable multivibrator state. | 82 |
| 3.16 | Glottal waveform produced with the IF model. | 84 |
| 4.1 | Initial sketch of the interface, with wings and spines. | 91 |
| 4.2 | Differential fibre optic deformation sensor. | 92 |
| 4.3 | Fibre optic multi-touch force and bend sensor grid. | 92 |
| 4.4 | Initial GGT interface prototype and Forton mould. | 94 |

| | | |
|------|--|-----|
| 4.5 | Plasticine model of the Ballagumi. | 95 |
| 4.6 | Sensors placed in mould after initial silicone base layer set. | 96 |
| 4.7 | The Ballagumi with sensors potted before final silicone pour. | 96 |
| 4.8 | The Ballagumi, silicone interface with embedded fibre optic sensors. | 97 |
| 4.9 | Glowing Ballagumi, showing sensor placement. | 99 |
| 4.10 | Sensor data for repeated “Apex Push” gesture. | 101 |
| 4.11 | The first three principal components of sensor data for “Apex Push” gesture. | 102 |
| 4.12 | Supervised clustering of gestures and their projection on a three-dimensional space | 103 |
| 4.13 | Unsupervised clustering of gestures and their projection on a three- dimensional space. | 104 |
| 4.14 | Block diagram of an adaptive system capable of mapping gestures onto sound parameters. | 107 |
| 5.1 | First generation piano prototype, presented in [113]. | 111 |
| 5.2 | First piano prototype with mirrors on backs of keys aligned with emit- ting and receiving fibres [113]. | 112 |
| 5.3 | Characterization curve of 1st-generation keyboard, key position in tran- sition region. | 115 |
| 5.4 | Yamaha P85 electronic weighted-key piano. | 116 |
| 5.5 | Yamaha electronic piano internals. | 117 |
| 5.6 | CAD drawing of optical fibre mount. | 118 |
| 5.7 | Annotated photograph of second keyboard prototype with optical fibres and mirror. | 119 |
| 5.8 | Photomasks to measure piano key displacement. | 122 |
| 5.9 | Optical quadrature sensing during key transition. | 124 |
| 5.10 | Triangular differential photomask sensing of a key. | 125 |
| 5.11 | Key sensing with continuously-variable gradient photomask. | 126 |
| 5.12 | Repeated key presses of a white key. | 127 |
| 5.13 | Repeated key presses of a black key. | 128 |
| 5.14 | Scatter plot of white key. | 129 |
| 5.15 | Scatter plot of black key. | 130 |

| | | |
|------|--|-----|
| 5.16 | Single key press of a black key. | 131 |
| 5.17 | Scatter plot of linearized black key. | 132 |
| 5.18 | Gradient field map of MRI phantom and 1st-generation keyboard. | 134 |
| 5.19 | MRI-compatible keyboard in MRI scanner. | 136 |
| 5.20 | Two keys played in succession using continuous gradient photomasks. | 137 |
| 5.21 | Gradient magnetic field maps (magnitude) of MRI phantom with 2nd-generation piano. | 138 |
| 5.22 | Gradient magnetic field maps (phase) of MRI phantom with 2nd-generation piano. | 139 |
| 5.23 | Gradient field map of subject with 2nd-generation keyboard at arm's length from isocentre (magnitude images). | 140 |
| 5.24 | Gradient field map of subject with 2nd-generation keyboard at arm's length from isocentre (phase images). | 141 |
| 6.1 | Annotated diagram of cello body. | 149 |
| 6.2 | Early assembly drawing of MRI-compatible cello interface. | 152 |
| 6.3 | CAD model of MRI-compatible cello body. | 153 |
| 6.4 | CAD model of MRI-compatible cello body mould. | 154 |
| 6.5 | Aluminium cello body mould during initial stages of layup. | 156 |
| 6.6 | Paper template preparation for cello body. | 157 |
| 6.7 | Cello body simulation of relative stress distribution with views of top and underside, showing higher-stress area at the heel of the neck. | 158 |
| 6.8 | Composite cello layup with first layers of fiberglass prepreg and balsa wood core. | 159 |
| 6.9 | Composite cello layup with second set of fiberglass prepreg. | 160 |
| 6.10 | Composite cello in vacuum bag and conformal protective envelope, ready for heat-curing. | 161 |
| 6.11 | First prototype of composite fiberglass cello body, with gut strings and transparent fingerboard embedded with optical fibres. | 164 |
| 6.12 | Fingerboard test fixtures with embedded optical fibres: composite fiberglass touch sensor). | 164 |
| 6.13 | Repeated touches on fiberglass touch sensor. | 165 |
| 6.14 | 3D CAD model of cello fingerboard mould. | 165 |

| | |
|--|-----|
| 6.15 Cello fingerboards (wood, fiberglass, and cast urethane). | 166 |
| 6.16 Transparent urethane test piece cast with 0.5 <i>mm</i> POF. | 167 |
| 6.17 Repeated touches on cast transparent urethane proximity sensor, measured with 0.5 <i>mm</i> POF. | 168 |
| 6.18 Transparent urethane fingerboard cast with POF mounted on cello body. | 169 |
| 6.19 Truncated portion of CAD model of 3D-printable cello fingerboard. | 171 |
| 6.20 Truncated 3D-printed fingerboard prototypes. | 172 |
| 6.21 Transparent 3D-printed fingerboard mounted on cello body with gut strings. | 173 |
| 6.22 Wooden wedge-style cello bridge on acrylic bridge plate, and peg box. | 173 |
| 6.23 CAD model of contact sensorized bridge. | 175 |
| 6.24 CAD model of non-contact sensorized bridge. | 177 |
| 6.25 Non-contact sensorized bridge with sensor head on D-string. | 177 |
| 6.26 Annotated diagram of cello bow. | 178 |
| 6.27 30 <i>cm</i> long MRI-compatible cello bow. | 178 |
| 6.28 CAD model of the bottom half of the silicone force sensor mould. | 180 |
| 6.29 Fibre optic force sensor. | 180 |
| 6.30 MRI-compatible cello bow with bow force sensors mounted at tip and frog, and 3D magnetic field sensor. | 181 |
| 6.31 CAD assembly drawing of optical flow sensor. | 183 |
| 6.32 CAD model (underside) of fixture used to secure and align POF with the optical flow sensor. | 183 |
| 6.33 Photograph of cello bowing trajectory. | 184 |
| 6.34 Faraday rotator schematic. | 186 |
| 6.35 Prototype Faraday rotator in aluminium lens tube. | 187 |
| 6.36 Proof-of-concept Faraday rotator moved past a niobium rare earth magnet with a bowing-like gesture (500 <i>Hz</i> sampling rate). | 187 |
| 6.37 Bowing-like gestures performed in 3 <i>T</i> MRI scanner measured with proof-of-concept Faraday rotator in aluminium lens tube (500 <i>Hz</i> sampling rate). | 188 |
| 6.38 Magnetic field sensor with optical elements and 3D-printed housing. | 188 |
| 6.39 CAD assembly drawing of 3D magnetic field sensor. | 189 |
| 6.40 CAD assembly drawing of cello support fixture with cello body in scanner. | 190 |
| 6.41 Material failure of 3D-printed ball joint. | 191 |

| | | |
|------|--|-----|
| 6.42 | Cello supported in MRI-compatible fixture. | 191 |
| 6.43 | Repeated finger-string-fingerboard contact with cast transparent urethane fingerboard. | 193 |
| 6.44 | Repeated finger-to-fingerboard contact with 3D-printed fingerboard. | 194 |
| 6.45 | Finger-string-fingerboard interaction, showing a shift from D^\sharp to F^\sharp on the D-string, as sensed with the 3D-printed fingerboard [239]. | 195 |
| 6.46 | Contact optical bridge pickup signal. | 196 |
| 6.47 | Non-contact optical bridge pickup signal. | 197 |
| 6.48 | Synchronous trace of a cellist performing a repeated bowing gesture in MRI scanner. | 198 |
| 6.49 | Characterization of toroidal force sensor. | 199 |
| 6.50 | Characterization of toroidal force sensor nonlinearities. | 200 |
| 6.51 | Bow-string interaction during down-bow. | 202 |
| 6.52 | Accuracy of bow position estimate using force sensors. | 203 |
| 6.53 | Linearity of bow position estimate using force sensors. | 204 |
| 6.54 | Bowing gesture captured with optical flow sensor. | 205 |
| 6.55 | MRI-compatible cello and cellist in MRI scanner. | 206 |
| 6.56 | Measurement of the MRI scanner's static magnetic field (Z-axis). | 208 |
| 6.57 | Gradient magnetic field maps (magnitude) of MRI phantom with cello and bow. | 209 |
| 6.58 | Gradient magnetic field maps (phase) of MRI phantom with cello and bow. | 210 |
| 6.59 | MRI field map of cellist with MRI-compatible cello (magnitude image). No image artifacts are visible. | 211 |
| 6.60 | Pitch shift on the D-string, captured with the non-contact sensorized bridge. | 212 |

List of Tables

| | | |
|-----|--|-----|
| 3.1 | <i>Fungible-1</i> Optoelectronic control board features. | 54 |
| 3.2 | <i>Fungible-2</i> Optoelectronic control board features. | 64 |
| 5.1 | MRI compatibility of first generation piano keyboard. | 134 |
| 5.2 | MRI compatibility of second generation piano keyboard. | 136 |
| 6.1 | MRI compatibility of optoacoustic cello. | 210 |

List of Acronyms

| | |
|--------|--|
| AC | Alternating Current |
| ADC | Analog-to-Digital Converter |
| ARM | A 32-bit reduced instruction set computing architecture |
| ASCII | American Standard Code for Information Interchange |
| BCI | Brain-Computer Interface |
| CAD | Computer-Aided Design |
| CLK | Clock signal |
| COM | Communication port, usually compliant with the RS-232 standard |
| CS | Chip Select signal |
| DAC | Digital-to-Analog Converter |
| DC | Direct Current |
| DMA | Direct Memory Access |
| DMI | Digital Musical Instrument |
| DOF | Degree-of-Freedom |
| DSP | Digital Singal Processor |
| EEG | Electroencephalography |
| EEPROM | Electrically Erasable Programmable Read-Only Memory |
| EMI | Electromagnetic Interference |
| ERF | Electrorheological Fluids |
| ERP | Event-Related Potential |
| FAT | File Allocation Table, computer file system architecture |
| FBG | Fibre Bragg Grating |
| FHN | FitzHugh-Nagumo, excitable cell model |
| FM | Frequency Modulation |

| | |
|-----------|---|
| fMRI | Functional Magnetic Resonance Imaging |
| fNIRS | Functional Near-Infrared Spectroscopy |
| FPI | Fabry-Pérot Interferometer |
| GGT | Goey Glob Thing, prototype interface |
| GPIO | General Purpose Input/Output, IC pin |
| HDMI | High-Definition Multimedia Interface |
| HID | Human Input Device |
| I2C | Inter-Integrated Circuit, communication protocol |
| IC | Integrated Circuit |
| IDE | Integrated Development Environment |
| IDMIL | Input Devices and Music Interaction Laboratory |
| IF | Ishizaka and Flanagan, vocal fold model |
| ISR | Interrupt Service Routine |
| LED | Light-Emitting Diode |
| LPG | Long Period Grating |
| LTC | Linear Time Code |
| M8C | 8-bit microcontroller core, based on the Harvard architecture |
| MEG | Magnetoencephalography |
| MEMS | Microelectromechanical System |
| MFST | Multiple Finger Sequence Task |
| MIDI | Musical Instrument Digital Interface, communication protocol |
| MISO | Master-In/Slave-Out signal |
| MNI | Montreal Neurological Institute |
| MOSFET | Metal-Oxide-Semiconductor Field-Effect Transistor |
| MOSI | Master-Out/Slave-In signal |
| MR | Magnetic Resonance |
| MR-COSBBE | Magnetic Resonance-Compatible Optical System-in-a-Box for Behavioural Experiments |
| MRI | Magnetic Resonance Imaging |
| NDK | Native Development Kit |
| NVM | Non-Volatile Memory |
| OS | Operating System |
| OTDR | Optical Time-Domain Reflectometry |
| P4A | Python for Android |
| PC | Personal Computer |

| | |
|--------|---|
| PCA | Principal components analysis |
| PCB | Printed Circuit Board |
| PET | Positron Emission Tomography |
| PGA | Programmable-Gain Amplifier |
| POF | Plastic Optical Fibre |
| POTDR | Polarization-Optical Time-Domain Reflectometry |
| PSoC | Programmable System-on-Chip |
| PWM | Pulse Width Modulation |
| RF | Radio Frequency |
| RS | Reed-Solomon, forward error-correcting codes |
| SAR | Successive Approximation Register, ADC architecture |
| SD | Secure Digital card, non-volatile digital media |
| SMPTE | Society of Motion Picture and Television Engineers, standard time code |
| S/PDIF | Sony/Philips Digital Interface Format, digital audio communication protocol |
| SNR | Signal-to-Noise Ratio |
| SPI | Serial Peripheral Interface, communication protocol |
| ST | Straight-Tip, style of fibre optic connector |
| TDCS | Transcranial Direct Current Stimulation |
| TGG | Terbium Gallium Garnet |
| TIA | Transimpedance Amplifier |
| TMS | Transcranial Magnetic Stimulation |
| UART | Universal Asynchronous Receiver/Transmitter, communication protocol |
| UNF | Unité de Neuroimagerie Fonctionnelle, Université de Montréal |
| USB | Universal Serial Bus, communication protocol |
| UV | Ultraviolet light |
| VBUS | USB power |
| VR | Virtual Reality |

The operating mechanism can even be thrown into action independently of any object to operate upon (although of course no *result* could then be developed). Again, it might act upon other things besides *number*, were objects found whose mutual fundamental relations could be expressed by those of the abstract science of operations, and which should be also susceptible of adaptations to the action of the operating notation and mechanism of the engine. Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expressions and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent.

—*The Right Honourable Ada Augusta, Countess of Lovelace*

Translator's Notes, Note A,
from, *Sketch of the Analytical Engine Invented by Charles Babbage*,
by Count L. F. Menabrea
Initially published in English in *Scientific Memoirs*, 1843

Chapter 1

Introduction

The neuroscience of music performance is increasingly making use of neuroimaging techniques to correlate localized neural activity with perceptual and behavioural tasks [10, 63, 295]. Ideally, a researcher would like to ask a musician, such as a cellist, to perform a musical excerpt on his or her cello while collecting performance and neuroimaging data using functional magnetic resonance imaging (functional MRI or fMRI). Timing of note onset and its regularity, as well as quantifying fingering mistakes and pitch accuracy are of keen interest and need to be acquired synchronously with brain imaging data. Differences in bowing and left-hand performance between amateur and professional musicians, and the transition from one to the other through practice, can be tracked and correlated with changes in both brain structure and activity [119, 120]. With a better understanding of the differential neural activity elicited by subtle, tightly controlled differences in auditory-motor integration tasks—the musician’s specialty—neuroscientists can develop more refined models of how the musical brain works. Additionally, many interesting possibilities are made available to the researcher if the auditory feedback provided to the musician can be removed, pitch-shifted, or delayed. It is clear that an acoustic instrument, and even a commercially-available electric emulation (e.g., an electric cello or an electronic piano keyboard), is insufficient for performance within the environment of the MRI scanner, and compatible and specially-augmented interfaces are required for neuroimaging research.

In order to undertake these fMRI studies, musicians perform on a specialized musical instrument while lying supine in an MRI scanner: a 60 *cm* diameter tube, wherein

there is little room to move and environmental noise during image acquisition can be in excess of 120 dB [73]. There are many additional design constraints that impact the compatibility of devices brought into the scanner that need to be addressed: a constant 3 Tesla (T) magnetic field—over 50000 times that of the earth’s—prohibits any ferromagnetic materials due to safety concerns; pulsed magnetic fields and radio frequency emissions cause electronics to misbehave, heat up, and possibly combust; conductive materials and electronics within the scanner can cause image artifacts; and with the MRI receiver coils so sensitive that they can detect changes in blood oxygenation within the brain—a correlate of brain activity—subject movement within the scanner blurs images.

Thus the goal of this work was to overcome the engineering problems associated with the design of musical interfaces that can be played within the MRI scanner while ensuring safety and compatibility and at the same time acquiring performance gestures electronically. Emulations of a piano keyboard and a cello were sought. Performance gestures specific to the piano included: which key is played and when, with a secondary goal of measuring the positional trajectory of each key as it is played. Performance gestures specific to cello performance included: left-hand finger-string-fingerboard contact position, string vibration, and right-hand bow speed and bow-string pressure. Keeping with previous neuroimaging experiments, a relative measure of these performance metrics, which can be tracked and compared across trials and experimental parameters, was sought without an explicit and absolute reference frame.

To overcome these challenges, we have designed and built prototypes of specialized optoelectronic interfaces which emulate acoustic musical instruments that can be played by musicians within the constraints and confines imposed by current neuroimaging techniques. The construction of these musical instruments entailed the design of fibre optic sensors and microcomputer systems, as well as computer-aided mechanical design.

1.1 Previous Work

The study of musicians’ brains during performance using fMRI has evolved from imagining musical performance [145], to asking the musician to pretend to play by moving their fingers in the absence of an instrument [161], to playing a mute electronic

piano keyboard devoid of any metal or sensors [177], to playing a mute electronic piano that directly tracks note onset and velocity optoelectronically [10], to a similar keyboard that has been used to produce real-time auditory feedback [116]. Where direct acquisition of musical performance was not available, researchers have used surface electromyography recordings [161] and expert analysis of in-scanner video to coarsely assess performance [177].

Various sensors and interfaces for use in the MRI environment have been developed [34, 113, 263], though the use of electronic sensing and signalling within the MRI environment requires significant testing to ensure compatibility [153, 230], adding to development time and cost. With researchers moving to higher-field scanners, there is no guarantee electronic equipment that was compatible with a 1.5 *T* scanner will be compatible with a 3 *T* or 7 *T* scanner due to the scanner's increased electromagnetic emissions and the need for increased field homogeneity within the scanner. Additionally, magnetoencephalography (MEG), another functional neuroimaging technique with finer temporal resolution, is also highly sensitive to similar material and electronic constraints as fMRI [19, 257].

The use of fibre optic sensing, as opposed to electronic sensing, does not affect image quality at any field intensity and has been applied to MRI-compatible interfaces, but only rarely in musical tasks [10, 17, 34, 116, 261, 262]. The neuroscience of music must expand its use of these techniques in order to progress. Fibre optic sensing in its simplest form requires an emitter, a detector, and optical fibre to convey light between the two. For our purposes, the emitter and detector, which are optoelectronic components, can be safely placed outside the scanner environment where the incoming light is converted into an electrical signal, amplified, and digitized to allow the measured performance gesture to be acquired in synchrony with the neuroimaging data. The fibre optic cable may be engineered to have some intrinsic sensitivity to the measured phenomenon (e.g., temperature or strain), or a beam of light, launched through a fibre optic cable at a sensitive extrinsic target, undergoes some measurable modulation which is conveyed to the detector with a second fibre optic cable. Intensity or amplitude modulation is the simplest to detect, and often other properties of light, such as polarization, phase, and wavelength, are converted to a measure of amplitude at the detector [275].

1.2 Design and Implementation

We have designed and constructed MRI-compatible musical instrument prototypes from which novel and interesting performance data may be obtained, and without which neuroimaging studies of musicians *playing* their instrument could not be conducted. Three instrument prototypes act as case studies for the development of the underlying technology: the Ballagumi, a novel interface for musical expression; an electronic piano keyboard; and an electro-acoustic cello.

Each instrument may be thought of as consisting of three distinct subsystems: the tangible human-computer interface, the sensors, and the acquisition and control electronics. When emulating an acoustic instrument, the physical interface must feel and react to the performer's touch as if it were an acoustic instrument. As much as possible within the constraints of the MRI environment, a realistic interface is essential. This may take the form of weighted keys of a piano, or a cello with a string length and neck dimensions that match its acoustic kin. The Ballagumi, a flexible silicone interface embedded with optical sensors, allows a completely novel musical experience as the musician has no preconceived notions as to its functioning or auditory feedback. The interface can be struck, bent, depressed, and generally deformed. Using gesture recognition techniques, nuanced two-handed musical performance gestures drive a synthesizer through a programmable mapping. The ability to manipulate the mapping from gestures onto sound allows an experimenter to investigate on-line auditory-motor integration, and allows a composer or performer to fine-tune their respective piece or performance.

We have designed various sensing elements by exploring the many configurations of inexpensive plastic optical fibres. Proximity, position, displacement, flexion, force, vibration, and magnetic field can all be transduced and measured as optical intensity using fibre optics with and without additional optical elements. We sensed flexion and deformation of a flexible interface using bending losses in a single fibre and transmission losses between pairs of fibres. We have designed sensors that measure not only which piano key is depressed and when, but the positional trajectory the key takes as it is depressed, through the use of photomasks. And we have measured cellists' left-hand fingering by embedding fibre optic proximity sensors within the cello fingerboard, and right-hand bowing gestures through the design of optical sensors that quantify bow-

hair displacement and the changes in magnetic field as the bow is moved within the MRI scanner.

Mixing traditional instrument building techniques and commercial products with computer-aided modelling and machining, as well as moulding and 3D printing, allows for a substrate which provides a more natural interface for musicians while enabling iterative design—the process by which a prototype is designed, built, tested, and refined in a recursive manner. Furthermore, embedding sensors within the instrument and making use of microcomputers, enables the direct measurements of musical performance not available with traditional acoustic (and even electric) instruments.

1.3 Contributions

The primary contributions are: the design and deployment of these unique instruments and technological tools that are directly enabling neuroscientists to further our basic understanding of how the brain functions and changes with music performance, and the development of novel sensors and embedded systems for new musical interfaces. These interfaces are being used in a variety of fMRI experiments, with MEG studies also being considered, which until now were impossible to conduct. Secondary contributions include the exploration of a variety of optical sensors and material processes, as well as field-tested modular embedded systems that can be applied to new interfaces, be they for musical expression, neuroimaging studies, or more general human-computer interaction. The applications of this research extend beyond the study of music and the brain to more basic motor control and motor-learning research and to clinical applications in neuro-rehabilitation.

1.4 Overview

The body of this thesis is organized into chapters: Chapter 2 briefly reviews the imaging techniques utilized by behavioural neuroscientists, discusses some of the literature surrounding auditory-motor integration tasks and the use of electronic musical instruments, and is followed by an examination of musical and non-musical behavioural neuroimaging experiments and their associated MRI-compatible interfaces. The second half of this chapter introduces fibre optic sensors, organized by sensing technique,

and suggests potential applications for their use in MRI-compatible musical interfaces. Chapter 3 discusses the rationale for building a modular, stand-alone, optoelectronic acquisition system and its use in behavioural experiments and musical instruments. The incremental design of the embedded systems is presented, encompassing the engineering of optoelectronic, analog, and digital circuitry, and the functionality of associated software and embedded firmware. Of note are the embedded gesture classification, mapping, and physical modelling synthesis algorithms. Making use of these algorithms is the Ballagumi, presented in Chapter 4, and a description is given including its conception, fabrication, and testing as a musical instrument. Chapter 5 gives a chronology of the MRI-compatible keyboard, its design, testing, and use in fMRI studies. Chapter 6 delves into the design of an MRI-compatible cello as a useful instrument to study on-line and continuous manipulation of auditory feedback and the cellists' resultant motor corrections. The design and construction of the optoacoustic cello, bow, and respective fibre optic sensors are presented, and their proper functioning within the MRI scanner is demonstrated. And lastly, Chapter 7 remarks on the utility of these interfaces, sensors, and electronics for future neuroimaging studies, musical expression, and other applications beyond.

Chapter 2

Background

In this chapter, we highlight two disparate, though germane, areas of research: functional neuroimaging and optical sensing. More so, this chapter serves as the basis for how the latter can advance the former; phrasing the questions, what are neuroscientists trying to accomplish? What roles do musicians and music play in this venture? How does one go about studying musicians, what tools are required, and what are the difficulties? How do optical sensors help to overcome these challenges? What optical sensing techniques are available and how can they be applied to musical interfaces? Instead of delaying the answers to these questions, perhaps a brief reply is fitting:

- Neuroscientists are not only interested in locating brain regions and identifying their involvement in cognitive processes. They are also trying to unravel the underlying mechanisms and network of connections within the brain in order to build models and test theories of the causal nature of our brain as an adaptive system with many inputs and outputs.
- Musicians are highly-trained and highly-skilled individuals with a keen sense of perception of what they hear and of how movements of their bodies (and their instruments as extensions of their bodies) affect the sounds they produce. They embody the finely-tuned adaptive dynamical system that integrates multi-modal feedback to evoke an emotional response through a motor output—something a robot has yet to do, at least convincingly and with any sort of prescience in anticipating fellow musicians within an ensemble.

- In addition to determining if an individual can perceive certain stimuli (in the presence or absence of other stimuli) through music listening, the study of music performance itself may provide a more direct insight into a musician's special abilities. Technology has enabled us to delve a little deeper into the working mind through non-invasive measurements of cognitive processes. A number of these techniques, including fMRI, can help to identify regions of neural activity in response to actions and percepts, though not all techniques can be applied in a naturalistic way to music as it is performed on stage.
- Instead, a musician must play their instrument—or an emulation thereof, if their instrument is prohibited from the experimental environment on the basis of safety, space, or lack of augmented functionality—in a setting that is foreign and in a way that is constrained in order to extricate the causal and differential neuronal activation from their respective behaviours and contexts. Electronic acquisition of performance gestures, in order to measure subtle changes in timing and accuracy, or for the purpose of generating arbitrary feedback (either conforming to or breaking the musician's expectations), is essential to probing behaviour and correlating it with neural activity. Electronic acquisition of behaviour through the use of physical interfaces and sensors in the imaging scanner is mired in complications of electromagnetic interference (EMI), space constraints, and the simple difficulty of measuring phenomena without interfering with the measurement at hand.
- Light has the unique capacity to be confined, conveyed, and contorted in ways that both mitigate the difficulties associated with fMRI and allow unfettered access to measures of vibration, displacement, and force through contact and non-contact transduction. Fibre optic sensors embedded in musical interfaces can provide gestural information through the modulation of light intensity, phase, wavelength, and polarization. Gestures can be analyzed and correlated with functional imaging data, or they can be used to generate feedback.

2.1 Neuroimaging, Interfaces, and Music

Music performance involves coordinating motor activity while integrating auditory, visual, haptic, and proprioceptive feedback. The underlying mechanisms implicated in this complex task are still poorly understood. Studies that investigate the neural activity of musicians and non-musicians as they perform musical tasks, such as audition and pitch discrimination or performance and motor learning, help neuroscientists better understand the function of specific brain areas and their connectivity.

Indeed, the neural activity needs to be correlated with measured behavioural changes if one is to link, for instance, motor learning with activation of specific neural centres. The most accurate way to compare motor learning with neuroimages is to acquire synchronized performance metrics electronically, thereby obviating the need for adjudication to be done by sight or other methods—a difficult process only exacerbated within the imaging environment [61, 161, 177]. Thus, a digital musical instrument (DMI) provides the means through which to acquire performance information electronically for off-line analysis, while at the same time providing real-time auditory feedback.

This section briefly discusses various neuroimaging techniques, provides a review of feedback manipulation in auditory-motor integration tasks, argues for the role of electronic musical instruments in neuroscience, presents a review of musical and non-musical tasks and interfaces employed within the MRI environment, and suggests the use of DMIs in neuro-rehabilitation and music-supported therapy.

2.1.1 Neuroimaging Techniques

There are several techniques available to probe the internal workings of the brain during behavioural tasks. None of them provide direct and omnipotent access to measures of interest, but each has its uses, advantages, and disadvantages. Here, we briefly describe some of the techniques available to neuroscientists to probe the working brain—that is to say the brain in action, separate from the structural images that a clinician uses to locate and diagnose a lesion.

The most direct measurement, and also the most invasive, is direct cellular recording. The pioneering work of Jasper, Penfield, and Feindel at the Montreal Neurological Institute (MNI) to investigate the functional anatomy and neural mechanisms of epilepsy made use of microelectrodes to measure the electrical signals of single brain

cells [68]. Electrocorticography has also been applied more recently in brain-computer interfaces (BCI) using microelectrode arrays implanted onto the surface of the cerebral cortex [60, 149, 162].

Non-invasive techniques, however, are necessary for the general study of individuals, and as such, electroencephalography (EEG) using surface electrodes to measure the electrical activity of the brain is a preferred method [9, 168, 226]. The impedances of the tissues and skull, and the ensemble of many neurons acting in unison (and/or independently and out of phase), cause a smearing of action potentials, making localization difficult. The high temporal resolution allows event-related potentials (ERP) to be measured as small disturbances in the time course of peaks in activity. Here, imaging techniques can be combined to obtain highly-localizable activity with temporal acuity [85, 266]. Along with EEG, transcranial magnetic stimulation (TMS) and transcranial direct current stimulation (TDCS) have been employed to temporarily knock-out or enhance performance of certain brain areas, with resulting changes in behaviour or evoked potentials [7, 29, 223].

Positron emission tomography (PET), which makes use of radioisotopes whose radioactive decay is measured and localized by the scanner, has been used in a number of neuroimaging studies to measure task-dependent changes in localized cerebral blood flow [201, 294]. The lack of an endogenous contrast agent, the administration of a radioisotope, and the use of MRI co-registration, have seen some types of PET studies superseded by other imaging methods.

A prominent technique available to neuroscientists is fMRI, which provides 3D images of the brain, enabling the correlation of action and planning with localized neural activity. The functional part of fMRI differentiates it from the anatomical MRI used clinically in medical diagnosis (although structural scans are currently being used to correlate changes in the size of certain brain areas with behavioural changes and the emergence of expert skill [119]), and requires many trials to build up a statistically significant response to a specific cognitive task, such as listening to music and tapping along to a beat.

In fMRI, changes in blood oxygenation, correlated with localized neuronal activity, is measured using a large magnetic and radio-frequency (RF) fields. Although the MRI scanner does not use ionizing radiation, the endogenous contrast, due to the differential magnetic permeabilities of deoxyhemoglobin and oxyhemoglobin, is delayed

“on the order of seconds” [26] with respect to the actual neuronal activity. The localized uptake of oxygenated blood follows a hemodynamic response, allowing for sparse sampling [79]: interleaving task and image acquisition, to reduce motion artifacts and the effects of extraneous auditory stimuli generated by the scanner. The MRI scanner environment, with its large magnetic field and RF radiation, poses a potential danger to participants, technicians, and equipment [113]: from ferrous objects becoming projectiles to RF heating causing burns. Ensuring that the scanner provides sufficiently detailed images without artifacts caused by other equipment—and ensuring that this other equipment is functioning properly within the electromagnetically-volatile MRI environment—is part of designing MRI-compatible interfaces.

Magnetoencephalography (MEG) is the magnetic equivalent to EEG and shares its high temporal resolution, though it provides much better localization due to the magnetic “transparency” of the skull [105]. The magnetic signals emitted by the brain are quite weak and so the MEG scanner has to be quite sensitive, which means that it is also susceptible to electromagnetic interference, similar to fMRI [19, 204].

Functional near-infrared spectroscopy (fNIRS) has the temporal disadvantages of fMRI and the spatial disadvantages of EEG, but it is a portable and relatively inexpensive way to measure shallow cortical activity. fNIRS has been used in behavioural neuroscience to measure changes in regional hemodynamic response associated with music listening [140, 281], though perhaps it is best suited for BCIs [208].

2.1.2 Feedback Manipulation in Auditory-Motor Integration Tasks

Manipulating feedback within a study is not the only way to probe the brain [10, 295], however, this method does allow differential examination of two interacting systems, by controlling one feedback modality while varying the other. In order to elevate the possibilities provided by electronic musical instruments, we first focus briefly on the types of tasks that have been employed in psychological studies, and how the experimenters were able to manipulate feedback. With an overlearned task¹, the subjects rely on their experience to anticipate their next action—utilizing internal feedforward models—while at the same time, residual conflict is resolved by reducing the error

¹An overlearned task is one in that the subject has mastered and subsequent training does not increase performance at the task (i.e., subject performance saturates).

between the desired state and the perceived state by the use of negative feedback mechanisms, through auditory, visual, proprioceptive, or somatosensory modalities [10, 100, 240, 273, 274, 282, 295]. This adaptation to the new feedback is then incorporated into their internal models.

One of the simplest ways to manipulate visual feedback, without electronic means, is with the use of a mirror. Fink et al. used Luria's bimanual co-ordination task, opening and closing the hands in-phase and out-of-phase with each other, while subjects were shown their left hand, or the mirror image of their right hand presented on their left side in place of their left hand [70]. Using a PET scan, the subjects showed significant differential brain activity, with respect to the control, during the out-of-phase portion of the task where non-veridical and incongruent visual feedback was presented: a response to the conflict between proprioceptive and visual feedback. This study, as a control, and another fMRI study by Christensen et al. used volitional and non-volitional movements of limbs to investigate interactions with visual feedback [43].

Studies involving feedback manipulation of the motor and auditory systems have usually relied on speech, and in some more musical cases, singing. The voice, in concert with the ears and brain, is the most embodied auditory-motor control system humans have. Delayed auditory feedback has been used therapeutically with people who stutter [2], however for those who do not stutter, delayed auditory feedback by a larger amount (200 ms) resulted in significant dysfluencies [256]. Stuttering is thought to be an effect of deficiencies in the feedforward system, whereas peripheral feedback mechanisms account for dysfluency in normal speakers with delayed auditory feedback [256, 273]. Tourville et al. in an fMRI study investigated the effect of shifting the first formant during speech and how subjects compensated for the manipulated auditory feedback [273]. A later neuroimaging study contrasted the manipulation of the first formant in speech by way of auditory feedback compared to somatosensory feedback by inflating a small balloon within the subjects' mouths [100]. Several speech studies have employed pitch manipulations in auditory feedback to investigate associated neural activity [78, 175].

Pitch manipulation in singers allows fixed or dynamic pitch manipulations within a trial. Jones et al. performed a behavioural experiment comparing adaptation to pitch manipulation in singers and non-singers, where the type of feedback (either shifted or unaltered) was fixed during each trial [130]. Subjects were first trained to sing a

target pitch as accurately as possible with unaltered feedback for ten trials, then the feedback was shifted down for 30 trials while they were to sing the same target pitch (by essentially ignoring the feedback), and finally unaltered feedback was presented while some participants sang the same target note and others sang a new target note. It was noted that singers were initially more resistant to compensate for the altered auditory feedback than the non-singers, though after many trials the singers also compensated. Once the feedback was returned to an unaltered pitch, singers continued to compensate and sang the target tone sharp, whereas non-singers did not. This implied that singers were relying more heavily on internal feedforward models and showed plasticity as they adapted to the pitch-altered feedback. Zarate and Zatorre performed an fMRI study of singers and non-musicians who were asked to sing target pitches with and without pitch shifted auditory feedback [292, 293]. In the former conditions, auditory feedback was shifted up or down by at most two semitones after a delay of 1–1.5 seconds, and subjects were asked to either ignore or compensate for the shift in auditory feedback (i.e., ignore the feedback and continue singing the target note, or compensate for a downward shift in the feedback by singing a higher pitch in order to match the pitch-shifted voice with the target pitch). Singers were more accurate overall and better able to ignore the feedback than non-musicians, eliciting different activation patterns for the “ignore” condition, and enhancement of activity in certain areas during the “compensate” task. The authors proposed that there exist two mechanisms for audio-vocal integration: a basic feedback system and an augmented system accessible through training, which allows for a strengthened connection between monitoring of acoustic feedback and motor compensation. A similar experiment containing only the “compensate” condition was carried out by Toyomura et al., using several pitch shifts within a single trial [274].

Pa and Hickok performed a study with skilled pianists to differentiate between auditory-motor integration, not by modifying the target acoustic information (novel piano melodies), but by asking participants to covertly hum (i.e., without vocalizing) or covertly play (i.e., without moving) the melody after hearing it [195]. Thus both conditions provided no somatosensory or proprioceptive feedback, and auditory feedback was purely imagined during the rehearsal phase. The authors linked certain brain regions to sensory-vocal integration, while another region was associated with sensory-manual transformations. Maidhof et al. published an ERP study involving

pitch expectancy during listening and performance tasks with trained pianists [168]. The bimanual performance task involved the blindfolded subject playing scales, while occasionally (once every 40th to every 60th note) the auditory feedback of only one note (in either the right or left hand) was lowered by a semitone. The perception task had the subjects listen to the same pre-recorded scales played by a non-participant with the same random distribution of pitch-lowered notes. There was no behavioural result during performance (i.e., on the temporal placement of successive notes) upon hearing the incorrect pitch, as the subjects were informed that they were to continue playing after feedback manipulation. A more pronounced mismatch negativity was found during the action task than in the perception task, arising from the stronger expectancy violation during performance.

2.1.3 Probing the Auditory-Motor System with Electronic Musical Instruments

Electronic musical instruments have the added benefit over acoustic instruments of the ability to decouple the action gesture on the manipulandum² from its acoustical correlate. That is, given an input device that measures some form of physical human-machine interaction, the resultant generation or modification of sound can be any arbitrary function of the gestural input signals, be it simple or complex, static or adaptive, stochastic or deterministic. A primary limitation of this flexibility, when presenting an instrument (the input device, mapping, and sound synthesis taken as a whole) to a subject for a psychological study, is the participant's ability to build a mental model of the operation of the device. As Stassen et al. note, for optimal human-machine performance, the human operator has to form an internal representation of the machine's statics and dynamics, the task at hand, and the disturbances for which he/she must compensate [251]. For an interaction to have any meaning, and thus be useful and provide repeatable results for a study, the subjects must be able to learn the results of their actions, thus a random mapping, especially one which is time-varying, may be quite useless. Given a congruent mapping, particularly one that is naturalistic, a subject will be able to learn to perform a task by interacting with the manipulandum and internalizing the feedback in order to embody the instrument as a whole. Just as humans

²*Manipulandum* is the term neuroscientists and motor control researchers use to refer to an interface that is manipulated by the subject.

learn to vocalize and acquire fine control over their own vocal motor end-effectors by incorporating auditory and somatosensory feedback, a subject will learn to manipulate the device while taking in whatever useful sensory information is provided.

The control of pitch is usually a function of position, place, or displacement, though it could be equally a function of velocity, force, acceleration, or torque. Western classical music has traditionally emphasized pitch, and thus it is usually mapped to the selection (either on a discrete or continuous scale) of an absolute finger contact position, for instance. It is much simpler to perform this type of gesture than the selection of a velocity or isometric force due in part to internal kinematic controls, but also to the fact that in order to apply a specific force or velocity, a finger must start at zero and ramp up to the desired measure. Whereas the finger can be moved to the proper position before the interaction gesture is performed (e.g., moving the finger into place before depressing a key, valve, or string).

In hand drum playing, for example, the performer uses pressure on the drum skin to increase the tension and thus increase the pitch of the instrument. This interaction could be mimicked using an electronic instrument. Studies consisting of single-handed gestures controlling both attack and pitch, or bimanual tasks, where one hand controls pitch while the other controls the attack, are both possible. Parametric variation of force-to-pitch functions, for proportional as well as inverse-proportional control, in continuous or discrete steps could be implemented. Extending the synthesis from percussive to continuous sounds enables the use of on-line manipulation of pitch in order to probe the subject's ability to ignore or compensate for the shifting of feedback. In this case, the "ignore" condition is a gestural target, whereas the "compensate" condition is more of an auditory target. Similar parametric studies involving not only pressure but position, velocity, and acceleration could also be employed with the same types of auditory feedback, or conversely the same gestures could be used with auditory feedback of a different timbre or pitch register.

An interesting auditory manipulation is the use of non-veridical yet congruent synthesized feedback. That is, with certain well-learned mappings it may be possible to provide the subject with sounds that are not actually controlled by the subject, yet the subject is lead to think he/she is in control. For example, pitches are heard to be more in tune than the subject had actually achieved through his/her actions, though the note onset would coincide with the action—for instance if they played the wrong key

on a keyboard, but the correct pitch sounded. This would perhaps provide a bridge between the purely perceptual listening tasks and the performance tasks where performers heard their errors and also heard randomly inserted pitch errors (which they did not cause through incorrect key selection), which was presented by Maidhof et al. [167, 168].

A general limitation of musical instruments is that while the auditory feedback can be arbitrarily mapped, the somatosensory feedback cannot. With the introduction of haptic and robotic devices, the statics and dynamics of the input device itself can be manipulated. Actuators and linkages could be used to manipulate the feel, shape, and response of the interaction. As this is not purely a motor-auditory interaction, somatosensory inputs must be taken into account (and perhaps even proprioception for larger gestures) when dealing with motor tasks involving sensory feedback. The addition of haptic feedback also allows the experimenter to produce volitional versus non-volitional movement conditions by inducing the subject to perform certain gestures robotically. As well, this can be done with and without auditory feedback.

2.1.4 Musical Tasks

Neuroimaging studies employing musical tasks have allowed neuroscientists to build upon their knowledge of sensorimotor integration and its relation to motor learning and neural plasticity [41, 253, 292]. fMRI studies involving music audition and imagery have led to questions involving music performance, and how each of these tasks differ [10, 40, 178]. Many so-called “musical tasks” used in neuroimaging studies, like finger tapping, are simplistic in nature in order to precisely identify involvement of the underlying brain structures. Music performance, however, is more often a complex sequence of movements coupled with visual, auditory, and somatosensory feedback. In order to undertake more ecologically-sound³ musical motor learning and auditory-

³ The term ecological validity is perhaps vague, so we keep with Guastavino’s usage: “Instructions were given to subjects to direct their response strategy towards everyday listening situations, so that they would react, to some extent, as if the[y] were in the actual situation...rather than in the abstract situation of a laboratory experiment” [97] and “to take contextual and environmental cues into consideration. An experimental protocol is ecologically valid if the participants react, to some extent, as if they were in a natural situation. In other words, the laboratory must be like life in regards to the context of the question asked.” [98] And so, we sought to study music performance from the standpoint that for musicians to play music meant that they required an instrument that, to some extent, felt and responded naturally and was comfortably playable. Even if their surroundings were vastly different from the conditions in

motor integration tasks, electronic MRI-compatible musical instruments are required. These interfaces must afford the feel of acoustic instruments, along with the ability to provide auditory feedback, and to synchronize behavioural data with functional imaging sequences. Presented here is a review of musical tasks performed in the MRI scanner as well as the associated interfaces.

Rhythmic and Musical Tasks

We first examine musical studies that used electronically-sensed input devices. There have been several studies of musicians that have incorporated “dummy” interfaces (a brief review can be found in [113]), usually an electronic piano devoid of any electronics, but in order to measure, log, and synchronize behavioural data with stimuli, feedback, and neural activity, an electronic and computerized system is required. The most basic of interfaces is the two-position switch, commonly referred to as a response box.⁴ These response boxes are usually employed for making forced-choice selections in response to stimuli [30, 31, 174], they are also used for determining timing accuracy in rhythmic finger tapping tasks [58]. Finger tapping tasks are conducive to the MRI environment because the minimal physical exertion reduces the possibility of movement-related image artifacts. As a musical task, finger tapping experiments help to elucidate the underlying systems responsible for auditor-motor integration, particularly for complex musical rhythms. Chen et al. employed a standard computer mouse for rhythmic tapping tasks [39] in a 1.5 T scanner. However, when a similar task was attempted in

which they were used to playing, at least the instrument would not further detract nor distract from their performance. That isn't to say that these emulations of musical instruments are meant for stage performance, thus comparing their playability is moot. If an emulation of an instrument is sufficient to perform the musical task at hand, and the subject approaches the performance of the emulated instrument as if he or she were playing its acoustic kin, then we can accept that the subject's behaviour and the underpinning neurological activity are, to some extent, the same. And as neuroimaging experiments rely on the differential neural activity and differential behavioural measures, the common-mode differences between stage performance and in-scanner performance are null. If the reader still has any doubts about the efficacy of using these emulated instruments within the MRI scanner in order to test musical performance, one may perform the same task outside the scanner with an acoustic instrument to assess purely behavioural differences.

⁴There are myriad commercial offerings of MRI-compatible optical computer keyboards, keypads, and response boxes. Commercial suppliers, however, do not usually give details about the inner workings of their devices. Companies include Resonance Technology Inc., Cambridge Research Systems Ltd., Current Designs Inc., NAtA Technologies, among others. These devices most likely use fibre optic amplitude sensing of diffusely scattered, transmitted, or reflected light.

a 3 T scanner, the computer mouse no longer functioned and an optical solution was developed [113, 252]. This is one of the main issues with using electronic devices in the MRI scanner: even if the device functions (i.e., the device works properly and there are no image artifacts) under certain conditions, there are no guarantees the device will function properly if the wires are moved slightly or if the study is performed in a different or more powerful scanner. This is especially the case if the electronic interface is brought within the imaging volume and near the isocentre of the scanner where the head is positioned and the magnetic field is the strongest and most homogeneous.

An Interlink force-sensitive resistor was used by Lewis et al. to measure tapping timing in an fMRI study [153]. The authors note that baseline noise was near 10% of peak force, however not all subjects pressed with sufficient force, and thus 5.9% of their data had to be thrown out. Though it is not stated in the article, one might assume they were interested in force information in addition to timing information but the sensor, when placed in the scanner, could not provide sufficiently clean measurements. It comes as no surprise that a later study published by the same research team used an MRI-compatible load cell⁵ during a finger opposition squeezing task [207]. While they made mention of its compatibility within the MRI scanner, the authors also mentioned the need to filter out the force signal, especially during MRI image acquisition.

A high-resolution angular velocity finger-sensing system was thoroughly tested for MRI-compatibility by Schaechter et al. [230]. The system utilized a commercial microelectromechanical system (MEMS) gyroscope, namely the Analog Devices Inc. ADXRS150. The advantage of using a MEMS sensor, while electronic, is that the signal processing is performed on-board the same chip as the sensing thereby eliminating high impedance outputs terminated with long runs of cable which could be more easily corrupted by electromagnetic interference. The effects of the scanner's static magnetic and RF fields on the operation of the device were tested, along with measures of decreased signal-to-noise ratio (SNR) of the MRI images. While the static magnetic field had no appreciable effect on the sensors, the sensor SNR was decreased to between 34.36 dB and 24.2 dB depending on the coil arrangement.⁶ The reduction in SNR for

⁵The commercial load cell was made by Novatech Measurements Inc., and from their product listings it seems to be a resistive strain gauge device, modified from their usual items by replacing ferromagnetic components with aluminium and brass [190].

⁶Again, while a MEMS or other electronic sensor may function properly for a specific application, changing scanning parameters, the placement of wires within the scanner, or migrating the task to a

a functional image sequence was from 30 dB to 29.96 dB, which, as they mention, is a negligible effect.

De Luca et al. presented an MRI-compatible device, using the fairly common optical quadrature method, to sense finger displacement in a bimanual tapping task [54]. Optical quadrature uses a periodic amplitude mask that moves relative to two pairs of sensing elements (i.e., two emitter-receiver pairs), which are arranged 90 degrees out of phase with one another [113]. Displacement is measured by counting the number of pulses, however, only relative position is attainable without calibration on start-up (and mechanical constraints for rotary encoders).⁷ The displacement direction is determined by the leading pulse phase. Fairly stable measures of velocity and acceleration may be acquired to build a kinematic profile. A Belgian research team conducted bimanual MRI tests using a shaft encoder [55], though few specifics are given. Non-ferromagnetic fibre optic rotary encoders are available, however, from Micronor Inc., specifically designed for MRI-compatible applications [179].

While perhaps not strictly musical in nature, James et al. tested two commercially available computer keyboards for MRI-compatibility in a 3 T head-dedicated scanner using a glass phantom, and found that two of the keyboards had decreases in the temporal SNR of 10% and 48% [124]. Furthermore, they note that field homogeneity would likely decrease if the interface were brought into the bore of a whole-body scanner. Therefore, without proper testing, the impact of using electronic devices in the scanner on MRI image quality is unknown.

Piano Keyboard Interfaces

Moving on to more musical tasks, Bangert et al. published an fMRI study that used a soundless, yet optically-sensed, piano keyboard [10]. The keyboard, removed from an acoustic grand piano and stripped of ferromagnetic components to maintain the “feel”, was augmented with fibre optic cables. The movement of each key was sensed with two stationary pairs of aligned emitters and receivers in combination with a moving optical interrupter mounted on the key. No auditory feedback was provided during this experiment.

higher field strength scanner will require additional testing, time, and cost.

⁷Absolute position is also achievable using “gray-code”—a binary counting method where only one bit changes at a time when transitioning from one state to the next.

We developed a similar fMRI-compatible device [113, 116], but instead of using two detectors for each key, a single detector and a two-threshold state-machine was employed. This keyboard measured the intensity of light reflected off a mirror that was raised into alignment with the fibre optic pair when a key was depressed. Control electronics connected via the fibre optic cables enabled the interface to comply with the MIDI (Musical Instrument Digital Interface) specification. No artifacts or magnetic field distortion were noted when scanning with the keyboard at arm's length from the imaging volume. More recent upgrades to this keyboard, discussed in Section 5.1, included a more compact message protocol, the addition of a message time-stamp, and a USB interface, in place of the previous use of standard MIDI hardware and messaging. This keyboard was used successfully in a rhythm and melody task [23, 24, 25]. Subjects were blindfolded in order to negate any effects of the visual scene, while auditory stimuli and feedback were provided using MRI-compatible in-ear monitors.

Another piano keyboard was used in an MRI study of jazz improvisation by Limb and Braun [157]. The keyboard is described as generating MIDI messages and being non-ferromagnetic, though no sensing methods were discussed. As well, no key velocity or force measurements were produced. The authors note the 35-key plastic keyboard was custom-built by Mag Design and Engineering. The company makes a wide range of MRI-compatible devices, most include electronic components [165]. The item listed as a MIDI piano keyboard is described as a commercial MIDI keyboard stripped of its internal circuitry and replaced with custom parts for increased reliability within the MRI environment. The company does not discuss any testing procedures that ensure their electronic devices were MRI-compatible.

Rehabilitation

Music has the ability to aide in the rehabilitation of patients who have suffered a stroke or to allow those with Parkinson's disease to move more freely [13, 33, 267]. It has been shown that music performance as part of therapy can improve motor skills more effectively than standard rehabilitation alone [3]. Electronic devices which not only engage individuals to follow their therapeutic routines at home, but also monitor their progress in clinical settings are required [77]. It is clear that given the current research in music-supported therapy for motor rehabilitation using MRI [223], researchers re-

quire MRI-compatible electronic musical instruments to allow behavioural measures of motor performance to be correlated with functional imaging data. To understand why music is an especially effective tool for rehabilitation, the underlying neural structures and mechanisms need to be better understood.

Music-supported therapy for the rehabilitation of fine motor skills after stroke is gaining support because of its ability to sustain patient interest and the evidence which demonstrates its effectiveness compared with conventional and constraint-induced therapies [234]. The setup used in this therapy by Schneider et al. included an electronic MIDI piano on which subjects practiced fine motor tasks: pressing keys in musical sequences on eight of the white keys with auditory feedback. As a metric of fine motor skill before and after the longitudinal therapy program, a computerized measurement system, comprised of ultrasonic transceivers, was used to acquire finger movements at a per-channel rate of 66 Hz during a self-paced tapping task. Subjects were asked to tap as quickly as possible while frequency, number of inversions of velocity profile (a measure of motion smoothness), and average angular velocity were collected. Another similar study included the addition of electroencephalography to correlate changes in event-related synchronization and behaviour, showing improved motor skill and neural reorganization after music-supported therapy [3]. A case study performed with the same therapy employed fMRI and TMS to show again the effectiveness of music-supported therapy through improved behavioural measures while demonstrating the neural changes associated with improved fine motor skill. Of interest is the fact that the fMRI tasks comprised of a finger tapping segment and a listening segment only. For the finger tapping segment, no movement analysis was performed nor was any auditory feedback presented. Clearly, the lack of MRI-compatible tools to provide direct measurement of movement and auditory feedback has limited the experimental possibilities.

2.1.5 Non-Musical MRI-Compatible Interfaces

Over the last decade, a host of MRI-compatible devices have been designed. Many of these devices were meant for clinical treatment and diagnostic use, however, more and more devices that offer some form of interactivity for functional neuroimaging studies are being developed. The MRI scanner, due to its intense magnetic fields, strong

pulsatile RF emissions, and extremely sensitive signal acquisition electronics, is a challenging environment for which to design safe and functional electronic systems and interfaces. This section presents current research and development of MRI-compatible devices for human-computer interaction; this includes input devices containing sensors, and haptic and/or robotic systems, which contain both sensors and actuators. The discussion will concentrate on devices that measure direct interaction, which may have applications in musical interfaces, as opposed to brain-computer interfaces where the scanner or surface electrodes are employed along with signal processing to extract patterns of neural activity that are then used for computer control or feedback generation [283, 288].

Force Measurements

From brain mapping to stroke recovery studies, MRI-compatible force sensors are prominent input devices that could be applied to musical tasks, including drum stick grip and fingerboard interaction. Experiments in which force is measured generally conform to grip (or squeezing) force tasks and isometric hand forces produced on a quasi-stationary joystick. A simple method of measuring grip force is with a pneumatic bladder, similar to the emergency alarm bell that comes standard with MRI scanners. In fact, early studies that measured neural activity with associated grip force used a sphygmomanometer, though this was not an electronic input device [270]. Liu et al. performed force measurements during an fMRI task while also recording surface electromyogram muscle recruitment. Their force sensor consisted of a hydraulic bladder connected to a water reservoir located outside of the MRI environment [49, 159]. The water reservoir contained a commercial electronic pressure sensor (Entran EPX-N1). Two different bladders were employed, a larger one to measure grip force, and a smaller one to measure finger forces [158]. As the part of the input device placed in the MRI environment did not contain any ferromagnetic or electronic components, MRI-compatibility was guaranteed.

Various researchers have used electronic strain gauges and load cells to measure force inside the MRI scanner. Cramer et al. used load cells manufactured by Interface Inc. (SSM-AJ-50 and SSM-AJ-250) for dynamometer squeezing and tapping studies of stroke patients [45, 46, 186]. The sealed load cells were made of anodized aluminium

and thus were well shielded, however, no details about MRI signal degradation was discussed. Hidler et al. designed a non-ferromagnetic joystick-like 6-degree of freedom (DOF) isometric wrist force-sensing device [110]. The device used a load cell manufactured by JR3 Inc. (35E15A). SNR in a 3 T MRI scanner was measured with and without the device using a function imaging sequence with a $NiCl_2$ spherical phantom. No significant differences in image SNR were found. They mentioned that the magnet did not affect the performance of the filtered load cell. A similar load cell (JR3 Inc.'s 45E12A-Z1) was used by Krainak et al. [141] to measure shoulder and elbow torque in an fMRI task. The load cell was made from non-ferrous aluminium, brass, and titanium. Measurements were made during scanning with and without filters (8th-order Butterworth low-pass at 15 Hz; 8th-order elliptic low-pass at 10 Hz). While the filters increased the load cell SNR (and reduced the signal bandwidth) significantly during image acquisition, they do not mention the difference between these two conditions and a baseline SNR without image acquisition. No significant reduction in MRI SNR or image artifacts were discovered. Isometric force measurements are conducive to fMRI tasks as they obviate motion artifacts [188], however, musical tasks often employ movement and dynamic forces, which would not be measurable with the aforementioned setup. A quartz piezoelectric force sensor was employed in [91] as part of an MRI-compatible ergometer, though the imaging volume was constrained to the quadriceps and not the brain.

An interesting design by van Duinen et al. in [277], presents an MRI-compatible force-sensing finger abduction device, made from electromagnetically-compensated strain gauges (MFLA-5-350-1L manufactured by Tokyo Sokki Kenkyujo Co. Ltd.) arranged in a Wheatstone bridge. The sensor was powered with a shielded lead-acid battery located within the MRI environment. The signal processing electronics for measurement and encoding were located on the device in a shielded box. The strain gauge signal was amplified with an integrated instrumentation amplifier and encoded with a voltage-to-frequency converter connected to an optical transmitter. Via fibre optic cables the frequency-encoded force information was sent to the control room, where it was decoded, further processed, and recorded on a computer. The sensor bandwidth was 250 Hz. No image artifacts were discovered with the device during scanning nor were there significant differences in baseline force signals during image acquisition.⁸

⁸While image artifacts arising from spurious frequencies in the RF band are one detrimental factor

This design is interesting because of its hybrid electrical-optical-electrical signal chain. Normally, a designer tries to limit conversion stages because of their inherent addition of noise to the signal. Commercial electronic sensors are less expensive and more widely available than optical sensors, however, optical fibre sensors are intrinsically MRI-compatible.

Several optical force-sensing techniques have been presented in the academic literature. Ehrsson et al. have used an optical force sensor in several studies to measure differences in neural activity as a function of varying finger and grip force [12, 61, 62]. No technical details were provided for the transducer, though an article citing this one by Fox et al. used an optical device for force sensing of button presses. This system measured the “graded interruption of an optical beam” [75]. An optical quadrature encoder was used by Yamane and Mitsumori for force sensing in rat gastrocnemius muscles [287] which could easily be adapted for musical purposes as a force and position sensor.⁹

Tada and his team have been working at developing an optical multi-axis force sensor since 2002 [260]. The first two-axis version using a single emitter and a four-quadrant photodetector, each coupled to a separate fibre optic cable mounted on a non-ferromagnetic elastically-deformable frame engineered with specific axes of deformation. The four receiving fibres were arranged to differentially accept changes in light intensity depending on alignment. This device was extended to a 6-DOF force/torque sensor, while improving the sensor accuracy and robustness, by substituting a light-emitting diode (LED) for the laser source and with an improved fixation of the fibres to decrease vibration-induced noise [262]. In order to further reduce the sensitivity to extraneous mechanical oscillations, a lens mounted on a deformable frame coupled

when considering MRI-compatibility, field homogeneity is also an important factor [192]. Higher field scanners (e.g., 3 T) are used for research purposes because of their higher SNR, and are also more sensitive than lower field scanners (e.g., 1.5 T) to field inhomogeneity [139]. Furthermore, image SNR is dependent on the specific imaging sequence, and making assumptions about device compatibility using one sequence may not accurately reflect the SNR of a different sequence. Unfortunately, there is no standard protocol for measuring fMRI-compatibility, and as a result, it is hard to assess and compare various devices in terms of compatibility. Field homogeneity is a determining factor in SNR and is independent of the task and imaging sequence, thus it is suggested as the best predictor of MRI-compatibility.

⁹Comparisons of MRI-compatible devices for studies with rats are less than optimal, as the scanners employed are vastly different. On the other hand, rat input devices include surgically-implanted force transducers, which might have applications in next-generation cyborg musicians [89]. That source also contains various references to ergometers for human magnetic resonance spectroscopy studies, though not specifically fMRI.

light from an emitting fibre to a receiving fibre. The advantage of this system is its use of multi-core fibres (217 cores per fibre) and micrometry which effectively optically magnifies the displacement. At the distal ends of the fibres (i.e., in the control room and away from the sensor itself) lenses were used to focus the light; in the case of the receiver, onto a four-quadrant photodiode. As the frame was deformed, light was coupled differentially to the spatially-arranged fibre cores, which translated into a variation in position (and thus differential intensity) of the beam being focused onto the four-quadrant photodiode at the distal end of the receiving fibre [259]. A more recent extension to their second-generation sensor incorporated two-axis sensing [261]. The device was intrinsically MRI-compatible, and the sensing strategy could be adapted to musical applications where single-point position, force, or vibration measurements are required. The sensor was relatively inexpensive and simple to implement, though optical alignment must be guaranteed. An all-optical torque sensor similar to the early work by Tada was implemented by Renaud and de Mathelin, except that the optical sensor was operated in a reflective mode, thereby reducing sensitivity to extraneous movement of the fibres [218]. Additional force and torque sensors are covered in this section, where they are applied to haptic devices.

Virtual Reality and MRI-Compatible Input Devices

This section deals with input devices for virtual reality (VR), and other position-based fMRI motor tasks. While there have been several studies using computer mice to measure cursor position [121], a mouse-like absolute position encoder (Felix, by Altra MicroPoint) was used by Balslev et al. [7, 8], though no MRI-compatibility issues were discussed. A modified Gravis gamepad was used in several virtual reality studies, along with Shape Tape (Measurand Inc.) [185, 271]. Shape Tape is an MRI-compatible optical bend and twist sensor, discussed in [113]. Mraz et al. used Shape Tape as a stylus for writing and as a data glove for VR navigation and to determine pointing and pinching gestures [143]. Other MRI-compatible stylus tracking and writing devices were developed by Reithler et al. [217] and Tam et al. [264]. The former used a distance-proportional resistive voltage divider with the stylus acting as the potentiometer's wiper, which needed to be heavily filtered to ensure functionality. The stylus was constrained to a single predefined path, and was not considered a writing tablet. Tam

et al. utilized a non-ferromagnetic resistive touch tablet (made by 3M) with additional shielding and electrical filtering to ensure compatibility. They note that an improved touch panel with better suppression of unwanted touches, pressure sensitivity, and multi-touch is desirable, though even the current device could translate well into a useful device for musical applications [291].

Eye tracking systems have been employed on several occasions to varying degrees of success. Felblinger et al. reported some success using electrooculography, a measurement of the resting potential of the retina, by placing electrodes near the eyes [69]. Not only was the system susceptible to RF noise, but the Hall effect due to pulsatile arterial flow is said to have increased the baseline eye measurement noise inside the scanner. An infrared fibre optic eye tracking system discussed in [137], details the use of LED emitters and photodiode receivers connected differentially to sense horizontal saccades. The use of fibre optics allows artifact-free functioning, though eye motion was only measurable in one axis. Conversely, a third eye-tracking system presented in [90] employs a video system to acquire gaze direction. The camera was situated within the MRI environment, about 3m away from the imaging volume, with fibre optic cables transferring data from the camera to the control room. No video or MRI signal degradation was found during scanning. Setup time, including alignment and focusing of the optics (i.e., infrared illumination, camera, lens, and mirror) as well as system calibration was less than ten minutes. Eye tracking systems have been used in musical applications on several occasions [118, 136].

Clinical Instruments

There are several examples in the literature of high-precision MRI-compatible sensors, which are used for clinical diagnosis or intervention, that could be adapted to MRI-compatible musical interfaces. As more and more surgeries are taking place within an MRI scanner, devices for patient monitoring and surgical probes are becoming commercially available. Pinet et al. (from FISO Technologies Inc.) note the need for patient temperature monitoring during MRI-based procedures due to the possibility of RF heating [203]. FISO Technologies Inc. sells MRI-compatible optical temperature, pressure, and displacement sensors, mainly based on Fabry-Pérot etalons (refer to Section 2.2.2), though one of their temperature sensors uses the temperature-dependent

transmission spectrum of a Gallium-Arsenide crystal.

A first prototype MRI-compatible needle deflection sensor by Park et al. used a single fibre Bragg grating (FBG, refer to Section 2.2.4) bonded at the base to the outside of a biopsy needle [196]. A second prototype was designed with three fibre sensors inserted into a needle [197]. Each fibre was inscribed with two gratings, positioned at 22 *mm* and 85 *mm* from the needle's base. The temperature-compensated sensors were able to resolve needle deformation and tip position with less than a millimetre error, even for two-point bending. FBGs were also used in a quasi-distributed temperature sensor for MRI-guided tissue cryo-ablation [227]. The fibre was inscribed with ten gratings to provide a temperature profile over the 6 *cm*-long probe, achieving a spatial resolution of 6.5 *mm* and a temporal resolution of 5 s. It was likely that the low temporal resolution was a limitation of the interrogation and acquisition system, and not due to any inherent property of the sensor, though no strain compensation was discussed. A FBG-based vibroacoustic sensor was presented in [272] with applications in MRI-compatible voice and heartbeat detection. The sensor was interrogated with a laser source and a photodiode, ostensibly converting the strain-induced Bragg shift into an intensity modulation, capable of detecting vibrations of up to 3 *kHz*.

A fibre optic MRI-compatible position and magnetic-field sensor was implemented by Bock et al. that measured the Faraday effect (refer to Section 2.2.3) in a high Verdet constant terbium gallium garnet crystal [16, 17]. A measure of field strength as a function of position and angle in a single plane during scanning was enabled by synchronizing the sensor acquisition hardware with the gradient field control signal. The laser launched down an optical fibre was then linearly polarized before being rotated by the crystal, and then analyzed with another polarizer before being sensed with a photodiode. The system was capable of sensing modulations of up to 30 *kHz*. A review of MRI position-sensing techniques can be found in [14].

Haptic Devices, Surgical Robotics, and Actuators

There are many MRI-compatible haptic devices in the literature, consisting of devices for fMRI motor and somatosensory tasks, rehabilitation, and remote or robot-assisted MRI-guided surgical procedures. Active haptic devices incorporate sensors, actuators, and a means of feedback control or mapping. Actuators, because of their active role in

converting electrical energy into mechanical energy, have inherent MRI-compatibility issues, as they often require high voltages or currents to operate. And the added constraint imposed by the MRI environment limiting the use of ferromagnetic materials reduces the efficiency of the electro-mechanical conversion process. Presented here is a brief overview of MRI-compatible haptic devices and actuators, though the reader is referred to [84, 113] for a more detailed survey of MRI-compatible actuators.

A shielded DC motor was used by Li et al. in the implementation of 1-DOF and 3-DOF MRI-compatible haptic interfaces for motor control studies [155]. The DC motor was placed about 2 m from the scanner, and though there was a small but noticeable drop in SNR when the system was in use during an MRI scan of a phantom, the authors mentioned “the interface does not interfere [with] the fMRI scanner.” Burdet et al. noticed that placing an active DC motor within the MRI room had no effect on a scan with a phantom but had significant effects with a human subject [27], as people often act as antennas for RF noise. Thus, they built a hydrostatic transmission system to actuate a robotic slave device. This device was augmented with an optical quadrature position encoder and optical torque sensor [83]. A similar torque sensor was used in combination with an ultrasonic (piezoelectric) motor for a haptic rotary device [72]. In order to provide additional torque control and the ability to back-drive the device an electrorheological clutch system was designed [35]. Electrorheological fluids (ERF) were also used by Khanicheh et al. in their hand-rehabilitation device to allow variable resistive torque to subject-applied forces [135]. A high voltage applied across the ERF clutch increases the force required to squeeze the device’s handle. An optical encoder and a load cell were used to sense position and force, respectively. Ultrasonic motors were also used in an MRI-compatible surgical robot to position a biopsy needle [147]. Optical encoders were used to determine the position of the device.

Suminski et al. proposed an MRI-compatible pneumatic actuator which produced torques on the subject’s wrist [258]. The joint angle was sensed with an optical encoder, while air pressure within the device was measured with a low-pass filtered commercial electronic pressure sensor utilizing microfabricated piezoresistors as strain gauges in a bridge arrangement (Honeywell 26PC). Commercially available non-ferromagnetic air vane motors from Deprag Inc. are compact, lightweight, and can be powered with a remote pneumatic source [57].

A three-phase electrostatic actuator was implemented in a 2-DOF haptic joystick

that used an optical force sensor, which measured reflected intensity [103]. Burton et al. used a stepper motor to provide embossed letters as tactile stimulation to blind and sighted individuals via a rotating drum [28]. An optical encoder was used to track the drum's position. Synchronization of the rotation of the motor with the scanning sequence minimized RF-induced artifacts, though compatibility testing was only carried out with a phantom. Riener et al. presented an MRI-compatible Lorentz force actuator that made use of an energized coil interacting with the magnetic field of the scanner to generate torques [219]. A control system was implemented to provide either active or passive force interaction with the use of an optical encoder to measure position and an optical force sensor, made from fibre optic cables in a simple differential intensity transmission sensing configuration. MRI-compatibility testing was performed only with scans of an MRI phantom. A Lorentz force actuator was also employed to deliver vibrotactile stimulation [92, 143], no artifacts were discovered in anatomical or functional human scans. An MRI-compatible piezoelectric vibrotactile actuator was presented by Harrington et al. [106]. While the coil-free design is less susceptible to MRI scanner interaction, the force produced by the piezoelectric actuator was relatively small compared to the previous electromagnetic designs, in spite of driving the device with a large voltage (around 100 V peak-to-peak).

Finally, electro-active polymer devices, including ionic polymer-metal composites and electrostrictive polymer actuators, have been discussed for use in MRI-compatible actuators and haptic devices [84]. Vogan et al. presented an MRI-compatible actuator to electronically adjust an MRI head coil antenna to increase SNR during scanning [280]. The device was comprised of a copper frame over which was stretched an acrylic film coated with compliant electrodes, made of silver and silicone oil. An applied voltage across the dielectric created a strain, compressing the polymer by up to half its original height. A similar device was designed by Carpi et al. using an accordion-like folded structure made of a silicone dielectric and silicone/carbon compliant electrodes [32]. No image artifacts or significant decreases in SNR were seen in MRI tests of a phantom.

Many of these actuation methods can be applied to MRI-compatible musical interfaces, especially for force feedback, haptic devices for virtual instrument design, and vibrotactile stimulation.

2.1.6 Discussion

This section presented the concept of musical instruments as composed of two components, a mechanical manipulandum and an acoustical correlate of the interaction with said manipulandum. This view is true of acoustic and electronic instruments, however, with the latter the acoustical correlate of the action is implemented as a concatenation of the mapping of gestural inputs onto synthesis parameters and the synthesizer that renders them as sound. Musical sounds usually have some variation in volume, pitch, and timbre. The control of these sonic properties within an electronic instrument may be any arbitrary mapping, with respect to the input gesture and within the limits of the input device and the synthesizer. The sound may be percussive or continuous, its pitch may be selected from a continuous span or it may be constrained to a discretized subset of pitches. On-line control of the pitch may be possible throughout the duration of the sound, without re-articulation, or pitch may be determined by the user only at note onset. Various mappings and musical targets may be presented and learned by a subject using the same interface, and thus differential analysis of behaviour and neuronal activation during auditory-motor tasks can be carried out. This ability would not be feasible, or perhaps even possible, without real-time electronic reconfiguration of the instrument. Haptic feedback could be used to reconfigure the mechanics of the instrument thereby probing the recruitment of somatosensory and proprioceptive modalities as well.

A review of MRI-compatible input devices, sensors, and actuation methods was presented. The issues arising from direct electronic sensing and actuation were discussed, as well as the limitations and benefits of certain devices. While the simplicity of implementation and the commercial availability of electronic sensors such as load cells and accelerometers results in a faster and less expensive implementation, thorough MRI-compatibility testing is required, offsetting any time or cost advantage. Unshielded devices are highly susceptible to RF interference, and may even cause artifacts, decreased magnetic field homogeneity, or decreased SNR. Fibre optic sensors are inherently MRI-compatible, though mechanical robustness and interrogation method must be engineered to ensure functional specifications are met. Likewise, actuators must be chosen to generate sufficient forces or torques, while ensuring MRI-compatibility. Remotely powered pneumatic and hydraulic systems ensure MRI-compatibility, though

they are often more complex and less flexible than electromechanical methods. Piezoelectric, electrostatic, and electro-active actuators are low-current (though usually high-voltage) devices, thus they generate smaller magnetic fields and thus interact less with the MRI scanner than electromagnetic devices. Electromagnetic devices may generate higher forces and torques due to their increased efficiency and driving power, though they are more likely to decrease MRI scanner SNR. Conductive coils, used to produce Lorentz forces in combination with the MRI scanner's permanent magnetic field, are susceptible to spurious actuation (caused by gradient field switching) and RF heating. As well, their efficiency is directly proportional to the magnetic field strength within the scanner bore, such that the generated force vector is limited by the position and orientation of the device within the scanner.

Furthermore, there is a serious problem with the lack of standardized procedures for testing MRI-compatibility. The difficulty in standardizing such a procedure results from the fact that testing an input device in a specific scanner with a specific protocol does not generalize well to other scanners, imaging sequences, or even a slightly different arrangement of wires or a more corpulent subject. There are two main factors when determining fMRI-compatibility from the imaging point of view: artifact-free functioning and high magnetic field homogeneity. While artifacts are generated by spuriously emitted RF frequencies or local distortions of the magnetic field, SNR of a functional image is a complex issue involving the scanning sequence, the subject, broadband RF noise, and field homogeneity. Newer scanners provide stronger and more homogeneous magnetic fields, translating into higher image SNR, and thus the significant correlation of even more minute phenomena. Decreasing the SNR by introducing a less-than-perfectly fMRI-compatible device into the scanner will only reduce the ability to find the very neural correlates the device and task are meant to uncover. Acquiring a magnetic field map and scanning with a phantom is a good first step toward ensuring compatibility. It is also advisable to use, as much as possible, materials and transduction techniques that are known to be inherently MRI-compatible toward the design of fMRI-compatible input devices, if these devices are to be used in multiple studies and scanners without compatibility testing with each use.

2.2 Optical Sensing

2.2.1 An Introduction to Fibre Optic Sensors

Fibre optic sensors are well-suited to MRI-compatible input devices due to their lack of metallic and ferromagnetic materials, current-free operation, immunity to electromagnetic interference, and low transmission loss over long distances. Fibre optic cables are structured optical waveguides comprised of an inner core with a higher index of refraction compared to the outer cladding. A buffer and/or jacket layer is usually applied to the outside of the optical fibre, however they play no role in transmitting light and are there only to protect and encapsulate the fibre. Fibre optic cables propagate light along the core through total internal reflection: when the incident angle is shallower than the critical angle with respect to the normal, all light is reflected (and none of it is refracted into the fibre cladding). If the cable is bent at a sharp angle, then the angle of incidence of a ray of light may exceed the critical angle and thus would be refracted into the cladding.

In order to make use of fibre optic cables in sensors, an optical system comprised of an interrogation scheme (utilizing an optical emitter and receiver), a transmission medium (namely fibre optic cables), and a means of coupling the measurand to a modulation of some optical property (such as the increased transmissivity of a medium under stain for force sensing). Fibre optic sensors are usually divided into two different categories: *extrinsic* (or hybrid fibre optic) sensors use fibre optic cables to convey light along the emitter-measurand and measurand-receiver paths, but utilize some other medium (e.g., an air gap or a compressible plastic) to transduce a change in the measurand, whereas *intrinsic* (or all-fibre) sensors use an optical fibre itself to sense the measurand as well as convey signals to and from the optoelectronic measurement equipment. Intrinsic sensors are attractive for industrial, commercial, and MRI applications where a compact single fibre sensor can be easily replaced or embedded in composite materials and proper alignment is guaranteed, though the cost and complexity of designing such sensors is perhaps more expensive and is currently an active research area [150].

As is discussed in my master's thesis [113], the sensing of changes in light intensity to measure position, force, flexion, or rotation based on bending losses, alignment of two optical fibres, or by modulating the transmissivity of an optical path is relatively

simple and inexpensive to implement. Single-ended intensity measurements, however, suffer from low sensitivity and selectivity compared to other optical techniques. Owing to their inherent measurement of relative amplitudes and reliance on imperfect emitters and detectors, the accuracy of these systems is limited by calibration and the ability to engineer robust solutions which compensate for endogenous and exogenous sources of error, such as internal noise and environmental illumination, respectively.

This section deals with interrogation techniques, other than pure light intensity measurements, to apply toward MRI-compatible devices. The constraint of fibre optic sensors is imposed in order to perform remote sensing with the measurement equipment in the control room, obviating the use of optoelectronic devices within the scanner environment. While other sensing techniques, optical or otherwise, are conceivable (such as video, optical motion capture, hydraulic tubes, or mechanical linkages), the following discussion will focus on the acquisition of force, pressure, displacement, acceleration, temperature, and magnetic field strength through the use of optical interrogation techniques based on the modulation of phase, wavelength, or polarization of light. While these measurements are also based on the use of photodiodes that measure light intensity to simplify the implementation (or render it tractable due to the dearth of inexpensive, high-speed optical spectrum analyzers), the theoretically defined signals provide absolute physical measures, in most cases. This section is organized by optical modulation technique; the theory and applications are discussed herein, and we will see that certain interrogation techniques enable the acquisition of different types of measurands depending on their spatial and temporal constraints.

Before we proceed, an important distinction between the terms distributed, quasi-distributed, and discrete sensors is in order. While discrete sensors are the simplest case, where a single measurement is taken at one point or across a sensitive region without regard or ability to distinguish its spatial location, a distributed sensor measures the continuous profile of the quantity in question (e.g., temperature or force) along the length of the sensitive region. Spatial resolution in the distributed case is mainly limited by the interrogation equipment and not the sensor. Quasi-distributed fibre optic sensors are comprised of many discrete sensors that lie along the same fibre spaced at (usually regular) intervals. The spatial information is extracted by interrogation of the fibre through multiplexing in time or frequency in order to separate the signals arising from each sensitive site. A special case of quasi-distributed sensing is as

before, but instead of an array of discrete sensors, several short distributed sensors are placed at regular intervals. This last case provides a trade-off of sensitivity and spatial selectivity, while also increasing the packing density and system complexity, as we shall see. General reference information can be found in [50, 94, 95, 111, 125, 134, 275].

2.2.2 Interferometric Sensors

Interferometry is the measurement of periodic constructive and destructive interference patterns which arise from two (or more) waves as they interact with each other. By splitting a coherent beam of light into two beams which run down two separate paths and then mixing them back together at the point of measurement, one can determine the optical path difference. In a laboratory, interferometry is carried out with lenses, beam splitters, and mirrors, however with the losses associated with light being coupled into and out of optical fibres a similar setup is prohibitive and not robust enough for practical applications. Coupling of two fibres together allows all-fibre interferometers to be manufactured. Two of the most common interferometer architectures are the Michelson and Mach-Zehnder interferometers. The former—used by the eponymous scientist and Edward Morley who famously could not find the æther—is a folded-over version of the Mach-Zehnder interferometer that has two outputs for differential measurements. The optical phase shift (ϕ) between the sensing arm and the reference arm, interrogated with a narrow-band coherent source (i.e., a laser), is a function of the optical path length (nL), itself a product of the index of refraction (n) and the physical length travelled by the beam, such that

$$\phi = 2\pi nL/\lambda = nkL$$

where k is the optical wavenumber in a vacuum and λ is the wavelength. The fibres are sensitive to elongation due to strain, or changes in temperature or electric field, which may induce a change in the index of refraction. The sensing-arm fibre can be coated with specialized materials, such as magnetostrictive or piezoelectric compounds, which induce sensitivities to other phenomena (magnetic and electromechanical, respectively). In fact, a piezoelectric actuator is often applied to the reference-arm fibre to allow scanning of interference patterns over a range of optical path lengths or to ensure quadrature, as discussed below. Shi et al. demonstrated an interferometric ap-

proach to sensing magnetic field strength using magnetostrictive coating on the sensing arm and a piezoelectric actuator to ensure quadrature on the reference arm [242]. The optical path difference at a photodiode translates into a differential photocurrent of:

$$i = \epsilon I_0 \alpha \cos(\phi_d + \phi_s(t))$$

where ϵ is the photodetector responsivity, I_0 the nominal laser intensity, α is the optical loss term, and where ϕ_d and ϕ_s are the drift induced and measurand induced phase shifts, respectively. An important result is that i is periodic for large changes in ϕ_s , thus the signal is non-monotonic and the relative monotonic dynamic range is limited to small changes in phase. A near-linear response is seen over an even smaller range, which is a fraction of the wavelength of light used for interrogation. Another important limitation is the ϕ_d term, which can be seen as a bias and can drift with temperature or polarization effects, changing the quiescent point and thus the sensitivity, linearity, or even monotonicity of the signal. This is known as the fading problem. For optimal signal detection, ϕ_d must be maintained at a quadrature point, where $\cos \phi_d = 0$. There are a number of ways to reduce the fading problem which are discussed in the literature [275], however the use of white light interferometry, or more properly low-coherence interferometry, has many advantages for absolute measurements (such as displacement and force) [42, 52, 212]. These techniques are insensitive to fibre losses and emitter power. They can make use of a less expensive LED rather than a laser, though they usually require a mechanically scanned reflector to perform the measurements, which increases system complexity and the sampling period. Inaudi et al. discussed a deformation sensor utilizing low-coherence interferometry that has 10 μm resolution over a large range of 60 mm [122].

The Fabry-Pérot interferometer (FPI) can be fabricated with fibre optic cables in a variety of ways: intrinsically by splicing in a semi-reflective etalon, extrinsically by separating two fibres by an air gap, or by mounting a reflector onto the end of a fibre with a small air gap¹⁰ in between [95, 275]. The FPI produces a periodic attenuation of the spectrum with a complex nonlinear transfer function, and thus provides only a relative measurement for small phase shifts that are highly dependent on the bias point. While measurements of temperature or strain are achievable, a simple pressure

¹⁰The air-fibre interface in these sensors acts as a semi-reflective mirror.

sensor can be fabricated by mounting a reflective diaphragm to the end of a fibre [1]. Due to the high-sensitivity and small range of displacement of the diaphragm, this type of sensor is used to measure acoustic signals in optical microphones and hydrophones. Microfabrication of optoelectronic devices has led to the use of the FPI in accelerometers [237].

The Sagnac effect is harnessed in another form of interferometry utilized mainly in optical gyroscopes. Consider a loop of fibre optic cable supporting two counter-propagating coherent beams of light emanating from the same point—assume the cable has a core index of refraction of 1 for simplicity. At rest both beams have the same path length determined by the length of the loop, $2\pi R$. As the loop is rotated in a clockwise fashion, with a rotation rate of Ω , the optical path length of the clockwise-going beam is increased by $\Omega 2\pi R^2/c$, and the optical path length of the counter-propagating beam is decreased by the same amount. The total optical path difference is then $4\Omega\pi R^2/c$, resulting in a wavelength-dependent phase shift, meaning a periodic shift in interference peaks in the spectrum in relation to the rotation rate. In addition to measuring rate of rotation, the Sagnac effect can be used to measure many high-frequency measurands, specifically strain and vibration. Udd listed in 1991 [275] some of the possible uses of Sagnac interferometry in acoustic sensor arrays for submarine warfare, listening devices for espionage, all-optical microphones, as well as optical fibre strings for optoelectronic guitars.

Intrinsic interference of modes in a multimode fibre interrogated with a coherent source produces a spatially-arranged random speckle pattern on the (rough) end-face of the fibre [94]. Coupling the output of a multimode fibre to a singlemode fibre effectively spatially-filters a single speckle. The amplitude of a speckle has the same transfer function as a two-beam interferometer, along with the same quadrature limitations. Thus vibrational information can be extracted, though low frequency random fluctuations in temperature, source intensity, and polarization limit its response around DC. An electro-optic guitar was developed which makes use of this spatially-filtered speckle pattern [74].

Overall, interferometric sensors are best suited to measuring small displacements and rate of rotation. End-coupled fibre FPIs are inexpensive and simple to make. Multiple FPI sensors can be fabricated on a single fibre, and can be multiplexed to provide quasi-distributed sensing. Properly designed interferometric sensors exhibit high sen-

sitivity (if perhaps a low dynamic range), though cross-sensitivity to temperature and other factors (such as random polarization effects) limit their use. While some purely optical feedback techniques have been shown to ensure the sensor is maintained in an optimal quadrature state, most methods employ electro-mechanical actuators making these systems less attractive. Fabry-Pérot etalons, as well as the Mach-Zehnder and Michelson interferometers, are commonly used to interrogate grating sensors and to achieve variable-wavelength optical filters. Interferometers are an important part of wavelength detection systems commonly employed with grating and scattering-based sensors, which are discussed in later sections.

2.2.3 Polarization Sensors

Considering a travelling plane wave, the electric field is directed perpendicular to the direction of propagation, say the z -axis. The electric field can be decomposed into components in the x -axis and y -axis, and the difference in phase and amplitude ratio determines the wave's polarization state and direction: if they are the same phase¹¹ they are said to be linearly polarized, with the direction dependent on the components' relative strengths; if they are 90° out of phase and of equal amplitude they are said to be circularly polarized; and otherwise they are said to be elliptically polarized. A dichroic material, such as a polarizer, transmits linearly polarized light. Polarizers are usually applied in pairs: one to ensure the light source is first linearly polarized before it is used to interrogate a sensor, and the other polarizer is used to analyze the resultant modulated beam before the detector. A simple rotational position sensor would use the aforementioned setup, with the analyzing polarizer coupled to the measurand. This results in an amplitude-dependent angle measurement with the signal decreasing as the polarizer axes are rotated out of alignment. The other materials used in polarization sensors are birefringent crystals, especially those which have photoelastic properties. Birefringence is the property of a material to retard the phase of one polarization direction more than another—in other words, an anisotropic index of refraction with respect to polarization. Photoelastic materials exhibit birefringence when strained, making them candidates for force and pressure sensors. Optical retarders (or wave plates) are birefringent materials, specially-aligned to exhibit orthogonal fast and

¹¹or completely out of phase, which is the same as an inversion of amplitude.

slow axes of polarization. The relative phase shift imparted through birefringence is highly wavelength-dependent, allowing spectrally-informed interrogation. Along with polarizers, wave plates can be used to improve sensor quality by minimizing error due to environmental effects and emitter power fluctuation [275]. Birefringence in optical fibres is both a limitation (thus the need for polarization-maintaining fibres) and a benefit where high-birefringent fibres are used as sensing elements [50, 150].

Most applicable to MRI-compatible interfaces, is the Faraday effect: a magnetic-field dependent rotation of the axis of polarization in a circularly-birefringent medium. A MRI-compatible Faraday rotator was implemented by Bock et al. that was capable of measuring magnetic field strength as a function of position and orientation (i.e., angle with respect to the main B_0 field) [15]. The authors note that the unstable output power of the source laser caused errors, which could be reduced by using spectral modulation techniques [275]. An improvement to the design by doubling the sensitivity of the device, as they note, could have been made by taking into account the non-reciprocity of the Faraday effect and using a reflector to pass the light back through the fibre, doubling the rotation angle of polarization.

Measurements of magnetic field strength in MRI-compatible interfaces could provide measures of position and acceleration inside the scanner bore, which could prove useful for measuring bowing gestures. Although pulsed gradient fields would disrupt the sensor readings of position¹², this type of sensor is well adapted for use with sparse functional scanning sequences already employed in auditory and musical tasks [41, 79, 295]. While the static magnetic field and the gradient components are well defined within the centre of the MRI scanner bore (a sphere with a diameter of about 24 cm), it is highly nonlinear outside this region [232].

2.2.4 Grating-Based Sensors

Distributed or quasi-distributed pressure measurements are useful in MRI-compatible interfaces, especially for fingerboard-based musical instruments where individual fin-

¹²The gradient fields could actually be used to determine position absolutely without mechanical constraints by synchronizing the sensor measurements with the scanner. The imaging sequence however would likely inhibit the use of the gradient fields for this application, at least where high-temporal resolution is required. As well, the switching of gradient fields creates audible noise, which is ideally minimized during musical tasks, thus discouraging this application.

ger pressures and contact points are of interest. Intrinsic grating sensors can provide high-sensitivity strain measurements and can be multiplexed to create quasi-distributed sensor arrays. Fibre Bragg gratings (FBG) and long period gratings (LPG) are two such sensing technologies that are discussed in this section. A grating is inscribed in germanium-doped silica fibre by exposing it to UV light using an interference pattern [111, 125, 134].

Fibre Bragg Gratings

By inscribing a periodic modulation of the index of refraction in an optical fibre, a notch in the transmitted spectrum is created. The wavelength corresponding to the reflected resonant peak is known as the Bragg wavelength, given by:

$$\lambda_B = 2n\Lambda$$

where n is the index of refraction, and Λ is the grating pitch (thus the grating period is about a third of the wavelength, given a nominal $n = 1.5$). The bandwidth of the Bragg signal is related to the length of the grating, which varies greatly from sub-nanometre to tens of centimetres and beyond, depending on the application [134]. The Bragg wavelength is sensitive to strain, both axial and transverse, through the modulation of the grating pitch and by photoelastic modulation of the index of refraction [111]. Transverse strain leads to birefringence, and thus bifurcation of the resonant peak [80]. The index of refraction's sensitivity to temperature, as well as the elongation due to thermal expansion enable FBGs to be used as temperature-sensitive devices, however, care must be taken to differentiate the effects of strain and temperature when measuring the Bragg shift. As the FBG response is represented by the spectrum, an absolute reference independent of source intensity or coupling losses is ensured. Using interferometric techniques to measure dynamic strain (as opposed to the slower scanning methods used for static strain), the Bragg wavelength shift is converted to a wavelength-dependent phase shift:

$$\Delta\phi = \Delta\lambda_B(2\pi nd)/\lambda_B^2$$

where nd is the interferometer's optical path difference. The path difference, however, must be less than the coherence length of the Bragg signal, otherwise the interferometric technique approaches that of the low-coherence case. In order to discriminate between strain and temperature-induced Bragg shifts, two co-located gratings may be employed. In one arrangement, two identical FBGs are arranged such that they both undergo the same temperature shift, but only one is mechanically-coupled to the externally-applied strain. In a second arrangement, two gratings (be they both FBGs or one LPG) having differential and separable changes in both measurands allows both the temperature and strain to be sensed independently after calibration. As FBGs reflect only a narrow band of the spectrum whose shift is also relatively small (on the order of hundreds of kilohertz), several FBGs inscribed along a single fibre can be engineered to have spaced spectral channels, which can then be interrogated using wavelength-division multiplexing. Time-division multiplexing can also be employed, though for closely spaced gratings of around a centimetre apart, for instance, the difference in reflection times are on the order of 30 ps—too quick for most photodetectors, let alone digitization equipment. Excitingly, FBGs have been manufactured out of plastic optical fibres, showing sensitivity to humidity, temperature, and strain [104].

An augmentation to the regular FBG is the chirped grating. By creating a modulation in the grating period or effective index of refraction (most easily accomplished by tapering the fibre core diameter), a wider variety of Bragg spectral responses can be produced. A smoothly varying broad-spectrum Bragg reflection can be created where the amplitude response of a single interrogation wavelength can vary linearly over a large range of strain. In addition, the temperature and strain characteristics of a chirped grating are separable in certain configurations. For example, a change in strain may broaden and shift the Bragg resonance, whereas a change in temperature may only affect the spectral centroid. Longer chirped Bragg gratings, on the order of centimetres for instance, can provide intra-grating distributed measures of strain and temperature. Local changes to the grating properties of a broadband chirped fibre induce deviations in the spectrum from its initial characterization. That is, for a fibre with a nominally flat Bragg response over a certain range perturbed in the centre, a shift would be seen in the centre of the spectrum, attenuating the reflection of the median wavelength and increasing the amplitude of slightly larger wavelengths. In order to implement a practical interrogation system without a spectrum analyzer or wavelength-scanning system,

the amplitude and/or the phase response (using a mismatched interferometer) can be measured with a coherent narrow-band source.

Chirped gratings have several advantages over their milquetoast brethren: the ability to engineer more nuanced Bragg spectra, separable strain and temperature response, intra-grating distributed sensing, and interrogation using the amplitude and/or phase response instead of the whole spectrum, among other benefits. Limitations include more complex interrogation techniques to extract a more detailed measurand profile and the general need for well-characterized sensor response. They are also more difficult and more costly to manufacture. In fingerboard instruments, vibrato sensing or other small tuning-sensitive measurements emerge as useful behavioural measures made available through the use of chirped fibre grating sensors.

Long Period Gratings

FBGs and LPGs allow many of the same measurements to be made (i.e., sensitivity to strain and temperature), however, LPGs have a more complicated spectral response that requires more careful design and interrogation. While FBGs are inscribed with a sub-micron grating pitch, the grating period of LPGs typically range from $100\ \mu\text{m}$ to $1\ \text{mm}$. This has the effect of coupling propagating modes of certain frequencies into the cladding, instead of reflecting them as with FBGs. LPGs can be written into fibres to have specific sensitivities to bending and index of refraction, in addition to positive or negative (or nulled) temperature coefficients. The more numerous and broader spectral resonances make LPGs more difficult to multiplex. With only a single grating and well-characterized interrogation, various measurands can be sensed simultaneously and separated accordingly [125].

2.2.5 Scattering and Luminescence-Based Sensors

Methods for distributed sensing that incorporate completely intrinsic or all fibre solutions include optical scattering techniques. Optical time-domain reflectometry (OTDR), the main interrogation technique for backscatter sensing, uses a very short laser pulse and a fast-acting photodetector to measure the reflected power of certain wavelengths of interest and the time-of-flight (i.e., propagation delay). The backscatter power-time curve of an unperturbed fibre decays with a known slope. A perturbation of the mea-

surand at some point along the sensor translates into a deviation of this curve, where the position of the measurand mediates the time-of-flight of the discrepancy. There are three main scattering effects seen in fibre optic cables: Rayleigh, Raman, and Brillouin.

Rayleigh scattering is an elastic process where small random fluctuations in the index of refraction cause some of the incident light launched into the fibre to backscatter toward the source. The reflected light is of the same wavelength as the incident beam, and its intensity drops off as a function of distance from the source. The Rayleigh scattering intensity is a function of stress, temperature, external refractive index, and fibre defect/fault. Using polarization-optical time-domain reflectometry (POTDR), birefringence of the fibre can be measured continuously along the fibre, leading to measurements of anisotropic pressure, electric, and magnetic fields. Spatial resolution of around 1–10 *m* is common.

Spontaneous Brillouin scattering exploits nonlinear effects of optical fibres when interrogated with an intense laser pulse. The interaction of energetic photons within the optical medium results in the ejection of photons of frequencies which differ (very slightly and contained within a narrow frequency band) from the incident laser beam as a function of temperature and strain. The power of the Brillouin signal is also, yet differentially, dependent on strain and temperature, thus by using measurements of both power and Brillouin frequency shift, strain can be decoupled from temperature [198]. Additional compensation using Raman and Rayleigh scattering has also been employed [18]. Unfortunately it is difficult to find devices in the literature which can obtain sub-decimetres spatial resolution.

Raman scattering is similar to Brillouin scattering in that the reflected photons are created through Stokes and anti-Stokes processes, representing photons ejected with decreased and increased frequency (energy) compared to the incident light, respectively. While mechanical interaction of phonons within the crystal lattice results in Brillouin scattering, Raman scattering is caused by molecular-level vibrations, thus the use of Raman scattering in spectroscopy. The Raman-backscatter signal is a broadband signal and extremely weak, about 60 dB less powerful than Rayleigh scattering, though sufficiently spaced in frequency to distinguish without interferometry (which is not always the case with Brillouin scattering). Raman scattering, however, provides only a measure of temperature by taking the ratio of Stokes and anti-Stokes powers, and is insensitive to strain. Again, spatial resolutions of fibre Raman scattering temperature

sensors is usually around 1 m with temperature resolution on the order of tenths of degrees centigrade, over a range of several kilometres.

Back-scattering is a highly sensitive distributed technique for measuring strain and temperature, among other measurands. It is hampered by low spatial resolution and very weak signals which require integration over many measurements, thus increasing acquisition time. While measurand cross-sensitivities can be resolved by measuring various back-scatter phenomena at once, very fast and highly-sensitive optics are required in concert with specialized processors. There are several advanced techniques, compared to OTDR and POTDR, employed to interrogate distributed fibre sensors using multiple lasers or by using surface enhanced and stimulated scattering processes as opposed to the spontaneous processes discussed above [203]. While these augmented techniques provide more accurate and stable results, spatial resolution is still not sufficient for most applications in musical instruments.

Using a distributed high-sensitivity temperature sensor, with sufficiently high temporal and spatial resolution however, could be a useful measure of human-interface interaction if the specific heat capacity is sufficiently low, allowing for fast temperature fluctuations to be sensed. As of yet, these techniques, have not matured to this state [95, 134]. In contrast single-point, quasi-distributed, and distributed fluorescence sensors formed in specially-doped fibres have shown promise in recent years for temperature and strain measurement [93, 138, 225]. Most recently Dalzell, Han, and Ruddock were able to sense strain and temperature with spatial resolution as high as 2 cm using two time-correlated counter-propagating beams [51]. Two-photon excited fluorescence is created at any point along the fibre determined by the variable time delay between the launching of the two sub-threshold counter-propagating pulses, which superimpose at a point to exceed the fluorescence threshold and emit easily distinguished lower-energy photons. This technique is limited in response time by the fluorescence lifetime of the dopant as well as the time needed to scan the whole length of the fibre, thus there is a trade-off between temporal and spatial resolution.

2.2.6 Discussion

Optical fibre sensors offer high sensitivity, immunity to electromagnetic interference, and function remotely without ferromagnetic materials, making them ideal candidates

for MRI-compatible sensing applications. MRI-compatible input devices may require a variety of measurands to be sensed including position, force, acceleration, vibration, and pressure. For useful real-time interaction, the time taken to perform both the signal acquisition and signal processing must be relatively quick (e.g., on the order of milliseconds). After a review of the literature on advanced fibre optic sensors presented here and elsewhere, there are a few techniques which show a lot of promise, namely: Fabry-Pérot interferometers for small acoustic vibrations, the fibre Sagnac interferometer gyroscope, magnetic-field sensitive Faraday rotators for position and acceleration sensing, in-fibre Bragg and long period gratings for quasi-distributed temperature-compensated force measurements, and two-photon excited fluorescence for distributed temperature and strain measurements. While these techniques offer increased measurand sensitivity over intensity-based methods, the system cost and complexity also increases to ensure cross-sensitivities are accounted for and noise sources are mitigated.

Chapter 3

Embedded Systems Design

This chapter presents the development of an optoelectronic acquisition system through several iterations. The embedded systems¹ presented here formed the backbone of the varied musical interfaces and instruments presented in the following sections. It played a dual role as an acquisition system for behavioural neuroscience and as a performance tool to interface with optically-sensed electronic musical instruments. Such a comprehensive and specialized system was, hitherto, unavailable.

As a system for acquisition of subject behaviour during fMRI studies, as shown in Fig. 3.1, the electronics were relegated to the control room and optical fibres were used to convey light from the acquisition system into the scanner wherein the subject lay. LED and laser light sources were used to interrogate the optical sensors that were embedded within the musical interfaces. A subject's interactions with an interface were sensed as modulations in light intensity and were relayed back to the acquisition system in the control room via optical fibres. The acquisition system then converted the light energy to electrical energy, amplified and digitized the signals, and sent them along to other systems for logging and feedback generation.

There was no single interface for which the system was designed, but rather a set of intersecting applications for which such a system was required. That is to say, subject response systems for behavioural neuroscience studies have an overlapping set of characteristics shared with digital musical instruments (DMI): low-latency response (e.g., a range of 1 *ms* to 10 *ms*), some mixture of analog and digital inputs and outputs, the

¹An embedded systems is an electronic and computerized assembly that acts to perform a set of specialized tasks, often as a stand-alone device.

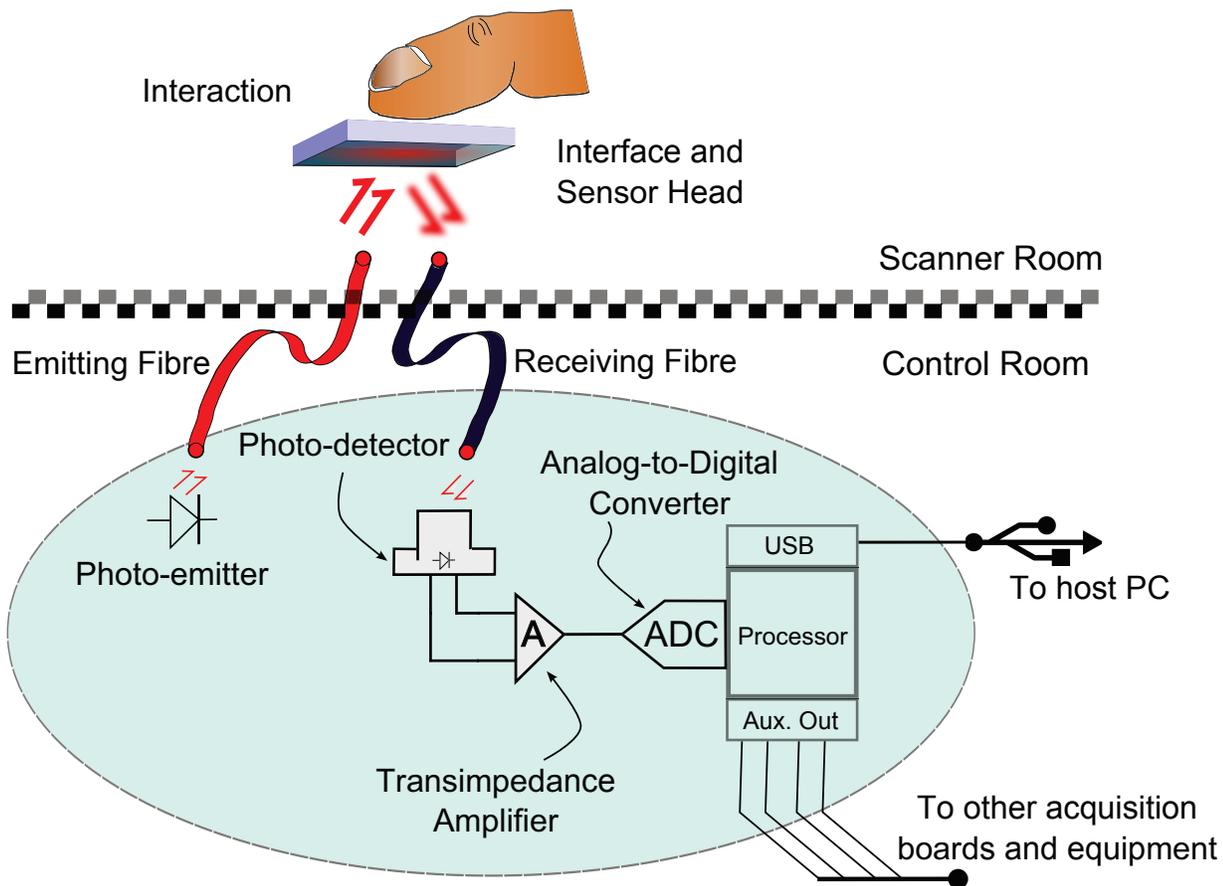


Figure 3.1: Block diagram of optoelectronic acquisition system.

need for a stand-alone system that is easy and quick to set up, the flexibility of an interoperable system that can communicate with legacy or general-purpose equipment, and furthermore a high-speed, high-resolution system that can present, capture, log, and transmit data synchronously with other systems.

There are, of course, many caveats and exceptions: a neuroscience study looking at changes in inter-onset-interval response timing may only be concerned with temporal resolution and latency variability. If feedback is not being presented, a fixed delay would not impact the difference measures of two similarly-spaced events. On the other hand, if a DMI's analog sensor (e.g., a linear potentiometer) were used to coarsely select the intensity of a processing effect between movements,² an audio-quality acquisition of the sensor's signal would not be necessary.

²Here, *movement* is meant as a delineated section within a musical composition.

Unsurprisingly, the confluence of musical interfaces and behavioural neuroscience equipment has similarly matched constraints, made only more specialized by the introduction of musical interfaces for neuroscience studies into the environment of the MRI scanner. The specialization, mainly of optical sensors and optoelectronic components to obviate EMI, does not preclude the use of such a system with DMIs, nor with conventional manipulanda and response boxes for behavioural neuroscience.

Musicians are said to be pickier about their gear than are psychologists, which is why psychologists want to study how musicians are so discerning and adapt so well to perturbations in what they hear, see, and feel. It is the psychologists, however, that want to quantify these perturbations and adaptations. Thus, we set out to design a system for live music performance, which met the more rigorous and exacting specifications to satisfy the musician, with the ability to acquire and analyze the components of musical gestures, which enabled novel neuroscience studies. The system required the flexibility and extensibility to interface with a variety of tangible controllers and sensorized instruments. For example, the system needed to accommodate interfaces with sensors of different types and quantities. While we strived for a stand-alone system, communication with musical equipment using MIDI, and with general purpose computers over USB, was a positive aspect.

While we endeavoured to provide on-board sound synthesis capabilities, studio-quality synthesis was not considered as a fundamental requirement for several reasons: many excellent commercial solutions exist as stand-alone or add-on components; DMIs have the intrinsic ability to decouple the control from the mapping and synthesis; many musicians and composers use preferred equipment to which they have grown accustomed; and neuroscientists often make use of simple synthesis, pre-recorded sounds, or other commercial synthesis equipment depending on the study. Demonstrating on-board sound synthesis as a *proof of concept* with limited hardware specifications engenders the idea of high-quality audio production on the inevitably more powerful, next-generation low-power processors.

One of the major difficulties in the implementation and use of this type of acquisition system is to enable the neuroscientist access to the system: to load whatever protocols or experiment-specific task parameters, to allow the use of pre-existing software for stimuli presentation or data processing, and to monitor the state of the system and progress of the subject. As embedded systems occupy more of our spaces, with or

without our conscious invitation, their proliferation provides ever faster network-ready distributed processors. With the capacity to off-load processing, data logging and visualization, as well as user-interface and control, the stand-alone ideal—of one device doing everything—can be subverted by remotely attaching to the specialized embedded systems through the use of general purpose devices. That is, the embedded system need perform only its time-critical specialized tasks while smart-phones, tablets, and ultra-portable computers can be used to interact with the acquisition system and the data it produces at a higher level.

3.1 Preliminary Efforts

The hardware and firmware, presented in [113, 116] as part of the initial MRI-compatible keyboard project, was a comparator-based design with manually-set thresholds and MIDI output. The extensibility of this design to allow additional input channels was in the form of daisy-chained microcontrollers over universal asynchronous receiver/transmitter (UART). The 8-bit microcontroller was a Cypress CY8C27443 programmable system-on-chip (PSoC).³ This first prototype was hand-soldered on a pad-per-hole circuit board, as shown in Fig. 3.2. This device was utilized in a neuroimaging experiment conducted by Steele and Penhune and presented in [253], with a fibre optic response box in the form factor of a computer mouse presented in [113]. In addition to the neuroscience applications, the design was used by Fabrice Marandola as part of a percussion performance, using a set of Roland V-Drums as input.

The hardware and firmware were first expanded with a serial port emulation over USB (using an FTDI FT232R, UART-USB converter IC), a more compact messaging protocol, and the addition of a time-stamp to all messages. These changes were made to improve the data throughput and to integrate the acquisition system with the experimental design software platform, *Presentation*, made by Neurobehavioural Systems Inc., which did not support the MIDI protocol. Thus, the acquisition system was connected, via USB, to the host computer running *Presentation*, which generated the auditory feedback. An issue of message corruption arose where certain serial symbols were aliased

³The Cypress PSoC family of devices has evolved, since this effort, into four classes of devices, each with a successively more powerful microcontroller core, yet all with internal reconfigurable analog and digital blocks.

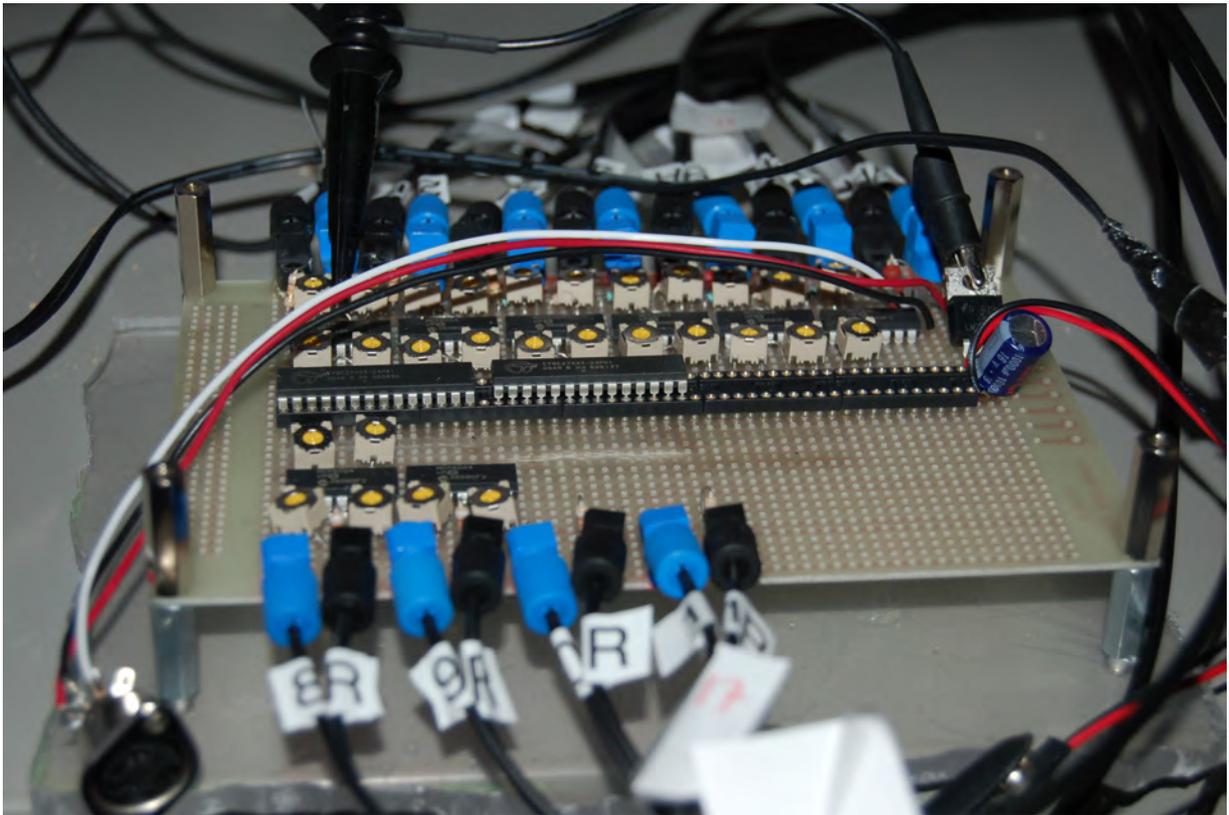


Figure 3.2: Acquisition electronics for first-generation MRI-compatible keyboard.

onto another valid, though not utilized, symbols. This issue was easily resolved by the addition of ferrite beads to the USB cables and a software routine to catch any residual symbol aliasing. The reason for the issue was due in part to a slight mismatch in clock rates exacerbated by EMI. This acquisition system was used as part of a neuroimaging experiment presented in [25], using the MRI-compatible keyboard as input [113, 116].

3.2 Hub and Node Hardware Platform

In [117], Joseph Thibodeau and I introduced the concept of a fungible interface as “not merely the portmanteau of a fun and tangible interface, but an interface with the ability to exchange one of its components with another.” This was a platform to allow the rapid prototyping of electronic musical interfaces, incorporating sensor signal acquisition, processing, mapping, and synthesis. The initial idea was to have

interchangeable and self-identifying node devices, each acquiring its own set of inputs and applying a primary layer of signal conditioning, connected to a central hub where data would be collated and further processed and sounds synthesized. We started with a bare-metal⁴ ARM7 microcontroller (LPC2468) for the hub and node devices using the Cypress CY8C29466 PSoC 1,⁵ though any devices that supported the I2C multi-master slave protocol, such as the Arduino, would suffice.

Joseph Thibodeau has continued to expand upon and apply this ideal with great success to his augmented trumpet and modular *Symbiote* platform [268, 269]. His platform has opened up, not only new extended techniques that take advantage of cognitive and motor “bandwidth” without detracting from classical techniques and an exemplary performance, but also a means to customize and adapt one’s own performative substrate in a standardized yet flexible manner.

For the divergent purpose of interrogating fibre optic sensors of interfaces discussed in Chapter 4, I built a prototype system using this hub and node architecture to form a stand-alone embedded system that included analog acquisition and digitization, signal processing, mapping, and physical modelling sound synthesis.

3.2.1 Analog Front-End

Each node was configured to modulate the light intensity of several LEDs using pulse-width modulation (PWM) while transimpedance amplifiers (TIA) and programmable-gain amplifiers (PGA) converted the multiplexed photodiode currents into an acceptable voltage range for the analog-to-digital converters (ADC). This configuration was implemented on the PSoC 1 in firmware and with a collection of external electronic components. A block diagram of the system is shown in Fig. 3.3. Analog signals were sampled with an 8-bit ADC at a run-time programmable rate that was maintained at around 100 Hz in order not to overwhelm the successive sub-systems. Digitized input signals were transmitted from the node devices over I2C to the hub.

⁴Bare-metal refers to the low-level programming approach that interfaces directly with the embedded hardware, such as clocks, registers, and other hardware blocks without an intermediary operating system.

⁵The PSoC 1 family of devices include an 8-bit M8C microcontroller core, denoted by the CY8C2XXX part number.

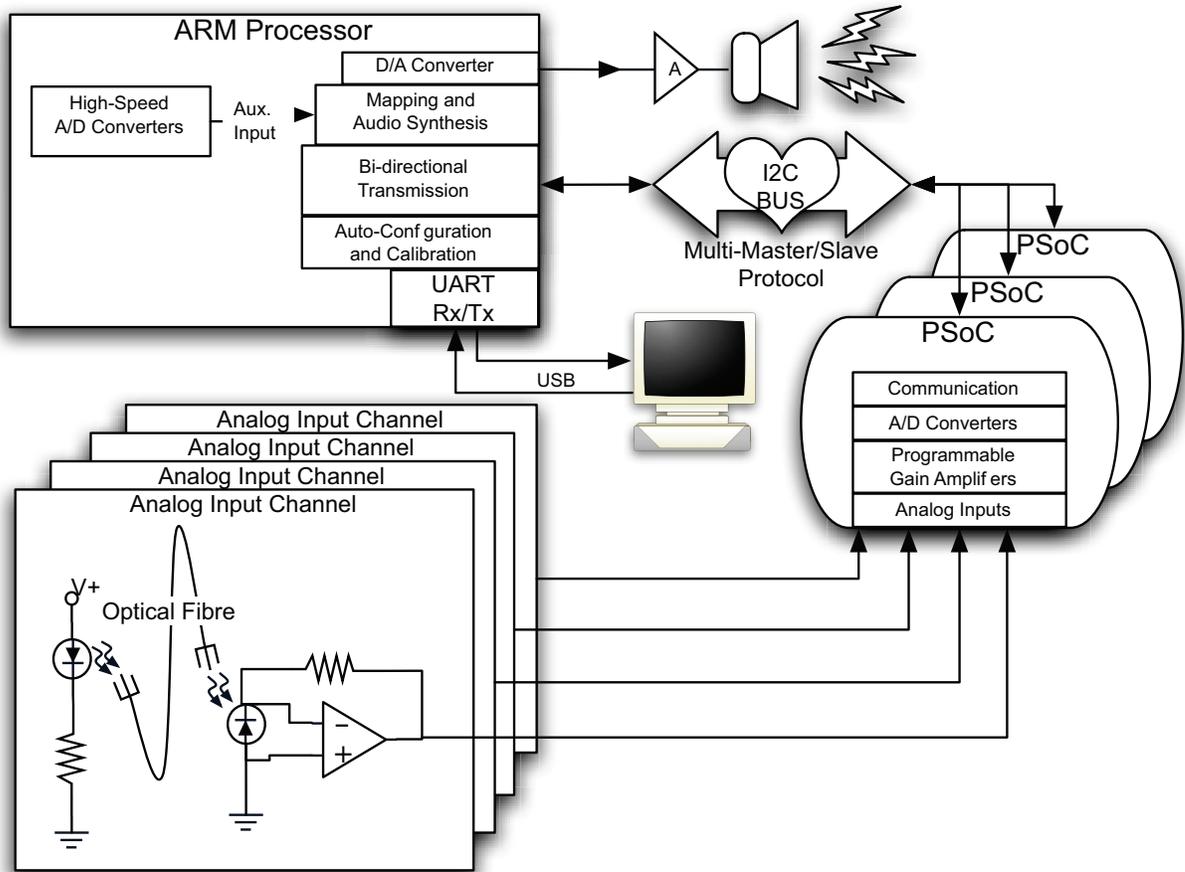


Figure 3.3: Hub and node hardware platform block diagram for optically-sensed interfaces [117].

3.2.2 Mapping

A static mapping structure with run-time parametrization was designed to constrain the processing load of the hub microcontroller, while allowing the flexibility of real-time mapping. Each source signal, represented on the hub as a floating-point value, was first scaled with a cubic polynomial function (for applying gain and possible linearization). Then each destination signal was computed as a linear combination of all processed input signals. The cubic and linear scaling parameters could be modified on-line through a command line interface to a PC (shown as the UART-USB connection in Fig. 3.3). The parameter names were human readable, and meta-data regarding the

source and destination signals, and mapping parameters could be polled, in order to ease manual, automated, or graphically-enhanced mapping design. During the prototyping phase of a DMI, for instance, the designer could experiment with mappings by remotely modifying parameters through the use of the real-time audio software MAX/MSP and/or the mapping tool, *libmapper*,⁶ designed and maintained by Joseph Malloch and Stephen Sinclair [170]. *Libmapper* is “a software library and protocol for providing network-enabled discovery and connectivity of real-time control signals” [172]. These mapping parameters were also stored in non-volatile memory and loaded upon boot or on command, allowing for *untethered* operation. The mapping loop, implemented in firmware on the hub, was executed with a run-time programmable rate coupled to the I2C transaction rate.

3.2.3 Synthesis

Decoupled from the mapping layer, a synthesis layer was developed on the hub to perform sample-by-sample or buffered audio output. The LPC2468 included a 10-bit digital-to-analog converter (DAC) with a minimum 1 μ s settling time which was connected to an amplifier and speakers for sound diffusion. Using the physical modelling synthesis algorithms described in Section 3.5 in combination with the mapping structure described above, sample rates between 5 kHz and 10 kHz were achieved with real-time control.

3.2.4 Discussion

The prototype system demonstrated a modular platform architecture for prototyping DMIs, specifically for real-time and stand-alone applications. The system integrated re-configurable node devices for sensor signal acquisition, and connected to a central hub processor with mapping and synthesis capabilities.

The LPC2468 integrated many useful features, including internally multiplied high-speed ADC, DAC, communication modules (e.g., I2C, SPI, USB On-The-Go, Ethernet, etc.), and a 72 MHz 32-bit ARM7 core. However, the overwhelming issue we faced when programming this microcontroller was a lack of clear documentation and an open-source development tool-chain. Basic operations like programming and booting

⁶The *libmapper* library is available at <http://libmapper.github.io/>

the chip were mired in incomprehensible binary files and cryptic boot files. Without proper access to the internal debugging system, debugging was performed with flashing LEDs and a serial console. Thus, it was to our delight that Cypress released a 32-bit ARM-core PSoC 5, with a unified Integrated Development Environment (IDE), compiler, and debugger; an environment similar to which I had grown accustomed with the PSoC 1.

3.3 Fungible-1

Continuing with the modular hub and node paradigm employed in the previous section, a system was implemented that was comprised of the design of a new optoelectronic acquisition board (referred to as the *fungible-1*), and a main control and communication board based on the Cypress PSoC 5 FirstTouch Starter Kit. A description of both boards and their connectivity is described in this section and in [114], as well as the application-specific configuration and firmware.

3.3.1 Optoelectronic Acquisition Board

The purpose of the optoelectronic acquisition board was to interrogate fibre optic sensors, perform basic signal processing, digitize acquired signals, and communicate data to the hub. A printed circuit board (PCB) was designed as there was no commercially-available system with the required specifications. The PCB incorporated several subsections including: power supply, analog front end, LED drivers, the programmable system-on-chip (PSoC), and a USB-to-UART integrated circuit for communication with a host PC. The board was based around a Cypress CY8C29566 PSoC 1 which consisted of an 8-bit microcontroller with integrated modular multi-function analog and digital hardware components that act like building blocks in a reconfigurable system [47]. Analog blocks come in continuous-time and switched-capacitor flavours. Continuous-time blocks include multiplexers, comparators, and PGAs, while switched-capacitor blocks may take on the role of filters, DAC, or ADC. Digital blocks are configurable as clocks or timers, PWM, and digital communication modules. See Table 3.1 for a summary of the board's features and Fig. 3.4 for a conceptual block diagram.

Table 3.1 *Fungible-1* Optoelectronic control board features.

| Feature | Description |
|-----------------|---|
| Microcontroller | 8-bit processor core, up to 24 MHz clock, with analog and digital blocks |
| Amplifiers | 8 off-chip opamps, 8 variable-gain instrumentation amplifiers, and up to 4 on-chip PGAs |
| Analog inputs | Externally-multiplexed onto 8 analog pins, internally multiplexed onto 4 analog buses |
| ADC | Many choices available, e.g., three-channel simultaneous 8-bit conversions at 3 kHz |
| Analog outputs | 4 buffered outputs, available input from analog inputs or DAC |
| LEDs | Eight buffered digital outputs, PWM drive |
| Power supplies | 5 V USB and 3.3 V (digital), 3.3 V and 1.8 V (analog) |
| Comm. | USB (UART), I2C, SPI, and others |
| Optical sensing | Photodiodes and high-intensity LEDs, ST connectors |
| Modularity | I2C for board-to-board communication on a hub and node bus |
| Reconfigurable | Run-time control of analog multiplexing and gain, on-chip analog and digital blocks configured through firmware |
| Other Features | 8 electronic potentiometers |

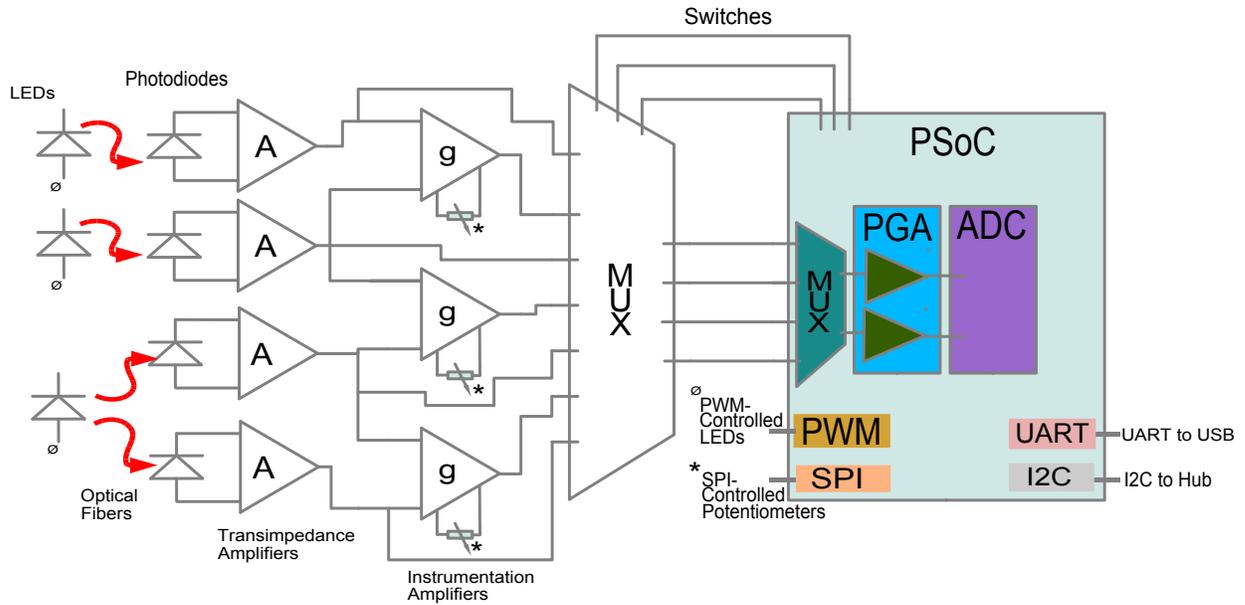


Figure 3.4: Analog front end, showing an example configuration utilizing four optical channels: three differential and four single-ended [114].

Analog Front End and LED Drivers

The analog subsection was comprised of eight opamps and eight instrumentation amplifiers. The gains of the latter were electronically-programmable. The outputs from these 16 amplifiers were multiplexed down to eight channels which connected to analog input pins of the PSoC. For optoelectronic acquisition, straight tip (ST) fibre optic housings were employed to allow for quick and robust connection to ST-terminated plastic optical fibre (POF) cables. Eight ST housings with photodiodes were connected to the opamps wired as TIAs. Photodiodes provided a high bandwidth linear current proportional to light intensity. Six different pair-wise configurations were connected to instrumentation amplifiers. This allowed for the multiplexing of eight single-ended and six differential measurements. The instrumentation amplifier gains were set by electronic potentiometers, controlled by the PSoC over a serial peripheral interface (SPI) bus. The layout of the four-layer PCB for the *fungible-1* optoelectronic acquisition board is shown in Fig. 3.5, while a photograph of the board is shown in Fig. 3.6.

Eight red (660 nm) LEDs in ST housing were selected to match the transmission spectrum of the POF [182]. MOSFET switches were used to drive the LEDs, controlled

by the PSoC. The effective LED intensity was controlled by modulating the duty cycle of a PWM signal of a sufficiently high frequency (i.e., between 100 kHz and 1 MHz). At the far end of the emitting and receiving fibres, optical measurements could be made. Four emitter followers were wired to buffer analog outputs from the PSoC to drive LEDs or actuators. Four additional PSoC pins with special connection to specific analog blocks were also brought out to a header on the board.

It should be noted that the board was fabricated with space to install additional passive components on the analog front end in order to accommodate resistive or voltage-based sensors, as well as active filtering. Two of the instrumentation amplifiers were left uncommitted and could be used for strain gauges or bio-signal electrodes (e.g., electromyogram, electroencephalogram, galvanic skin response). Fig. 3.4 shows a sample configuration with three PWM-controlled LEDs coupled with fibre optic cable to four photodiodes, providing four single-ended measurements and three differential measurements.

Internal PSoC Hardware and Firmware

Internally, the PSoC has several reconfigurable analog and digital hardware blocks in addition to the 8-bit processor core, all programmable through firmware. Thus, depending on the application, the board could be reconfigured to suit the user's needs: one configuration might require one differential analog input to be bandpass-filtered and digitized with a 14-bit ADC at 10 kHz in order to drive a single-DOF haptic interface using a DAC, while another configuration might require the programmable amplification and digitization of several slowly-varying analog signals that are then sent off-board for logging or further processing for an interactive multimedia installation. The system allowed for both of these hardware configurations to exist simultaneously or toggled at run-time.

For the acquisition of single-ended and differential optoelectronic signals, a single internal hardware configuration was implemented with a programmable analog gain and multiplexing structure to allow for a flexible trade-off between the number of acquired signals and the acquisition speed. The higher-level settings including the selection and ordering of channel acquisition, data reporting selection and speed, LED intensity, internal and external gain stage settings, and calibration could be modified

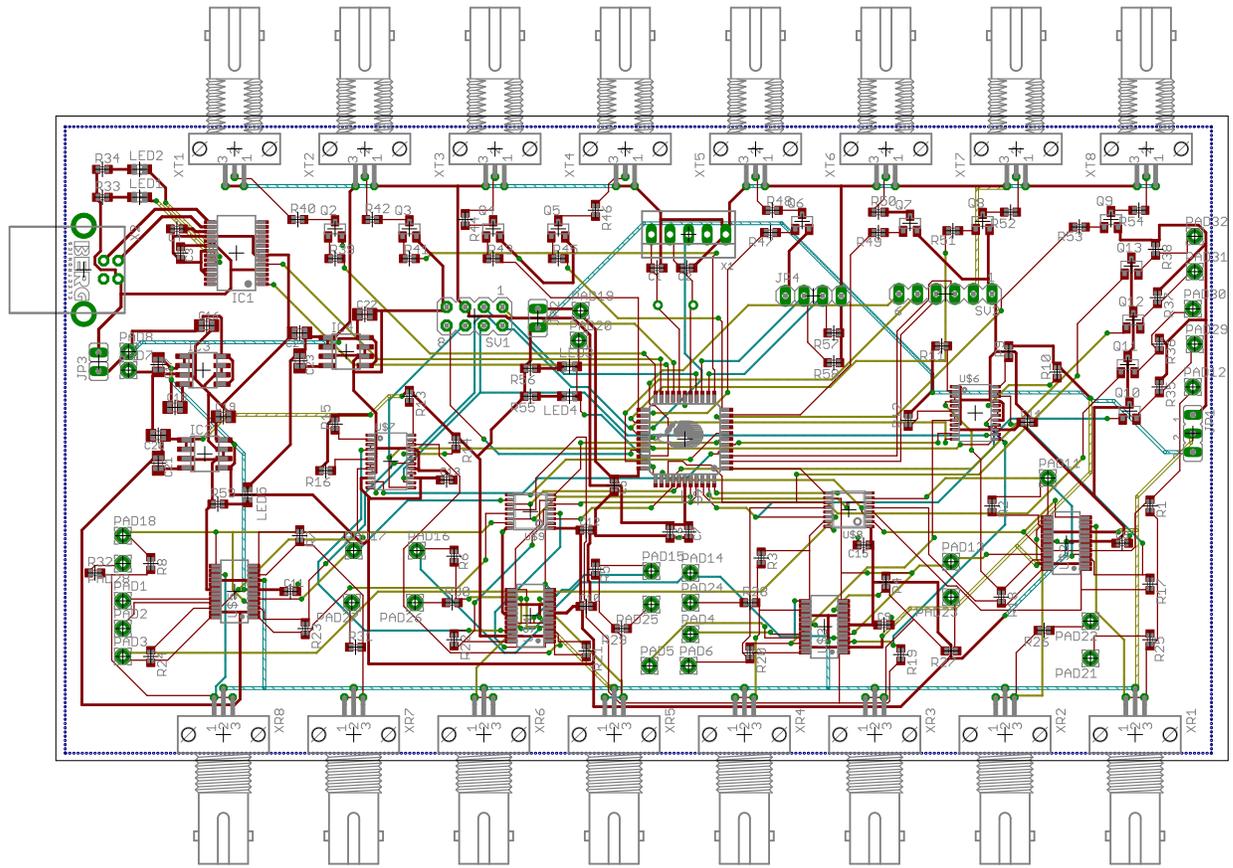


Figure 3.5: PCB layout of *fungible-1* optoelectronic acquisition board.

at run-time through commands sent to either the USB port or over I2C. Auxiliary PSoC pins were brought out to headers and reserved for additional digital communication (e.g., secondary SPI master and slave buses) and in-circuit serial programming.

The application-specific internal hardware configuration enabled the multiplexing of the eight analog inputs onto three PGAs before being synchronously digitized with a three-channel 8-bit ADC running at 3 kHz. Six of the eight LEDs were PWM-controlled, allowing for variable intensity output. An SPI master block was used to control the electronic potentiometers which set the instrumentation amplifier gains. An I2C slave block was used to send acquisition data and receive commands from a hub board, while a UART block was connected to the external UART-to-USB chip to allow for an auxiliary communications channel. Non-volatile memory (NVM) was used to store and load settings and calibration data. Additionally, a 16-bit counter block was used to

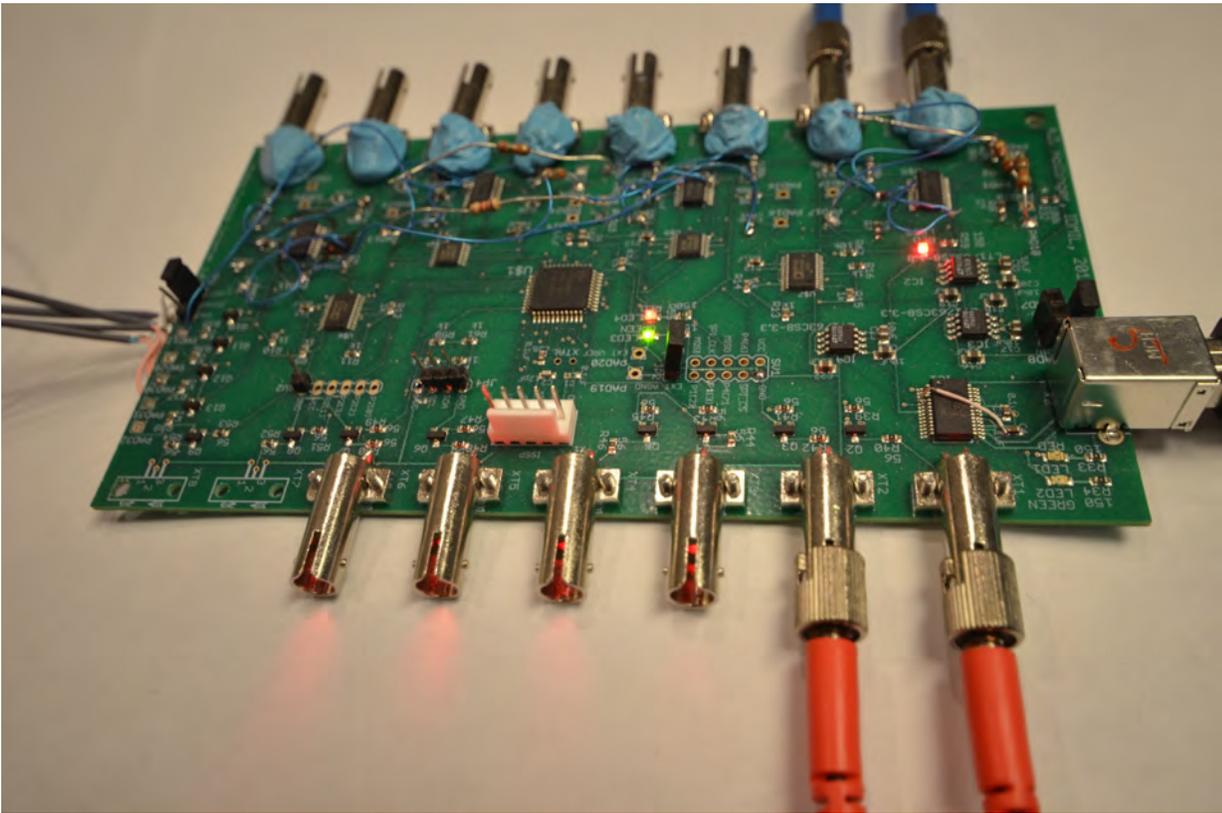


Figure 3.6: Optoelectronic acquisition and control board with fibre optic connectors.

control the timing of certain interrupt service routines (ISR).

On boot-up, the system and all its hardware blocks and registers were initialized. If valid, settings from a prior configuration were loaded from NVM. The ADC acquisition ISR, written in assembly language, stepped through the selected subset of possible input channels (referred to as an acquisition frame), setting respective multiplexer switches and gain settings before each three-channel parallel acquisition was initiated. Upon acquisition of an entire frame, the data was loaded into the I2C buffer to be read by the I2C master device. If enabled, data could be printed without polling to the USB port at a rate and in an order selectable at run-time. Otherwise data could be polled over I2C or USB.

An auto-calibration procedure was implemented which first maximized the LED outputs so as not to saturate the photodiodes, then the dynamic range of the differential channels was maximized by setting the instrumentation amplifier gains, and finally

the dynamic range of the input to the ADC was maximized by setting the PGA gains. In order to properly calibrate the analog signal chain, a run-time command was available to pair specific LEDs with their respective photodiode channels. Ambient light compensation was implemented by taking the difference between a signal captured with its respective LED on, versus with the LED off. As noted, calibration values and other settings including I2C address and LED-photodiode pairings could be saved to, and loaded from, NVM using commands sent over USB or I2C.

3.3.2 Hub Board

The purpose of the hub board was to collate acquired signals from multiple optoelectronic boards, centrally configure the nodes, provide higher-level signal processing, time-stamp incoming data from multiple sources (including MRI synchronization pulses), and act as a communication link between the interface and the host computer,⁷ if required by the application (e.g., for running the behavioural experiment software to present stimuli and log responses). The hub board was based on a 32-bit ARM Cortex M3 microcontroller, specifically the Cypress PSoC 5 FirstTouch Starter Kit. The I2C bus was used as the main communications link between the hub and the node boards. The hub board was also used to run mapping and synthesis algorithms as discussed in Sections 3.2.2 and 3.5. The development kit hardware was modified to accommodate a B-type USB connector, using Cypress' USB-UART serial emulation module library.⁸ For critical data collection during fMRI experiments, forward-error correction with Reed-Solomon (RS) encoding [215]⁹ was applied to serial data to minimize message corruption. The ADCs on-board the hub were also used in applications where high-speed, high-precision converters were required (e.g., for audio-rate, direct-to-hub acquisition). Fig. 3.7 shows a prototype optoelectronic board with four efferent fibres and eight afferent fibres connected in a differential configuration, along with an I2C

⁷The node boards were also able to communicate directly with the host computer. Although this channel was initially reserved for manual configuration and debugging, it was found to be of use for simplified configurations.

⁸The development kit shipped with the on-board USB module functionality disabled because the USB-reserved pins were soldered to the debug chip, though they were not utilized. This was circumvented to re-enable USB device functionality.

⁹Implementations for RS encoding are available, including the version that was utilized in this project, written by Phil Karn: <http://www.eccpage.com/>

connection to a hub board and a USB connection to a laptop computer.

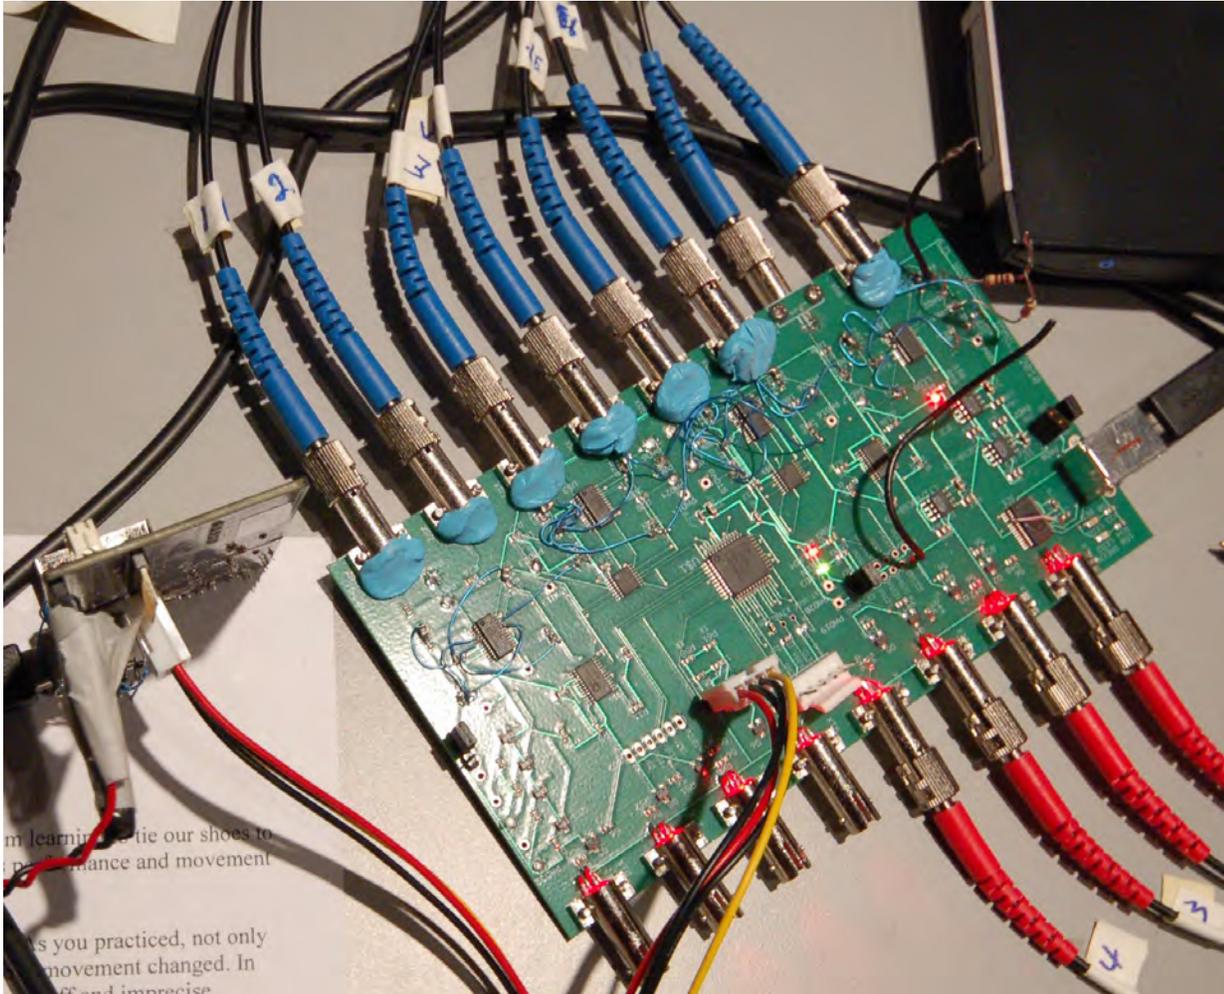


Figure 3.7: Optoelectronic acquisition board connected to fibre optic interface, hub board, and PC [114].

3.3.3 Discussion and Applications

The modular and reconfigurable optoelectronic boards were used with and without a hub board. They enabled a single hardware platform to be used with different physical interfaces and optical sensing techniques with little to no change in hardware or firmware, using run-time commands to reconfigure the platform to suit the needs of the specific application. To replace the photodiodes with resistive or voltage-based sen-

sors, they would need simply to be connected to the opamp inputs directly, rather than the TIA inputs. LEDs and small actuators could also be connected to the buffered analog outputs for visual or haptic feedback, respectively, driven with an arbitrary control signal using the DAC or directly with an amplified analog input signal.

This system was used successfully in applications that included pilot and MRI studies with the MRI-compatible keyboard, presented in Chapter 5, and with a novel musical instrument, the Ballagumi, presented in Chapter 4.

Limitations

With its 100 – 400 *kHz* clock frequency, I2C was adequate for applications with relatively low data throughput, thus SPI is suggested for higher-speed data transfer. Since I2C is slew-limited due to its open-collector signalling, it does not radiate much EMI and board layout is simpler. SPI, on the other hand, can operate at much higher frequencies (i.e., above 1 *MHz*).

The instrumentation amplifiers (part number, ISL28474) on-board the *fungible-1* were selected to achieve low-noise, high-gain amplification of differential signals. Many of the differential signals we dealt with, however, were orders of magnitude larger than could be handled by the instrumentation amplifiers, due to a fixed TIA gain. Thus the gain of the instrumentation amplifiers was reduced below the intended specifications of 100, leading to oscillations due to a large gain-bandwidth product. Thus the amplifiers were compensated with external capacitors to maintain stability at low gain settings. The lack of differential headroom also caused cross-talk between channels of the monolithic (integrated two-opamp, two-instrumentation amp) IC, requiring additional resistive padding. This was unexpected, as standard instrumentation amplifier simulation models use a different internal architecture—thus highlighting the need to simulate with manufacturer models—and amplifiers are considered to be independent, which was not the case. Although this component was not ideal for some applications, the uncompensated instrumentation amplifier would provide fine performance for differential strain gauges and bio-signal electrodes.

3.4 Fungible-2

This section describes a simpler, faster, and more integrated solution to the problem of optoelectronic acquisition and control. While the *fungible-1* system was intricately designed with state machine control of analog signal paths and run-time configurable organization of data, to end users,¹⁰ the layers of redirection might have seemed more like layers of misdirection. After meetings with colleagues at the IDMIL, Joseph Thibodeau and Joseph Malloch, we came to an agreement on the requisite features and form factor of a board for rapid prototyping of DMIs. In the end, we wanted access to most (if not all) of the features of the Cypress PSoC 5 [48] and as many pins available from the IC as possible. The form factor was to be sufficiently small to fit inside a T-Stick¹¹ [169], with a width not more than 3.4 cm. Thus, I designed the *fungible-2* from the start to be a development board for DMIs, with the dual role as optoelectronic acquisition system.¹²

The hub and node architecture was still central to the modularity and extensibility of the system. I aimed at increasing the throughput between nodes and the hub, and simplifying the analog front end, without asserting a final use case. The increased performance of the PSoC 5 over the PSoC 1 (up from an 8-bit, 24 MHz processor core, to a 32-bit, 80 MHz core,¹³ not to mention the improved clock cycles per instruction), as well as the more flexible internal analog routing, and experience with the IDE were all important factors. Maintaining I2C capabilities was still warranted, however, a shift to SPI communication between hub and nodes was required to transmit more data, more quickly.

With the advent of the Raspberry Pi single board Linux PC [276], an inexpensive

¹⁰It has always been my hope that the boards would be used by others, specifically within the Input Devices and Music Interaction Lab (IDMIL), for incorporation into stand-alone DMIs. While the neuroscientists were also end users, I did not expect them to delve into the firmware and hardware layers.

¹¹A T-Stick is a digital musical instrument conceived and designed by Joseph Malloch, having a cylindrical form and containing a variety of electronic sensors, including capacitive touch, inertial measurement, and grip pressure.

¹²Joseph Thibodeau was instrumental in putting together PCB footprints for components not within the standard and manufacturer libraries. He also brought a wealth of ideas and encouragement during conception and design review.

¹³Cypress downgraded this specification to 67 MHz, upon volume release of the chip, with several other amendments to the initial data sheet.

platform became available that would act as the hub device, and with it came the support of a large community, ample documentation, and a trove of open-source tools and software.¹⁴ What made this type of PC different from the desktop and laptop computers found in homes and offices around the world was easy access to low-level hardware: general purpose input/output (GPIO) pins, UART, I2C, and SPI.

3.4.1 Optoelectronic Acquisition Board Overview

The *fungible-2* optoelectronic acquisition and control board was designed with several subsystems and features: power regulation and distribution; optical emitters, their drivers, and optical detectors; the PSoC 5 and peripheral components (e.g., programming header, crystal oscillator); electronically-controlled potentiometers; status and power indicator LEDs; direct access to GPIO and uncommitted pins through break-out headers; a robust mini-USB connector; micro-SD card slot; and jumper-configurable signal routing. A summary of the features are given in Table 3.2.

Board layout was accomplished using 6 layers, with special attention paid to isolating analog and digital domains, in order to reduce cross-talk and noise on the analog routes, through spatial partitioning and the use of separate ground planes. The portion of the board reserved for the optical connectors, shown at the top of the PCB layout in Fig. 3.8, was connected to the rest of the board through copper traces located only on the top and bottom layers, allowing the board to be scored and broken to reduce its width for use within the T-Stick, without causing short-circuits on the internal layers. Critical paths, such as analog signal lines, power buses, and the differential USB signal pair, were given priority in routing to optimize clearance, trace thickness, and path lengths. A photograph of the populated board is shown in Fig. 3.9.

3.4.2 Power Regulation and Distribution

Separate voltage regulators were employed to isolate analog and digital domains. Two low-dropout 500 mA regulators were used to power the digital domain, one 3.3 V and one 5.0 V (Linear Technology Inc., family part number LT1763). A separate 5 V

¹⁴A large number of other, more powerful single board computers were available, but the entry cost—monetary and in set up time—were quite reduced in part to the large community and educational nature of the Raspberry Pi *movement*.

Table 3.2 *Fungible-2* Optoelectronic control board features.

| Feature | Description |
|-----------------|--|
| Microcontroller | 32-bit processor core, up to 67 MHz clock, with analog and digital blocks |
| Amplifiers | 4 on-chip PGAs, 4 comparators, and 4 opamps |
| Analog inputs | 8 dedicated pins for optical connections, other pins brought out to header |
| ADC | 2 SAR and 1 Delta-Sigma, many configurations available, e.g., multiplexed Delta-Sigma at 12-bits, 10 kHz |
| Analog outputs | Flexible analog routing to pins, 4 DACs |
| LEDs | 4 on-board ST connectors, routing for 6 |
| Power supplies | 5 V USB or external power, Separate 5 V and 3.3 V (digital), 5 V and electronically-adjustable 1.22–4.7 V(analog) |
| Comm. | Mini-USB (UART, MIDI, HID), I2C, SPI, and others |
| Optical sensing | Photodiodes and high-intensity LEDs, ST connectors |
| Modularity | SPI for board-to-board communication on a hub and node bus |
| Reconfigurable | Run-time control of internal analog multiplexing and gain, on-chip analog and digital blocks configured through firmware |
| Other Features | Micro-SD card slot, 55 pins brought out to headers, S/PDIF digital audio output, 8 electronic potentiometers |

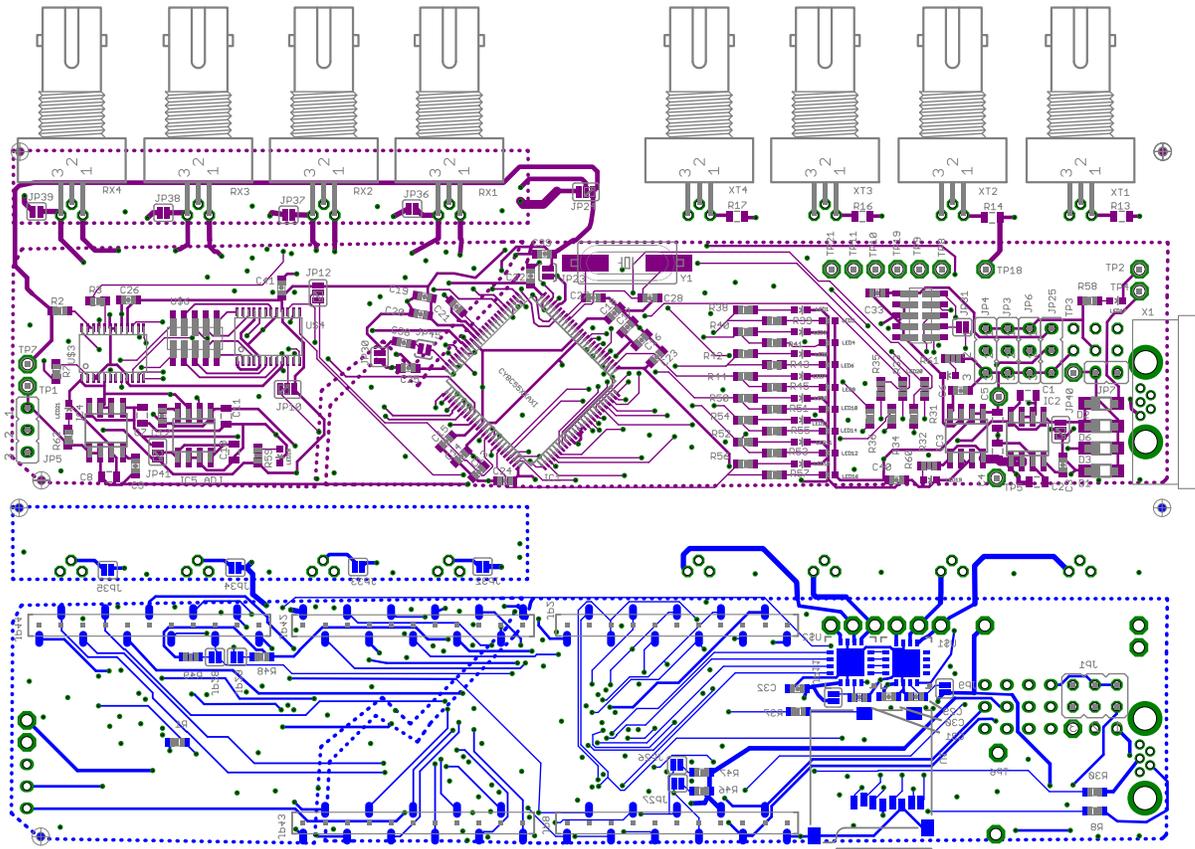


Figure 3.8: *Fungible-2* optoelectronic acquisition board PCB layout, top layer (magenta) and bottom-layer (blue) traces.

regulator and an additional adjustable-voltage regulator were reserved for the analog domain. The adjustable regulator was controlled electronically with a microcontroller-programmable electronic potentiometer, over a nominal range of 1.22 V to 4.7 V.¹⁵ The voltage sources for each of these regulators was user-selectable between off-board sources, 5 V USB bus power (i.e., VBUS), or on-board 5 V regulators for the 3.3 V and adjustable supply.

Reverse polarity protection was implemented with Schottky diodes, though the user was responsible for ensuring proper voltage selection to the PSoC sub-domains (i.e., analog, main microcontroller core, and peripheral quadrant supplies) through the use

¹⁵It was possible to increase the maximum output voltage, up to its own supply voltage (maximum 20 V) less a 300 mV dropout, however, voltage limits on other components, specifically the electronic potentiometers, would need to be considered.

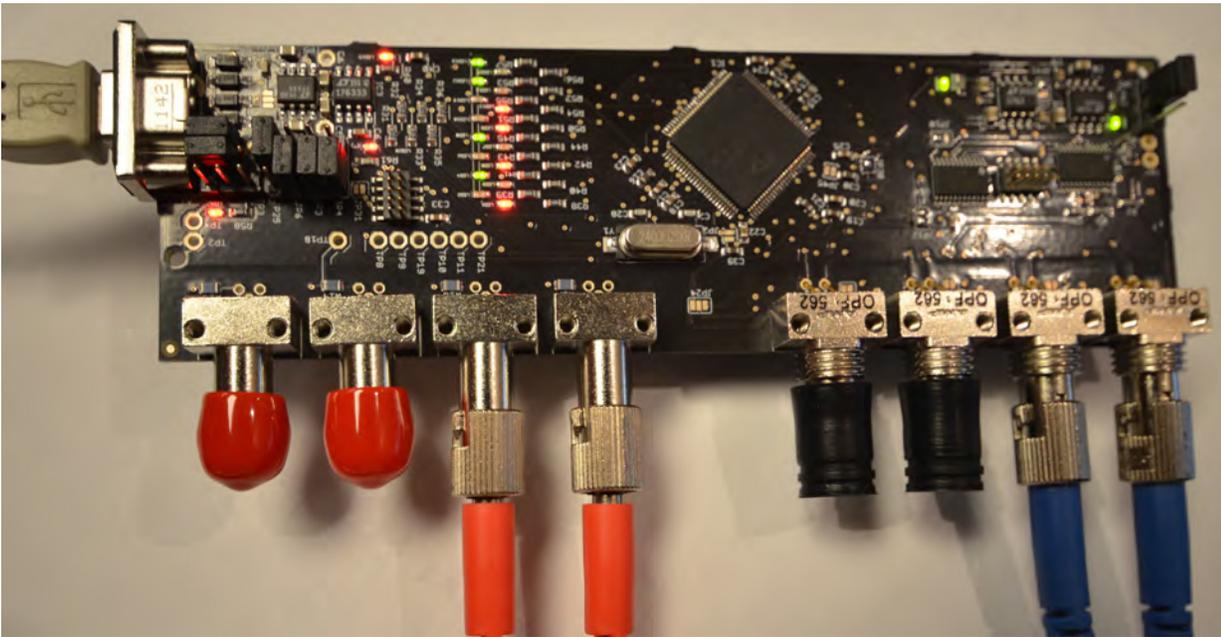


Figure 3.9: *Fungible-2* optoelectronic acquisition board.

of on-board jumpers. User-selectable voltages for each domain allowed interchangeable interfacing with peripheral devices and the upstream hub controllers. For instance, the SD card, Raspberry Pi, and off-board optical flow sensor (see Section 6.3.6) required 3.3 V, though analog components and status LEDs were usually running at 5 V. For simplified or lower-voltage operation, the analog domain could also be powered by the digital domain regulators with a slight increase in the noise floor.

3.4.3 Optical Subsystem

The optical subsystem included the emitters, emitter drivers, and detectors. The flexible routing and jumpers allowed the use of single photodiodes, common cathode quadrant photodiodes, and photodiodes with on-die TIAs. Individual photodiodes could be wired with reference to ground, to an on-board or off-board voltage reference, or with both anode and cathode floating for internal biasing within the PSoC.

Space on the board accommodated eight ST fibre optic connectors, four emitters and four detectors, however analog routing allowed for at least eight single-ended detectors and through-hole connection of six driven emitters. Red LEDs and diode

lasers ($\lambda \approx 660 \text{ nm}$), were both considered as emitters.

Two daisy-chained three-channel electronically programmable constant current LED drivers were selected (Allegro Microsystems, part number A6281) to provide a high-current range (10–150 *mA* per channel) with independent control and a minimum number of microcontroller pins. Unfortunately, a mismatch in the PCB footprint and the device pads meant that the part could not be populated. Several work-around were successfully tested: direct constant optical emitter powering; independent PWM control of external, discrete, transistor-driven emitters; ensemble-programmable voltage control using the buffered wiper output of the electronic potentiometers; and ensemble-programmable voltage control using the on-board SPI-programmable adjustable voltage regulator.

The normalized output of this final option is shown in Fig. 3.10, where 255 incremental steps of the electronic potentiometer, which in turn controlled the output voltage of the adjustable regulator, is plotted against the output voltage of the regulator that was driving an LED (through a series resistance) and the photodiode voltage (converted from a current using a TIA). The LED was coupled to the photodiode using a short length of POF terminated with ST connectors. Both voltage signals were acquired by the PSoC using on-chip ADCs while synchronously stepping the potentiometer resistance via SPI. The voltage regulator showed good linearity over its range of 1.22 – 4.7 *V*, with the photodiode registering a linear output as well—except for the turn-on region around the LED’s minimum forward voltage drop.

For another test, the adjustable voltage regulator, meant for maintaining a constant, though adjustable, stable voltage supply was adapted—by reducing the value of capacitors meant to filter out ripples under transient changes in load current—for use as a high-current (i.e., maximum 500 *mA*) voltage DAC. In addition to modulating medium-power optical emitters, this DAC could also be used to drive speakers or other electromechanical and electromagnetic actuators at frequencies well into the audio range.

3.4.4 Internal PSoC Hardware and Firmware

As the internal hardware of the PSoC 5 is reconfigurable, there are too many possible configurations to detail here. In this section the hardware configuration and related

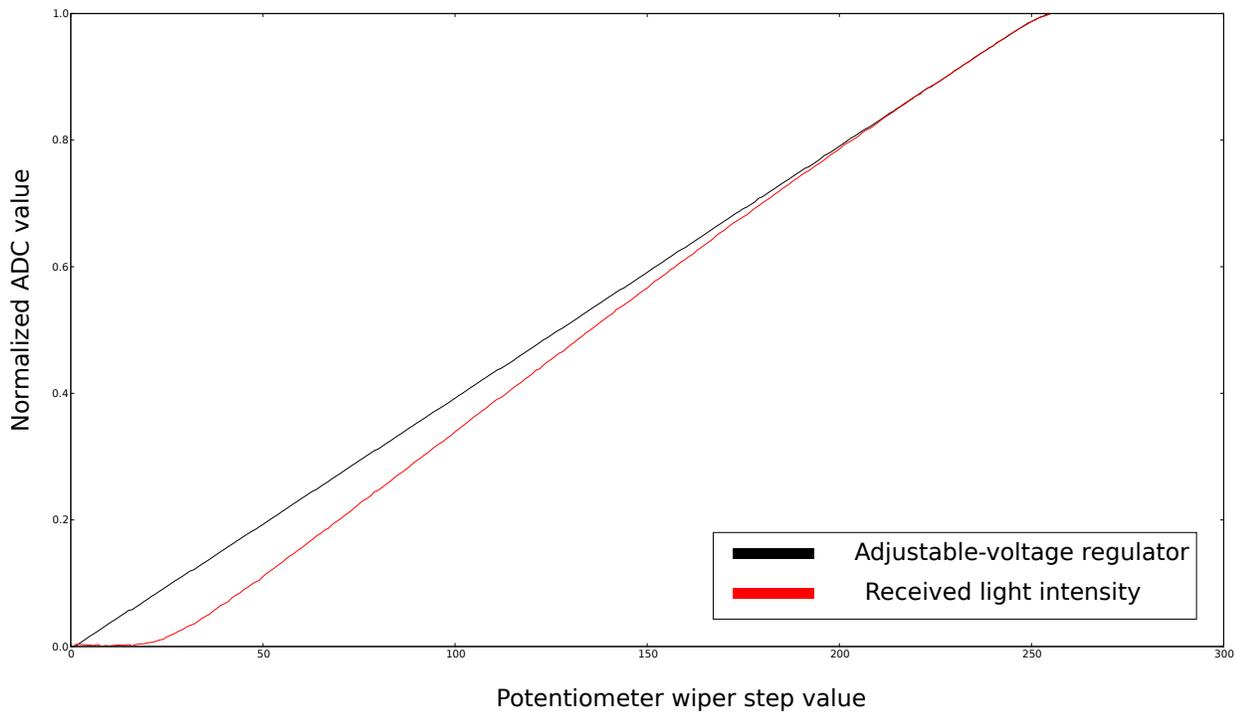


Figure 3.10: Normalized ADC and photodiode signals with LED voltage control (255 steps) using adjustable voltage regulator and SPI-controlled electronic potentiometers.

firmware routines of a single use case are briefly detailed to elucidate the system's flexibility and utility. As an aside, it is advantageous to offload certain operations, like counting of precision timing events and recurrent block memory transfer to hardware modules in order to lessen the processor load and maintain clock-synchronization. Hardware interrupts are also essential for maintaining tight synchronization and obviating the variable delays of firmware-pollled events. The example presented here makes use of both techniques.

On-Chip Analog Module Configuration

The analog hardware configuration is shown in Fig. 3.11, where four zero-biased photodiodes were connected across the inputs of their respective TIA (on the left). The gain and bandwidth of each TIA was independently configured at boot, and run-time programmable through a serial link (i.e., using either the USB emulated serial port to a PC or the SPI bus to the hub). A four-to-one analog multiplexer was synchronized

with the delta-sigma ADC using an interrupt routine, where the multiplexer would sequentially step through its four inputs and save the converted values to a ring buffer whenever the previous conversion completed. The negative input of the differential delta-sigma ADC was connected to the TIA reference voltage. The delta-sigma ADC was used as a 16-bit converter, with a programmable sampling rate between 1 *kHz* and 10 *kHz*,¹⁶ and a run-time programmable dynamic range of $\pm 0.512\text{ V}$ to $\pm 6.144\text{ V}$. Other analog components included a unity-gain opamp with its output routed to a physical pin, and with input from the multiplexed selection of an 8-bit DAC¹⁷ or the output of TIA_1. Two single-ended 12-bit successive approximation register (SAR) ADCs were operated between 0.0 and 2.048 *V* with a 55 *kHz* sampling rate¹⁸ were connected to dynamic memory access (DMA) blocks. The DMA hardware blocks, though configured in firmware, ran independently of the processor to allow direct buffering of SAR ADC samples to the S/PDIF (Sony/Philips Digital Interface Format) block shown in Fig. 3.12, the digital hardware configuration.

On-Chip Digital Module Configuration

Digital components included: a “full speed” USB block; a MIDI UART block; a FAT file system (emFile_1) for SD card reading and writing (with implicit SPI block); various clocks, counters, status and control registers, and hardware interrupts; a PWM block to control the optical emitters; SPI master blocks to communicate with the electronic potentiometers and with the off-board optical flow sensor (see Section 6.3.6); an SPI slave block to communicate with the Raspberry Pi using DMA blocks to buffer incoming and outgoing data; a S/PDIF digital audio output block with DMA-controlled buffering; and an EEPROM to store and load persistent configuration settings at run time.

¹⁶With an equivalent per channel sampling rate of 250 *Hz* to 2.5 *kHz*, over four channels.

¹⁷Suggested with a wire, though not connected in this implementation, was the use of buffered DAC output as the TIA and delta-sigma ADC reference voltage. With the removal of the wires connecting the Vref inputs of the TIA, the TIA voltage references could be asserted off-chip and independently from each other.

¹⁸ADC_SAR_1 conversion was triggered with a 10 *kHz* clock while ADC_SAR_2 was “free-running” at 55 *kHz*.

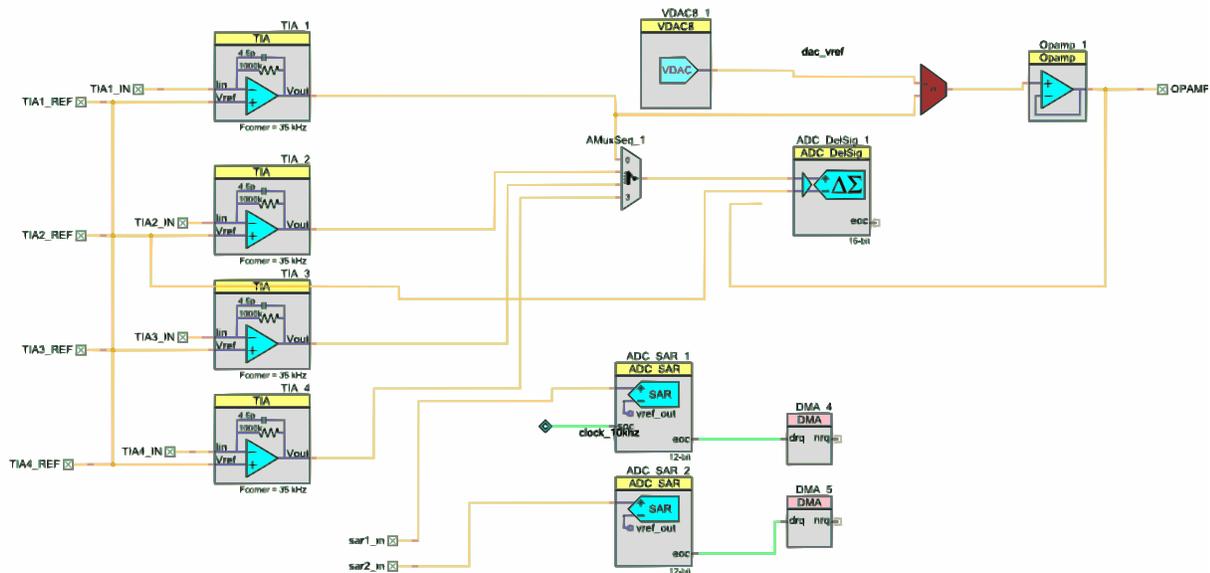


Figure 3.11: Sample layout of *fungible-2* on-chip analog hardware configuration.

Composite USB Device and Serial Port Mode

Of considerable interest to end users was the ability to switch between USB modes at run time by sending commands to the device. The three modes that were implemented were the USB serial port emulation, human input device (HID) keyboard and mouse, MIDI over USB, and USB audio. In neuroscience experiments, and even where DMI prototyping is concerned, an interface set up phase often precedes the actual testing. The command line interface that was developed to set up and query the system was available through the USB serial port emulation mode. This mode was also used for data collection when a PC was connected directly to the *fungible-2* board via USB, with acquired data printed to the serial port at a run-time selectable rate. For experiments where data streams could be replaced by simply reporting occasional messages or triggers, such as a key presses or MRI synchronization pulses, MIDI messages or computer keyboard key press messages (also known as key scan codes) are easily incorporated. Commands issued using the command line interface allowed the user to change the USB device from the serial port mode to MIDI or HID keyboard modes. Further work

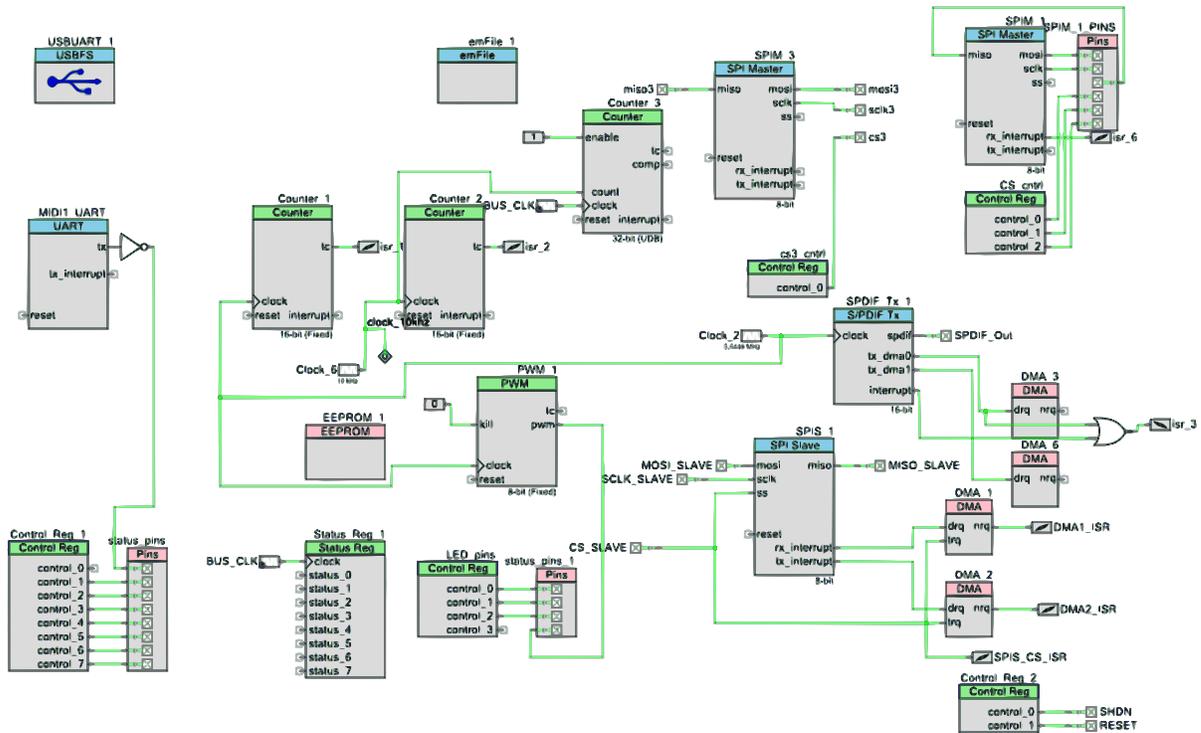


Figure 3.12: *Fungible-2* sample embedded digital hardware.

needs to be done to optionally superimpose the composite device modes (i.e., to allow the *fungible-2* to be seen by the operating system as separate serial, MIDI, and HID devices all at once).

USB MIDI Mode

The USB MIDI mode was useful where hardware or software synthesizers were already part of the experimental set up, but direct connection with a MIDI connector was not agreeable. Two on-board MIDI mappings were created: the first used a common MIDI channel number, where sequential input channels were mapped to respective note numbers (the standard MIDI use case); and a second mapping where all messages were sent with note number 60 but were spread across channels (useful for playing back a limited number of pre-recorded sound samples within a software envi-

ronments where varying the MIDI note number would otherwise automatically pitch shift the sample). The note (or channel) numbers were sequential, with a run-time programmable per-board offset. A state machine maintained the *on* or *off* state of each acquired channel, with transitions occurring when respective input signals exceeded thresholds with hysteresis. The *low* and *high* thresholds were calculated from 20% and 80% of each channel's dynamic range,¹⁹ obtained through a calibration routine asserted during the set up phase through the command line interface. In order to remain polarity-agnostic, to satisfy the connection of unknown sensors, a run-time toggle of the interpretation of *on* and *off* states with respect to *high* and *low* voltages was also implemented. In order to exit the USB MIDI mode, a specific message was sent from the host PC to the device. The firmware was easily expandable to allow the device to interpret other received MIDI messages, for example, to change parameters or actuate other functionality.

USB HID Mode

The aforementioned state machine transition reporting was extended to provide emulated key presses using the USB HID keyboard mode. Reporting was limited to note *on* events which were mapped to alpha-numeric scan codes, again with a run-time programmable offset. This allowed inputs from one board to be mapped to keys 'a' through 'd' and inputs on another board to be mapped to 'e' through 'h', for example, with an auxiliary input from the MRI scanner to send a '5', as is standard at the MNI. The drawback of the HID messaging, in contrast to the serial port, is that events are polled more slowly by the host operating system and may be received with some temporal variability [238].

As an HID keyboard is an output device, reverting it back to the USB serial port mode was not so simple. On modern keyboards, however, the *Caps Lock* state is mediated by the host operating system, not by the keyboard itself—ensuring that with two keyboards connected, they both share the same *Caps Lock* state. Often HID keyboards have LED indicators on the *Caps Lock* key, which is toggled through messages sent by the host PC to the keyboard. This upstream channel was thus used to switch the *fungible-2* from HID to serial port USB operation whenever another keyboard's *Caps*

¹⁹These were the default threshold ratios which could also be modified at run-time.

Lock key was depressed (or when the equivalent command was executed by the host PC's operating system directly).

The HID mouse and joystick functionality were tested briefly as a proof of concept. Using the optical flow sensor (see Section 6.3.6) a remotely-sensed fibre optic mouse was implemented, which allowed control of the cursor on a PC monitor. Joystick buttons and sliders were implemented to provide state and analog values, respectively. This mode enabled seamless interfacing with the operating system, and could easily be applied to adapted gaming controllers for those with disabilities [284]: on-board signal processing could be used with a somewhat conventional controller to reduce motion artifacts caused by tremors or spasticity for those with motor impairments, while disparate interaction modalities such as eye blinking or tongue movement could equally be mapped onto control messages. Haptic, visual, or auditory feedback could also be presented directly to the adapted joystick—not a new idea, but one that is feasible and easy to prototype with this technology.

USB Audio Mode

A USB audio mode was implemented on the *fungible-2* to allow each board to appear to the host PC as a standard five-channel audio interface. A 16-bit, 8 kHz per channel encoding was selected²⁰, as it was a standard lower-fidelity commercial audio rate consistent with the multiplexed capabilities of the delta-sigma ADC. Four channels were allotted to the optoelectronic inputs, while the fifth was an auxiliary input signal sampled separately with a SAR ADC to provide a means of synchronization with external equipment. This audio mode was highly effective for rapid integration of several *fungible-2* boards into a unified high-throughput acquisition system, enabling the use of pervasive audio software tools to ease the process of synchronous recording.

S/PDIF Digital Audio Output

Normally, audio interfaces (e.g., a portable audio recorder or consumer sound card) incorporate a high-pass filter to block the DC component of the input signal, which might otherwise damage the input amplifier or cause it to saturate. This limits the

²⁰In order to achieve the 8 kHz per channel sampling rate, the dynamic range and actual conversion bit-depth was reduced to between 12 and 14 bits per sample.

ability of an audio interface to record slowly-varying (i.e., < 20 Hz) signals. Many audio interfaces, however, incorporate digital audio inputs, specifically S/PDIF. A single S/PDIF module on the PSoC can output a 24-bit stereo pair at up to 192 kHz. The advantage of transmitting a signal from the *fungible-2* to an audio interface using a digital representation, is the ability to record a slowly-varying signal (down to DC) on an audio interface synchronously with other audio-rate sources, with sample-rate conversion and time stamping performed by the audio interface. This was used with several *fungible-2* boards to capture sensor signals conveying cellist gestures (see Chapter 6) while synchronously recording the audio from the cello. In contrast to directly shuttling data from the SAR ADCs to the S/PDIF module using DMA transfers, the four-channel delta-sigma ADC data taken as 12-bit values can be packed into the two 24-bit S/PDIF channels and unpacked on a PC for later analysis—time division multiplexing is another possibility to pack more data channels into the S/PDIF stream.

3.4.5 SPI Communication and Mr-Cosbbe

The Magnetic Resonance-Compatible Optical System-in-a-Box for Behavioural Experiments (MR-COSBBE²¹) was developed as a stand-alone system that would enable neuroscientists to conduct experiments using an integrated computer system, where experimental software could be loaded onto one computer that would have direct hardware access to the acquisition subsystem. This section highlights the communication protocol linking the Raspberry Pi hub and the *fungible-2* nodes, as well as some of the software tools developed for the hub. Figures 3.13 and 3.14 show photographs of the underside of *fungible-2* boards mounted in the MR-COSBBE enclosure, and a full MRI equipment setup in the scanner control room.

SPI Communication and Timing

While acquired data streams could optionally be logged to the SD card for off-line analysis, it was insufficient for on-line processing and feedback generation. The USB serial port emulation was used to communicate data streams to a host PC in real-time, how-

²¹Thanks to Joseph Thibodeau for his masterful help with this acronym. The box itself, is perhaps not MRI-compatible, but is meant for use in behavioural experiments with optically-sensed MRI-compatible interfaces.

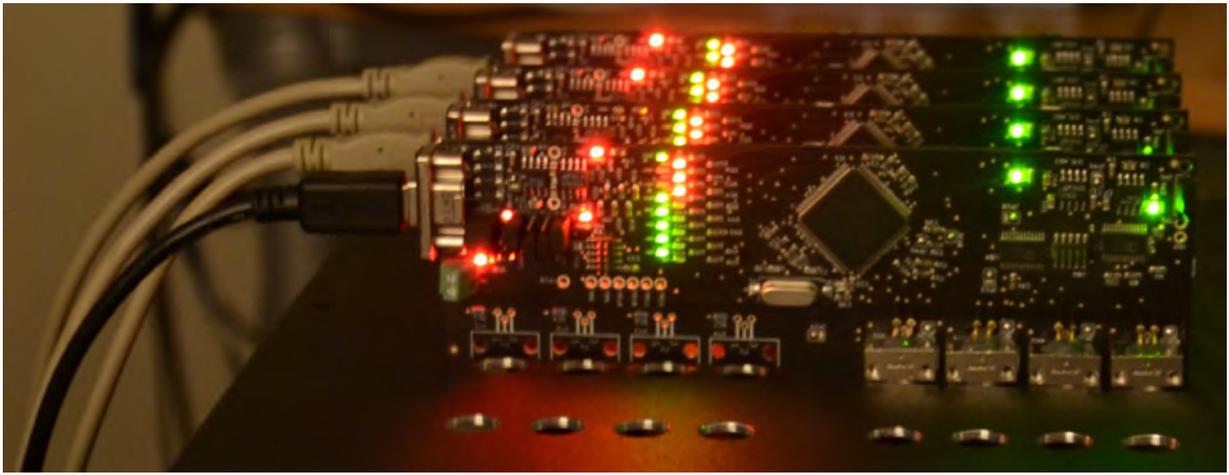


Figure 3.13: MR-COSBBE internals with several *fungible-2* optoelectronic acquisition boards mounted together.

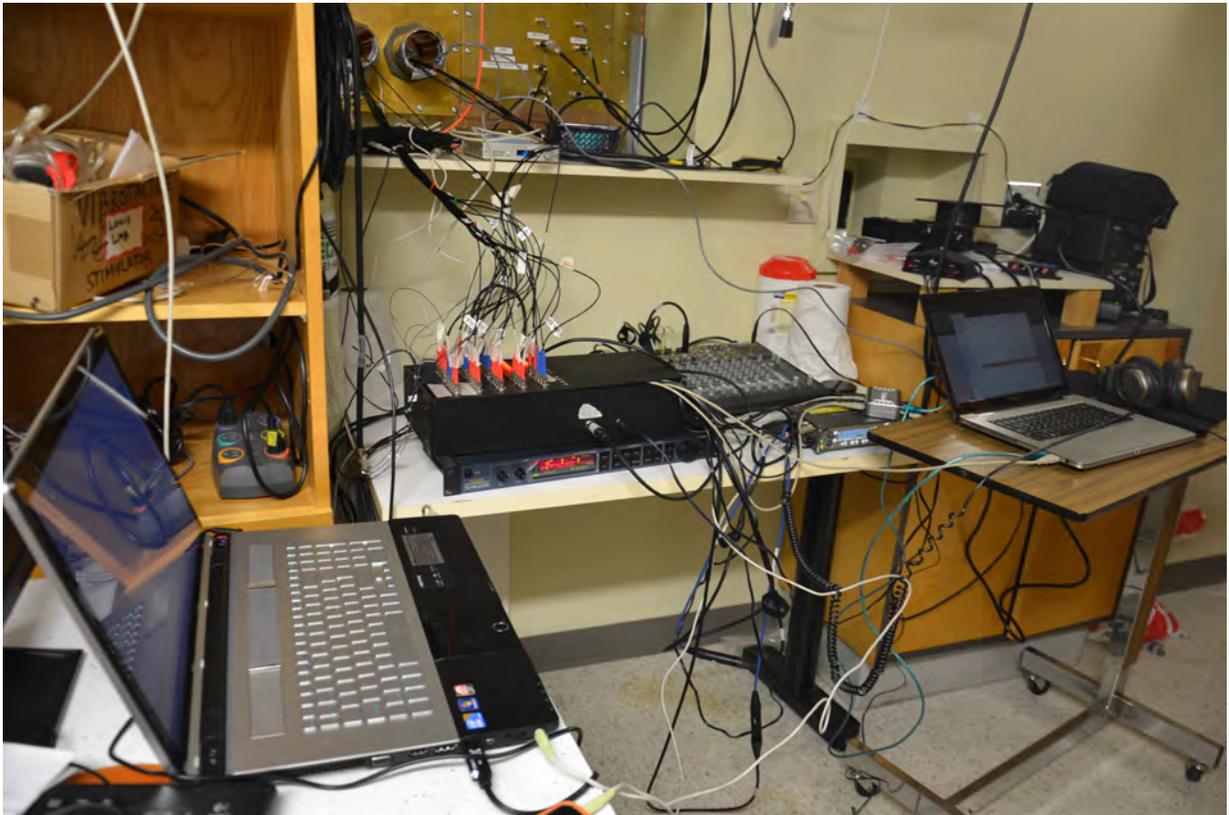


Figure 3.14: MR-COSBBE and other equipment set up in the MRI control room.

ever the protocol was not ideal for high-speed transfers from multiple boards.²² And as a stand-alone system was sought, with a hub to collate and log signals acquired from multiple *fungible-2* boards in synchrony with presented experimental stimuli and feedback,²³ a means to communicate between hub and nodes quickly and synchronously was needed. The SPI bus was chosen as it is electronically simple, can operate into the MHz range with no protocol overhead or handshaking, allows full-duplex data transfer, and does not mandate any standard message passing protocol.²⁴ A daisy-chained configuration was selected for this task as it meant that only four conductors were run between any two devices. The example implementation presented here used six *fungible-2* boards as “slave” devices and the Raspberry Pi as the “master”. The Raspberry Pi was selected as the hub, though the implementation using standard C libraries could be applied to any Linux-based PC with an SPI port, with little to no change to the code.

First, the SPI slave module on the PSoC was configured to receive and send data using separate DMA blocks, DMA_1 and DMA_2 from Fig. 3.12, respectively. DMA_1 was responsible for buffering incoming data from the Raspberry Pi and other *fungible-2* boards, while DMA_2 was responsible for copying data frames from its internal buffer to the SPI bus as well as forwarding incoming data frames from other *fungible-2* boards. Each data frame was 255 bytes long, containing a framing character, the board address,

²²Apple’s OS X operating system was found to have poor reliability for emulated serial port devices. Specifically, the computer would stop responding or disconnect from the peripheral device when data was overloading the PC’s receive buffer if the stream was not read sufficiently quickly.

²³Or in the case of a musical instrument, the hub was to map acquired signals and synthesize audio.

²⁴The SPI protocol uses four signal wires: clock (CLK), chip select (CS), master-out slave-in (MOSI), and master-in slave-out (MISO). The “master” device initiates all data transfers by asserting the CS line—the chip select line is normally high, so asserting CS is equivalent to pulling it low. With its CS line selected, the “slave” device begins to *listen*. Both “master” and “slave” devices each maintain two shift registers, one to shift data into (i.e., receive buffer) and the other to shift data out of (i.e., transmit buffer). This shifting in and out of bits, from the transmit buffer of one to the receive buffer of the other (and vice versa) is performed synchronously on the transitions of the master-generated clock signal. The clock transition phase and polarity are determined by the designer in advance. With each transition, one bit is clocked from the “master” to the “slave” on the MOSI line, and another bit (representing a disparate piece of information) is clocked from the “slave” to the “master” on the MISO line, thus the full-duplex nature of the bus. Bits are continually shifted between devices on each clock period as long as the CS line is maintained. It is standard practice to shift incoming bytes out of a slave device (using a single shift register for the transmit and receive buffers) when the number of bytes received exceeds the length of its internal shift register, thus allowing data to pass through undisturbed—convenient for daisy-chaining multiple slave devices. A transaction terminates when the CS line is de-asserted, which latches the contents in the devices’ shift registers.

time stamp, the contents of the delta-sigma ADC ring buffer, other meta-data (including the number of newly converted samples and the location of the ring buffer write pointer), as well as Reed-Solomon error-correction parity bytes.

Before each transaction—or equivalently, after the completion of each transaction—on each node, meta-data and the contents of the ring buffer were copied into the SPI transmit buffer and RS-encoded, in place. On the hub side, six 255-byte frames were prepared,²⁵ each destined for its respective node board. Each frame, similarly, contained a framing character and was RS-encoded, however its content could either represent an ASCII string, to be interpreted by the node’s command line interface, or a data buffer to be used by the node. For example, during initialization, the hub sent the command !A1 to board number one and !A2 to board number two, and so on, in order to automatically enumerate each board’s address on the bus. Sending a stream of data from the hub to a node could be useful for synthesizing waveforms on a DAC, or using the node as a digital audio device to retransmit the buffer over *S/PDIF*. As each node also buffered the data from other up-stream nodes, data could be shared between nodes and acted upon without input from the hub.

With the SPI clock running at 8 MHz, 1530 bytes representing a complete set of frames from six boards were shifted through the chain in less than 200 μ s, though RS-decoding on the nodes took on the order of 7 ms. Thus, the host program running on the hub, for this configuration, was set to poll the nodes at 10 ms intervals, providing a constant signal acquisition latency. In turn, each hub was able to operate with a maximum ADC sample rate of 10 kHz, or 2.5 kHz per channel (for four channels)—one to two orders of magnitude better performance than the previous *fungible-1* and data transfer over the USB emulated serial port.

Time stamp information was provided by an internal 32-bit counter, or alternately, using ADC_SAR_1 (refer to Fig. 3.11) as an audio input connected to a standard external linear timecode (LTC) source. LTC, implemented as per the standard put forth by the Society of Motion Picture and Television Engineers (SMPTE), is an 80-bit code usually embedded within an audio track of professional audio-video recording equipment, which allows temporal alignment and playback of multiple tracks when recorded on different pieces of equipment. For instance, a stand-alone audio recorder that provides

²⁵The 8-bit symbol RS-encoding scheme limited the block length to $2^8 - 1 = 255$, though shortened codes, where not all bytes are transmitted, was also implemented.

an LTC source, can record and time stamp audio generated by a PC, while several *fungible-2* boards acquire their own set of analog inputs (possibly, at a different sample rate) synchronously with the LTC from the audio recorder, ensuring that data sent to the hub and written as an audio file can all be overlaid for later analysis. The initial *fungible-2* implementation buffered and thresholded the LTC and sent it as a bit-stream along with other frame data to the hub, where it was then decoded using the open-source *libltc* library [221]. Porting this library to the *fungible-2* directly, in order to improve LTC decoding reliability, is planned.

Mr-Cosbbe and the Raspberry Pi

The Raspberry Pi (or *Pi* as a shorthand) has video (HDMI), audio, USB, and Ethernet connectivity, allowing it to be used as a standard PC, and/or accessed remotely as a “headless” system without a monitor. The software running on the hub was developed to be modular and allow portions to run on the *Pi* itself or offload certain functions to other networked computers.²⁶ The audio and video interfaces allowed the *Pi* to present experimental stimuli directly (through the use of a monitor and powered speakers) to a subject.

The *Pi* was configured with Raspbian, an optimized form of the Debian operating system. The *Pi* is underpowered compared to most PCs and operated best without the graphical user interface loaded. Thus, a small program using the Linux “frame-buffer” functionality was written to display images directly to the monitor, without the processing and memory overhead of a windowing system. The program equivalently mapped an area of memory to the screen, allowing the timely presentation of pre-fabricated bitmap images or animations. The experimental software, for instance, could be executed and monitored remotely over a network connection by the experimenter, whilst the direct graphical output of the *Pi* containing the visual stimuli would be visible to participant. Audio libraries and interfaces for presenting auditory stimuli and feedback, as well as developing low-level audio synthesizers are available.

The software written for the *Pi* included a program to configure and poll the SPI port (SPI Transfer), a command line program which took user input (Command Sender), and

²⁶This distributed processing is perhaps more applicable to interactive multi-media systems than to behavioural experiments, but it does allow other computer systems to take over in applications where the Raspberry Pi’s resources are overburdened.

a program that acted as a bridge to log and send data to other systems (The Bridge). After acquiring each data frame from the SPI port, SPI Transfer sent the data over a network socket to The Bridge. Here, the data was RS-decoded and the new samples within the ADC ring buffers from all boards were reorganized into a more useful data structure. The data was optionally logged to a file for off-line analysis, or retransmitted for on-line feedback. One method to generate auditory feedback, using a state machine similar to that described in Section 3.4.4, created a virtual MIDI device to send messages to an attached MIDI synthesizer. The design of an MRI study using the MR-COSBBE system in this way is under development.

The Bridge also incorporated calls to the *libmapper* library, described in Section 3.2.2, thus publishing the acquired signals to the network and allowing them to be arbitrarily and dynamically mapped onto synthesis parameters for auditory, haptic, or visual feedback. Interestingly, The Bridge could be run from a standard PC and communicate over the network with SPI Transfer application running on the *Pi*. This functionality, combined with *libmapper*, enables an experimenter to run a centrally managed multi-participant multi-modal fMRI experiment with multiple MR-COSBBES. Conceivably, an experimenter in one location could conduct an experiment with two participants in MRI scanners in other locations, each with optically-sensed MRI-compatible interfaces, allowing participants to interact directly.

3.4.6 Discussion

The *fungible-2* optoelectronic acquisition board and MR-COSBBE system are the culmination of several prototype iterations. The *fungible-2* has been shown to work reliably under the Linux and Windows operating systems. It worked sufficiently well under Apple's OS X to carry out an experiment, though care must be taken when used as an emulated USB serial port device. The board has been used successfully to interrogate fibre optic sensors, as shown in Sections 5.3 and 6.4, and can operate with a signal-to-noise ratio of 83.5 dB even when USB bus-powered.

MR-COSBBE is capable of capturing 24 16-bit streams of data each sampled at 2.5 kHz, with a latency of 10 ms. Latency may be further reduced by lowering the sample rate and modifying the Reed-Solomon encoding scheme. Through the use of a faster hub, additional performance gains may be achieved. Where forward error correction is not

necessary, the USB audio mode may be used to acquire the 24 data streams, each at 8 kHz (with 1 ms latency and some trade off in bit depth and dynamic range). MR-COSBBE is a system well-suited to fMRI experiments where participant behaviour is measured continuously with optical sensors, and where real-time feedback is presented. The system can also be used with non-optical sensors and has potential applications in DMIs, robotics, multi-media installations, home automation, assistive technology, and rehabilitation therapy.

3.5 Physical Modelling Synthesis

Audio synthesis may take many forms, including wavetable or sample-based, additive, formant, granular, physical modelling, and FM synthesis, among others [180, 248]. Physical modelling synthesis differs from other synthesis methods in how the audio samples are computed, specifically through the numerical solutions to dynamical systems which represent physical phenomena. Due to the nature of the equations which govern these phenomena, they tend to produce organic sounds with micro-variability in amplitude, pitch, and timbre. Some systems have many parameters, often modelled as mass-spring-damper (or resistor-inductor-capacitor) networks, where no single parameter independently controls the generated sound's characteristics. That is, all parameters may affect the amplitude, pitch, and timbre in some way, but none does so exclusively. Similarly, modifying a single synthesis parameter may impact some or all of the amplitude, pitch, and timbre to varying degrees. Parameters must be constrained to maintain the stability of the system, so that it does not "blow-up" numerically. The way in which the system is solved (e.g., the size of the time step and the precision of the numerical representation) may also cause the solution to misbehave.

Producing a variety of musically interesting sounds from a disparate array of parameters requires an interface and a mapping that together span the gamut while affording nuanced control of micro-structure. While much research has been focused on accurate and efficient physical modelling [132, 246, 247], interfaces that specialize in their control are sorely lacking [21]. In Chapter 4 an interface is presented, through which I attempted to provide such functionality. In this section, however, we discuss the implementation of two different physical modelling synthesizers and how they might be used within embedded systems toward the creation of a stand-alone DMI.

The synthesizers were initially developed in MATLAB, then ported to C and used as synthesis libraries on PCs and on various embedded systems. Physical modelling synthesis, due to its higher computational complexity, has for decades taken a back-seat to more economical methods, though with the advances in computer processing power, many of these algorithms can now be performed on consumer processors in real-time.

3.5.1 FitzHugh-Nagumo Excitable Cell

Some years ago, during the inter-match banter at a game of table tennis, Bartłomiej (Bartek) Borek, a friend and scientist, gave an explanation regarding his work culturing embryonic chicken heart cells [20]. These excitable cells can be modeled with the FitzHugh-Nagumo (FHN) model, itself a simplification of the Hodgkin-Huxley squid giant axon model [71]. The model has both monostable and astable behaviours: meaning that an input pulse can trigger a transient output that then returns to the stable state, or the system can continue to oscillate indefinitely. After being introduced to the model and discussing the possibilities of its use as a synthesizer, I got to working on a simulation. The iterative time-domain solution that was used for this implementation, presented in [117], was given as:

$$v_{t+1} = v_t + dt \cdot \delta v = v_t + dt(-w_t - v_t(c_t - a)(v_t - 1) + I) \quad (3.1)$$

$$w_{t+1} = w_t + dt \cdot \delta w = w_t + dt(b \cdot e \cdot v_t - g \cdot e \cdot w_t - d \cdot e) \quad (3.2)$$

where v and w refer to the membrane potential and the recovery variable, respectively. I represents the stimulus current, dt the integration time, and a , b , d , e , and g are used as other synthesis parameters. This model provided a simple nonlinear oscillator with coupled control over fundamental frequency, amplitude, and timbre. Fig. 3.15 shows the output of one FHN cell, parametrized to oscillate continuously. Several cells were connected together in series and in parallel to provide more interesting results.

Spatial Distribution of Excitable Cells

An extension was implemented which incorporated a two-dimensional grid of excitable cells coupled by the Laplacian heat diffusion equation, giving rise to a more complex synthesizer. The modified equation for two-dimensional coupling is given as:

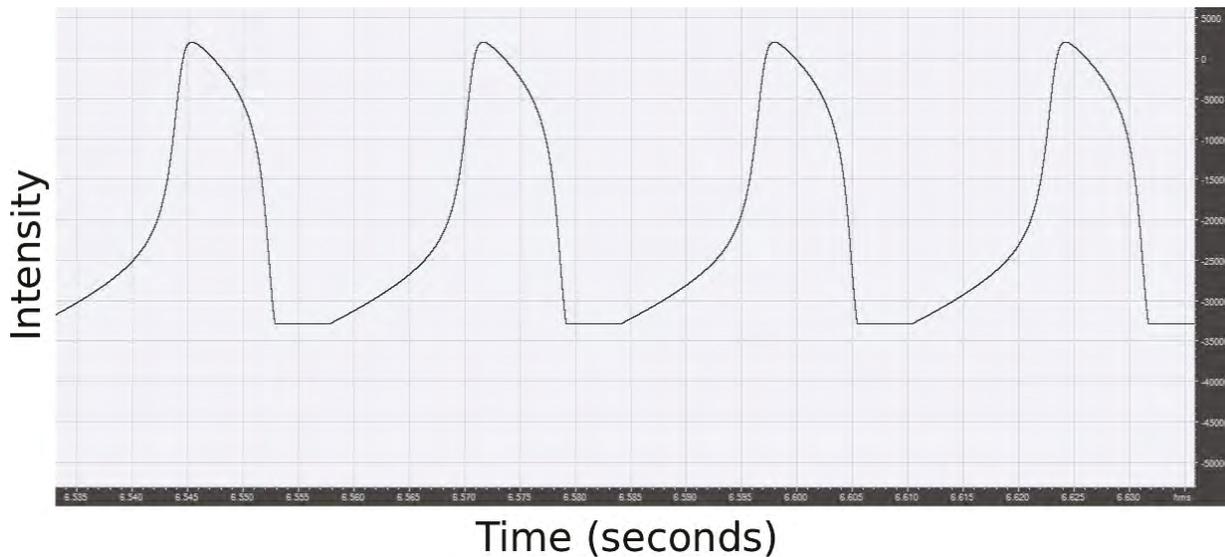


Figure 3.15: Waveform of FHN model in astable multivibrator state.

$$v_{t+1} = v_t + dt \cdot \delta v + dt \cdot \mathbf{D} \left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} \right) \quad (3.3)$$

Where \mathbf{D} is the diffusion coefficient (a scalar for isotropic diffusion or a matrix for anisotropic diffusion), and where the partial derivatives are computed as discrete second-order central finite differences, equivalent to the scaled sum of differences in local membrane potential between neighbouring cells. An explicit analogy to a planar control surface was accomplished by spatially varying the model's parameters and mapping excitatory input at specific locations to stimulus current at corresponding cells on the grid. The sound output is taken at one (monaural), two (stereo), or several locations (through some averaging or using a multi-channel audio system). The boundary conditions can be configured to clamp or reflect²⁷ the output, or wrap the output to the other side of the grid. With a sufficiently large grid and reflecting boundaries, rhythmic patterns emerged.

²⁷The reflection coefficient can even be frequency dependent. Care must be taken, however, to ensure system stability.

3.5.2 Ishizaka-Flanagan Vocal Fold Model

For decades, the voice has been studied by those wanting to synthesize our most primordial instrument [44]. The voice is usually approached as a source-filter model—with the vocal folds producing the glottal waveform which is then filtered by the vocal tract. While the glottal waveform is often implemented as a pulse train, there have been a number of physical models put forth to explain the dynamics of the vocal folds [5]. The high dimensional model of Ishizaka and Flanagan (IF) allows for more subtle vocal qualities, not available to parametric and lower-order models, such as “breathy voice” and “vocal fry”. Essentially, each vocal fold is symmetrically modelled as a two mass-spring-damper systems coupled with a nonlinear spring and driven with a differential air pressure. Federico Avanzini gives the dynamical equations describing the vocal fold model in [5] as well as diagrams and explanation.

Solving the nonlinear system of differential equations is non-trivial due to the presence of a delay-free loop: the pressure and position variables are inter-dependent for the same instant in time. Thus, to simplify the otherwise iterative convergence problem and compute the solution in a real-time fashion, the current positions of the masses were calculated using previous values of pressure, area, and displacement. That is, all of the displacement-dependent and pressure-dependent coefficients for the time t_n were calculated using values from time t_{n-1} . These coefficients were used to update the vocal tract areas at time t_n , and then solve for the glottal flow. The quadratic equation which related the flow to the displacement (through the glottal area) was solved and one root was propagated as the flow. The root that was thrown out was furthest in value to the flow at the previous time step. Only the real part of the flow was retained, as some parametrization resulted in complex flow. One set of glottal pressures were then calculated using the flow, while another set of pressures were calculated using displacement. These pressures were averaged and stored to be used during the next time step. To deal with the nonlinearities arising from the folds coming into contact with each other, different equations were used depending on the displacement of each mass. The implementation added a restoring force whenever the folds were in contact with one another. A sample waveform produced with this IF model is presented in Fig. 3.16, showing the transition to a more stable state.

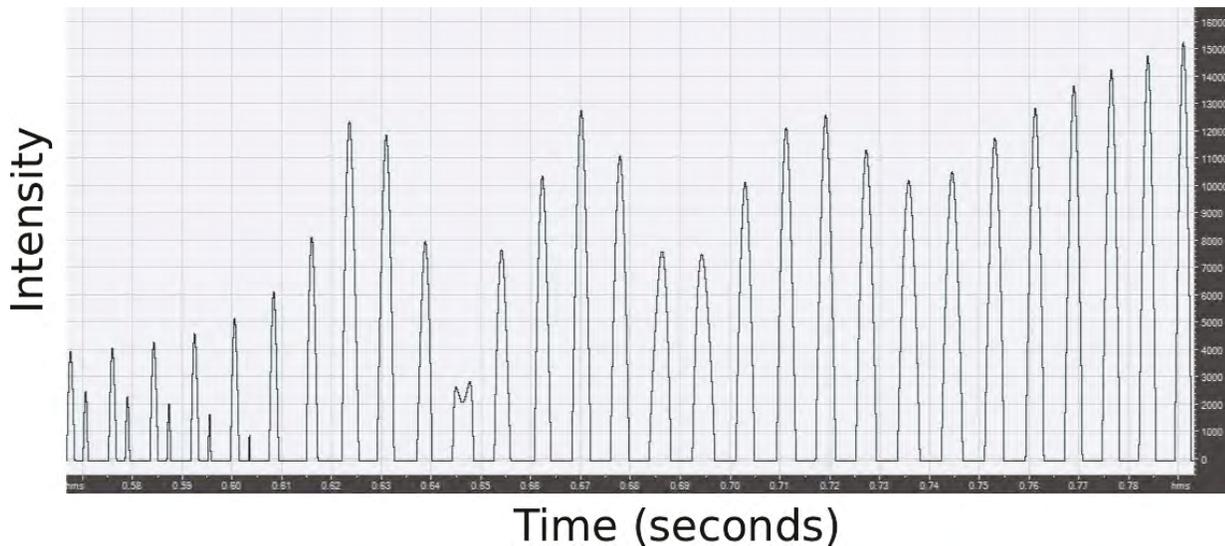


Figure 3.16: Glottal waveform produced with the IF model.

3.5.3 Real-Time Synthesis

Prototyping on a PC

The initial models were prototyped in MATLAB to get a better sense of the behaviour of FHN and IF synthesizers, to debug and improve the implementation, and to find a useful range of parameters. The IF synthesizer was directly ported to a C language implementation, while the FHN code was implemented in the FAUST functional signal processing language [194]. The FAUST compiler generated a C implementation and a wrapper for MAX/MSP.²⁸ The FAUST wrapper was also used as a template with the IF synthesizer to generate a MAX/MSP external. These MAX/MSP externals were used to further refine the models' parameter bounds and to generate limits to clamp the output to keep the models from “blowing-up” and automatically resetting the synthe-

²⁸The FAUST FHN implementation, excluding bounds checking, was written as:

```
new_v(v,a,I,dt,b,d,e,g)=
    (((I-(v-1)*(v-a))-(new_w(_,v,b,d,e,g,dt)))*dt)+v),I,new_w(_,v,b,d,e,g,dt);
new_w(w,v,b,d,e,g,dt)=((((dt*(e*((b*v)-(w*g)-d)))+w)<:(_,_))~(_,_)):(_,!);
process =(new_v)~(_);
```

sizer when they did. Mappings for these synthesizers were prototyped using MAX/MSP with control input from *fungible-1* boards via USB. The control input stream was parsed using a script written in the *Python* programming language, as a cross-platform application was required.

Another instantiation of the synthesis algorithms written in *C* were compiled as dynamically-linked shared libraries. A *Python* script was written to compute samples using the shared library and output audio using the *PyAudio*²⁹ cross-platform audio library. Another *Python* script was written to take the parsed serial input and using the *PyBrain* cross-platform modular machine learning library³⁰ to map control inputs onto synthesis parameters. Augmentations to this system included using *libmapper* to map data from the parser onto the machine learning module, and again to map the classified data onto synthesis parameters. The synthesis algorithms performed well, though latencies in parsing and shuttling data between programs meant the system was less responsive than had the system been written in one compiled program. The system was sufficient for prototyping.

Embedded Synthesis, Mapping, and Classification

The *C* code was further refined and integrated with code to manage the audio buffers and the DAC on the embedded systems described in Sections 3.2.3, 3.3.2 (as well as on the equivalent PSoC 5 microcontroller of the *fungible-2*). The embedded mapping architecture, as described in Section 3.2.2, allowed for static mapping, however, manual-parametrization of the models was tedious.

A naive Bayesian classifier [152], including on-line learning and decoding modes, was implemented to implicitly generate continuous mappings given a set of input vectors (representing an acquired gesture) and their respective set of synthesis parameters. The Bayesian classifier was employed as it is a simple supervised learning technique that can be computed on-line, essential for a real-time embedded system. The classifier's learning mode simply computed a running mean and standard deviation for the set of inputs for each class, as determined by the user.³¹ The decoding mode computed

²⁹The PyAudio library is available at: <http://people.csail.mit.edu/hubert/pyaudio/>

³⁰The PyBrain library is available at <http://pybrain.org/>

³¹The user could also adjust the prior weights, such that the rest state was more likely to occur in

likelihood ratios for each class, which were then used as weights to interpolate between respective synthesis parameters. This embedded mapping algorithm was used with the Ballagumi (see Chapter 4) and IF synthesizer to enable the implicit mapping of a complex set of inputs onto a complex set of outputs. The whole acquisition–classification–mapping–synthesis system performed adequately with a reduced synthesis sampling rate (5 kHz), and decoupled control rate (100 Hz to 1 kHz), and twelve gesture classes.

As the audio output from a single FHN synthesizer cell was underwhelming, and a large grid of cells needed more memory than available, the synthesis algorithm was found to be underwhelming on this hardware.

Audio Synthesis on the Raspberry Pi

The IF synthesis algorithm was ported to the Raspberry Pi using the low-level OpenMAX IL audio client [187]. This has yet to be integrated with the MR-COSBBE system and *libmapper*. The synthesizer performed as expected at a sample rate of 8 kHz. The limits of this system for real-time audio synthesis have yet to be explored.

Audio Synthesis on Android Devices

Another “open”³² embedded platform is Google Inc.’s *Android* operating system. While the hardware specifications of devices running *Android* are not standardized, newly released devices have increasingly more memory and higher processor throughput. The NVIDIA Tegra 4 and Qualcomm Snapdragon 800 have more powerful architectures than the *Pi*, running at about three times as fast and with up to 16 times the memory [191, 211, 214]. The portable embedded platforms provide a solid foundation for real-time audio synthesis, however, Java is the primary programming language. The *Android* Native Development Kit (NDK) provides C compiler support for low-latency audio programming, though a Java layer is still required for applications.

Thus, a route was sought that would circumvent the standard Java interface while enabling low-level synthesis on an *Android* device. The IF algorithm was ported to *Android* by integrating the C synthesizer code with calls to the low-level OpenSL ES

order to suppress spurious evocations.

³²While the source code is made publicly available to download and compile, *Android* does not share the same community-based free and open-source ethos as the Raspberry Pi Foundation.

embedded audio library [187], and compiled as a dynamically-linked shared library. Using the *Python-for-Android* (P4A) scripting layer [209], the user could spawn the synthesizer and control its parameters. As a proof-of-concept, the ambient light sensor and accelerometers of the *Android* device were mapped onto synthesis parameters, allowing for gestural control of sound. While P4A supported a graphical layer, a working prototype has yet to be completed. Additionally, by loading a complete Debian/Linux operating system within the *Android* environment, it was possible to compile and install *libmapper* on the device. Audio synthesis performed as expected, though there was a noticeable delay in fetching the sensor data, as this was performed by the operating system at a much lower rate than the audio loop.

3.5.4 Discussion

The ability to perform acquisition, mapping, and synthesis on embedded systems is an exciting prospect as high-performance mobile processors continue to improve and become more pervasive. Not only can a single person or device make music in isolation, but many people and devices networked together can interact and create. It is important to remember, however, that while these digital marvels include increasingly numerous and capable sensors, a phone need not be the interface through which the performer interacts. And for applications where virtuosity is prized, an interface with a complementary acquisition system is required.

3.6 Conclusion

This chapter presented the design and implementation of a collection of hardware, firmware, and software for electronic musical interfaces. The main applications for this equipment, as discussed, are neuroimaging experiments and DMI performance. Both of these applications share the need to acquire user input and generate a response perceived by the user with minimal latency. I have endeavoured to create a stand-alone system for use by musicians and behavioural neuroscientists to exercise their craft. I have also discussed some of the problems I faced and future directions of interest. The system will continue to be refined and used in further experiments and with additional interfaces.

Chapter 4

Ballagumi

In this Chapter, we present a novel electronic musical instrument, the Ballagumi,¹ that consists of a flexible physical interface cast in silicone with embedded fibre optic sensors [171]. The instrument was approached with a firm set of constraints and goals: the instrument should be composed of a physical interface generating a collection of analog signals with the purpose of controlling a physical modelling synthesizer, the interface should be sensed in a holistic fashion encouraging a complex mapping with nuanced control of acoustic properties, the instrument should be self-contained and not require a personal computer for audio synthesis, and in addition, the physical interface should be compatible with MRI scanners in order for it to be used in neuroimaging studies.

The novelty of this instrument for MRI studies is quite exciting as the subject has no preconceived notions of how to interact with the instrument nor the auditory feedback presented, thus providing a totally unique opportunity to study motor learning of musical tasks. In an attempt to approach the playability of an acoustic instrument meant for expert performance, physical modelling synthesis based on a dynamical system was employed (as discussed in Section 3.5). The complex interaction of synthesis parameters on its acoustic properties (i.e., pitch, amplitude, and timbre) allowed for an organic and engaging experience, though perhaps difficult to control in a musical fashion without sufficient practice.

¹The name Ballagumi, is a *portmanteau* of the words *balagan* and *gummi*—two words that have similar meanings and spellings in a variety of languages (including Polish and Hebrew) and is translated as a mess of rubber.

4.1 Interface Design

4.1.1 Design Concept

The interface was to have nuanced control of timbre, pitch, and amplitude, and to that end many sensors were used to collect information about the deformation of the instrument as a whole. While redundant sensing leads to coupled sensor responses, the coupling of acoustic properties of the sound is inherently coupled within the dynamical system, making the mapping section an equally important part of the equation. In order to draw users to the instrument, an organic form and a viscerally exciting material was necessary. The initial sketch of the interface is shown in Fig. 4.1. While the winged bat or ray-like form is an idiosyncratic design, it provides the performer with affordances for striking, pushing, pulling, plucking, twisting, and stretching gestures. It was meant to sit on the lap to enable bi-manual control and allow the legs to provide support and leverage for manual manipulation. The silicone rubber material, as noted by Morris et al. in [184] provides an engaging and continuous sensation of haptic feedback owing to its viscoelastic properties. Silicone is inert and MRI-compatible; as well, it can be cast using standard moulding techniques.

4.1.2 Optical sensor design

Fibre optic sensors embedded in silicone were used to sense strain, flexion, force, proximity, and displacement. Simple sensors were created with an LED, a photodiode, and 1 *mm* plastic multimode fibre². Fibre optic bend sensors were created by “sensitizing” the fibre along its length by removing the cladding from the side of a fibre to reduce the amount of total internal reflection at the sensitive site [53]. Thus, with the cladding removed on one side of the fibre at a specific location, the photodetector measured a decreased amount of light when the sensitive region was bent convex-up, whereas the amount of light increased when the sensitive site was flexed concave-up. Fig. 4.3 shows a grid of treated fibres spaced at approximately 4 *cm* intervals to sense flexion along multiple axes.

A pressure sensor was created to sense the light transmitted from the end of one fibre to another either through transmission or back-scattering. By embedding fibres

²All optical fibres in this chapter are similarly 1 *mm* POF, unless otherwise stated.

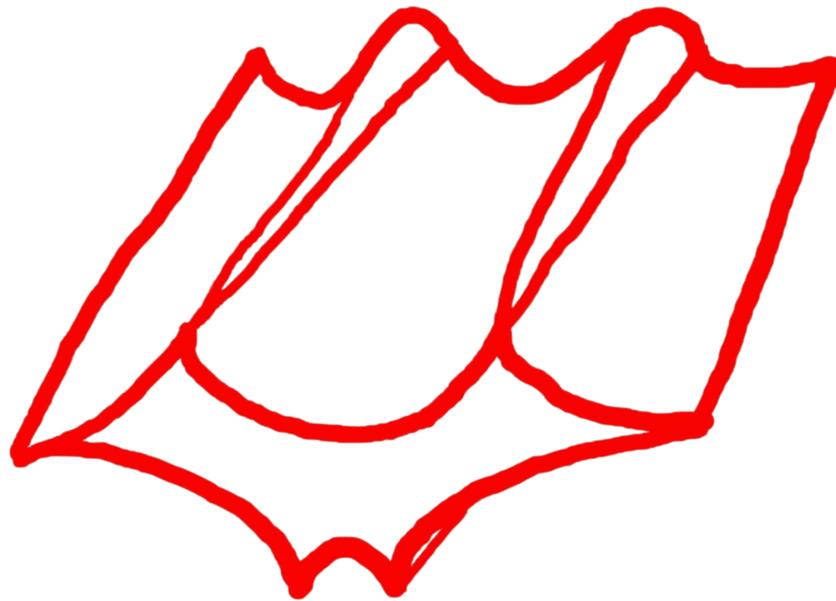


Figure 4.1: Initial sketch of the interface, with wings and spines.

in a translucent elastomeric compound, and allowing the light to propagate through a portion of the interface on its way from an emitting to a receiving fibre, the interface itself became part of the measured signal path. Fig. 4.2 shows a differential transmissive fibre sensor which can sense compressive and shear forces.

4.1.3 Initial Prototype

The initial prototype interface, known as the GGT³ [117], was made using a small mould created with the Forton⁴ casting system (a combination liquid polymer, modified gypsum, and fiberglass) with the original model made in plasticine. The initial prototype was cast in silicone (10A shore hardness). The interface, shown in Fig. 4.4, was embedded with a combination of optical bend sensors and three-fibre end-coupled back-scatter sensors. While the bend sensors performed well enough with large SNR (approx., 45 dB), one of them snapped during use turning it into a transmissive end-

³An abbreviation for the *Goey Glob Thing*, suggested by Vanessa Yaremchuk.

⁴A large thank you to the wonderful sculptor, Mark Prent, for providing guidance and assisting me with this first prototype.

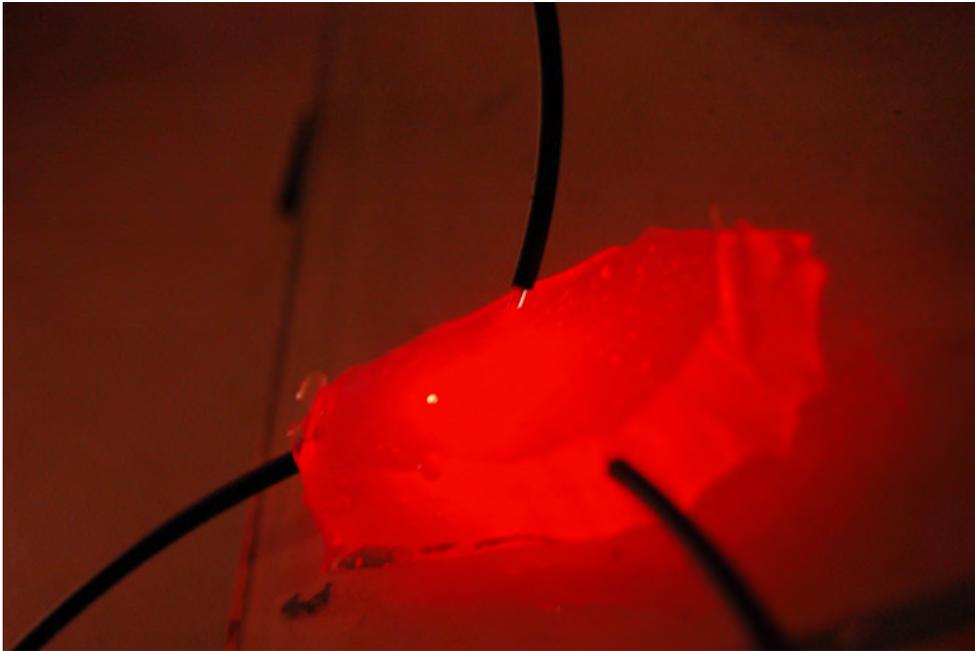


Figure 4.2: Differential fibre optic deformation sensor.

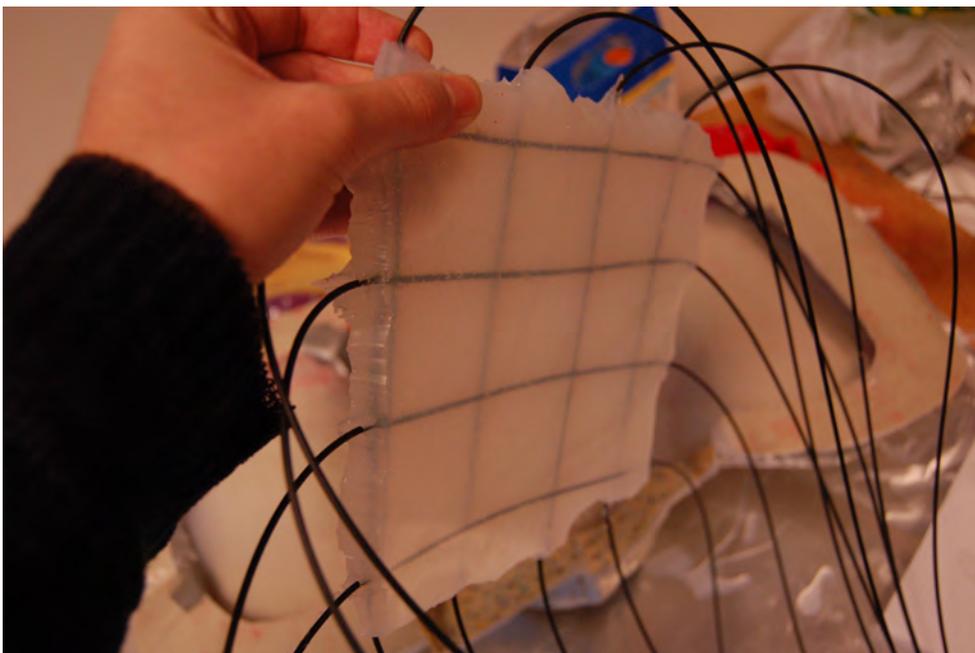


Figure 4.3: Fibre optic multi-touch force and bend sensor grid.

coupled sensor. The fibres have a minimum bend radius beyond which their optical properties deteriorate and then break. It was found that the fibre that broke was treated in such a way that a significant amount of the cladding and core were removed, and that the silicone was not sufficiently thick to limit the minimum bend radius of the fibre. Subsequent bend sensors were treated more delicately by removing a smaller amount of the cladding, embedding the fibre in a thicker silicone substrate, and retaining the protective jacket as a stress reliever on the opposite side of the treated area.

Back-scatter sensors were used to measure twist and other deformation at the middle cusp of the interface. One fibre emitted light while a pair of receiving fibres measured the reflected light. These sensors performed admirably, though suffered from a lower SNR than the bend sensors (approx. 20 dB), and were sensitive to fluctuations in ambient light. As the emitting and sensing fibres pointed upward, they responded to non-contact reflection as well as deformation, but were also quite sensitive to shadowing of ambient light (which was an interesting, yet unintended form of interaction). Performing differential sensing and time-domain compensation of ambient light greatly reduced the effects of ambient light. Initial synthesis and mapping tests were conducted using the system described in Section 3.2, showing it was possible to control synthesis parameters in real-time [117].

4.1.4 Sculpture, Mould-Making, and Casting

The Ballagumi prototype was first sculpted in plasticine, as shown in Fig. 4.5. The form was refined and the surface finished to a smooth texture. Several coats of high-gloss lacquer were spray-painted onto the plasticine to seal the model and give it an even shinier texture, as without these steps the finish would have been more similar to that of the first prototype, which did not have the feel and fine texture that was sought. Several coats of release were applied to the model before coating it in a liquid plastic moulding material. Several coats were needed to cover the entire model with a sufficient thickness. De-moulding was a painstaking process, as removing plasticine from the deep recesses of the mould proved difficult.⁵ The mould was processed with

⁵More experienced sculptors would have used a two-part mould to properly cast the undercuts. Several holes in the mould were accidentally made while de-moulding due to insufficient mould thickness of portions of the underside of the object. The plastic moulding material was quite brittle, cracked easily



Figure 4.4: Initial GGT interface prototype and Forton mould.

a Dremel tool to remove excess material and to facilitate final casting and embedding of sensors.

Fibre optic sensors were prepared separately and tested before being potted in the interface. A four-by-four lattice of fibres forming a multi-touch force and bend sensor, as shown in Fig. 4.3, was created by removing portions of the fibre jackets and cladding and then embedded the fibres in a block of silicone. Six transmissive end-coupled deformation sensors, as shown in Fig. 4.2, were prepared by removing the jackets and sliding the fibres into pucks of silicone. Four of the sensors each incorporated three receiving fibres, while two of the sensors consisted of only two receiving fibres—all having only one emitting fibre. These sensors showed much higher SNR over the back-scattered version in the initial prototype (approx., 35 dB), as the emitting fibre-end faced the receiving fibre-ends. This arrangement was much less sensitive to ambient light as they were embedded more deeply in the silicone and did not face upwards. Eight additional bend sensors were arranged along the perimeter of the interface's wings.

and caused cuts to the hands while de-moulding.



Figure 4.5: Plasticine model of the Ballagumi.

Casting commenced by pouring a base layer of silicone into the mould and letting it cure slightly. This ensured that the surface of the Ballagumi would be free of defects and the sensors would not protrude from the interface. Next the pre-fabricated sensors were positioned in the mould, as shown in Fig. 4.6, and were potted in place by filling the mould with silicone, as shown in Fig. 4.7. Retaining walls built out of plasticine were added to the plastic mould to increase the depth of the cast object, providing extra strain relief for the optical fibres exiting the interface. The Ballagumi was left to cure at room temperature before being de-moulded. The cured and de-moulded silicone interface with embedded fibre optics is shown in Fig. 4.8.

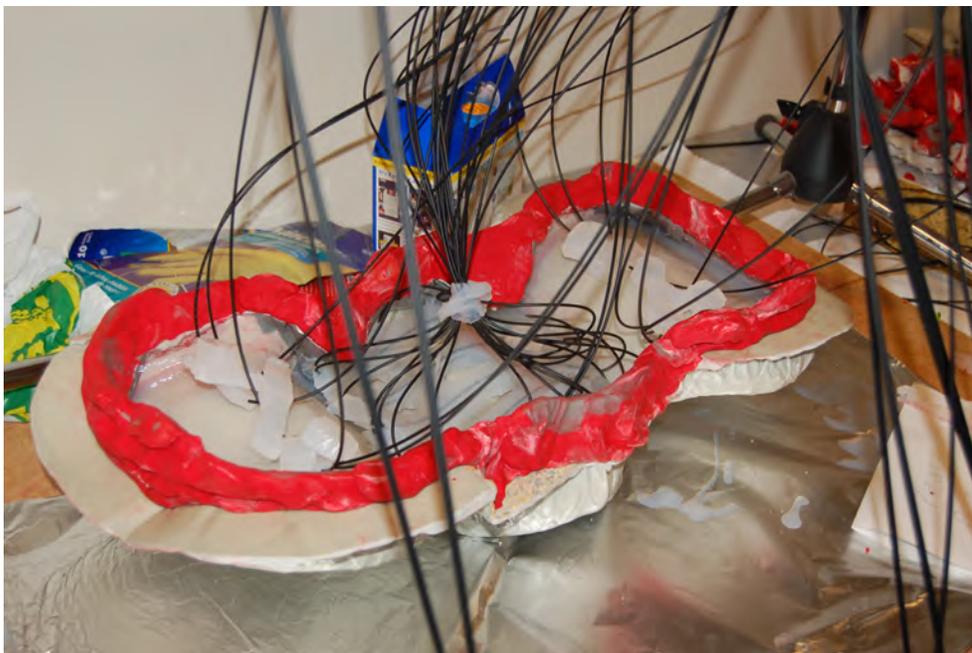


Figure 4.6: Sensors placed in mould after initial silicone base layer set.



Figure 4.7: The Ballagumi with sensors potted before final silicone pour.



Figure 4.8: The Ballagumi, silicone interface with embedded fibre optic sensors.

4.2 Signal Acquisition, Mapping, and Synthesis

Acquisition was accomplished using the system described in Section 3.3. Four *fungible-1* boards connected to a hub over I2C were used to digitize input signals, map control signals, and synthesize audio. As the ADCs converted the signals from three input channels at once, this allowed for the synchronous comparison of differential signals as well as triple-ended measurements in the same time-step. Ambient light compensation was accomplished, as discussed in Section 3.3.1, by toggling emitters while making successive measurements of the same channel. Mapping and synthesis were accomplished on-chip, as detailed in Sections 3.3.2 and 3.5.3. Both explicit many-to-many mappings and implicit mappings using a naive Bayesian classifier were implemented, allowing for real-time control of pitch, amplitude and timbre of physical modelling synthesis.

4.2.1 Qualitative Study

Additionally, data from the *fungible-1* acquisition boards was communicated directly to a PC via USB in order to analyze and map the signals in a more accessible way.

A setup, described in Section 3.5.3, was developed for Mahtab Ghamsari-Esfahani, to parse and collate the serial data streams in *Python* and publish them to the network using *libmapper*. This setup was used to perform a qualitative analysis of the impact of different mappings on a performer's experience while exploring a novel instrument [88, 166]. Ghamsari found that the weight of the instrument and the pressure required to excite the synthesizer led to player fatigue. In order to alleviate this issue, the next prototype of the instrument is to be lighter and make use of a more flexible silicone.

The problems encountered in Ghamsari's study, including fatigue due to the lack of responsiveness, were likely due to the lack of proper sensor calibration, utilizing only a small subset of available sensors, unmitigated and uncharacterized latency in the signal processing, mapping, and synthesis chain, and a mapping and synthesis strategy that exacerbated the perceived latency. In [166], the authors wrote "[i]ndeed, the additional difficulty in playing the instrument with the continuous energy input mapping is augmented by the latency. Improving the signal acquisition and parsing would in turn improve the overall playing experience." While a parsing script written in C as opposed to *Python* would have improved the data throughput, the acquisition system was certainly not the limiting factor, as synthesis performed on-board an underpowered microprocessor (see Section 3.5.3) showed no such latencies nor lack of responsiveness after proper configuration and calibration. Their experiment, in contrast, was performed on a more-than-capable Apple desktop computer, though the performance of the Mac OS X emulated USB serial port driver was questionable.

4.2.2 Gesture Analysis

As a deformation of any single point on the flexible interface simultaneously implicated responses from surrounding sensors, there was a high degree of correlation between sensor signals. However, due to the nature of the sensors, the responses were not strictly linear in response to flexion and pressure, nor even monotonic. The long-exposure photograph in Fig. 4.9 shows the interrogated sensors as brightly-lit regions. As the on-chip Bayesian classifier was performing its calculations based on the analysis of many channels of possibly redundant data, there was the possibility to decrease the computational load of the classification algorithm by reducing the dimensionality of the dataset.

Using the *fungbile-1* acquisition system and logging time-series data to a PC, further gesture analysis was performed. To limit the computational complexity, only eight channels of data were acquired from various sensors on the right side of the Ballagumi. 24 gestures were performed, including a rest state, and analyzed using MATLAB with the purpose of extracting the meaningful and non-redundant data. This analysis was meant to further inform the design of a data reduction scheme to simplify gesture classification and mapping. Each of the gestures was repeated 12 to 20 times and saved to a file with a descriptive label, (e.g., “base push,” “apex pull,” “bend sideways concave-down,” etc.). The gestures were chosen to reflect a range of interactions afforded by the Ballagumi, including bending the wings up and down, pulling the spine left and right, and pushing into its surface. These gestures were performed at various points on the right side of the interface. While the data presented here originates from only eight sensor channels, a fairly manageable number, we must be wary that the Ballagumi has upward of 32 input channels (depending on the interrogation configuration).

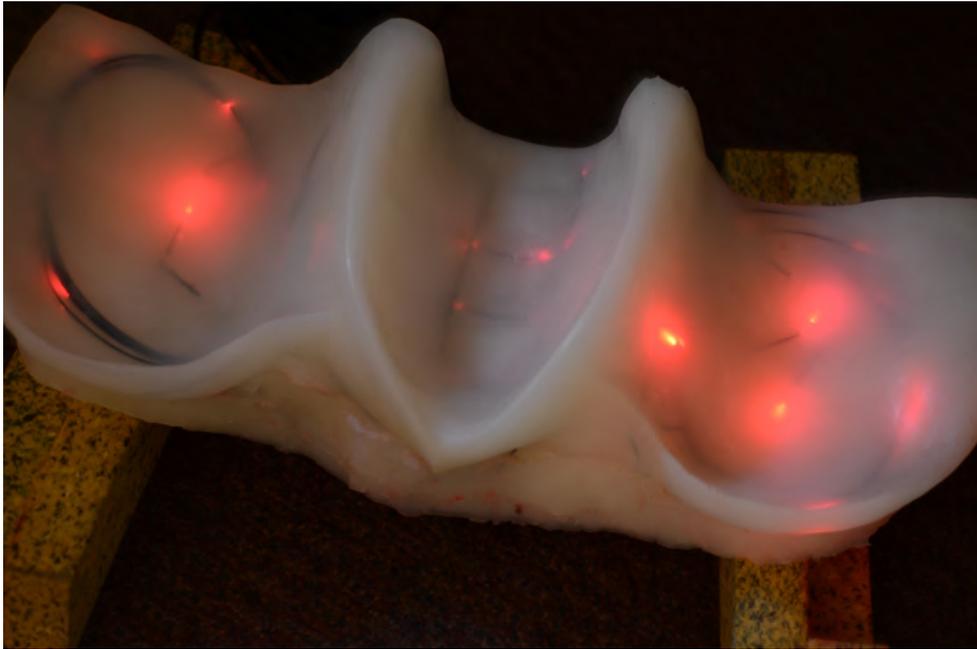


Figure 4.9: Glowing Ballagumi, showing sensor placement.

For a specific gesture or deformation, we can identify which sensors are actively perturbed and which ones co-vary with the others. Principal components analysis (PCA) is one of the most widely-used tools for reducing the dimensionality of a dataset

in order to extract the most salient features [129]. For instance, consider Fig. 4.10, which shows the time-domain traces of eight sensor channels for a particular gesture, in this case pushing down on the apex of the Ballagumi. The gesture is repeated 15 times, and between each repetition, the interface is allowed to return to its original shape.⁶ As can be seen, many of the signals follow the same trajectory, with slight distortions on some channels. Using PCA we can reduce the number of de-correlated signals to a manageable few. In Fig. 4.11, we see the first three principal components of the same gesture as in Fig. 4.10, where 91% of the variance in the data came from the first component. We have therefore reduced the number of meaningful signals from eight to one, as the residual components are either from noise sources or nonlinearities of the signal.

This can be repeated for all of the collected gestures, and we can find that for any one gesture, there are only one or two significant components that would suffice to represent the gesture. A calculation of the correlation between the original sensor signals and the rotated PCA vectors provides a measure of which signals contribute the most amount of information to specific principal components. Not surprisingly, the non-monotonicities and other nonlinearities of certain sensor channels contribute to secondary signals.

More interesting is the ability to visualize all of the gestures within a dimensionally-reduced space, in order to determine the similarity of different types of deformations as represented by the sensor data. First, the time-series data for each gesture was automatically segmented into rest, transition, and full-deformation (i.e., *deformation* data) categories. Next, the *deformation* data from all gestures was concatenated and processed with the PCA algorithm to produce a rotation matrix (i.e., the eigenvectors of the covariance matrix, truncated to three principal components). Finally, the initial *deformation* data was then projected onto the three-dimensional rotated space, as shown in Fig. 4.12, where each colour corresponds to a different deformation class (taken from the original gesture labels). It is apparent that some types of deformations are distinct and separable, while others have more within-class variability or show similarity to other classes.

To examine how a computer might classify these gestures based on the clustering

⁶Or close to it—as the instrument is played on the lap, the rest position is a noisy and vague “gesture” in its own right.

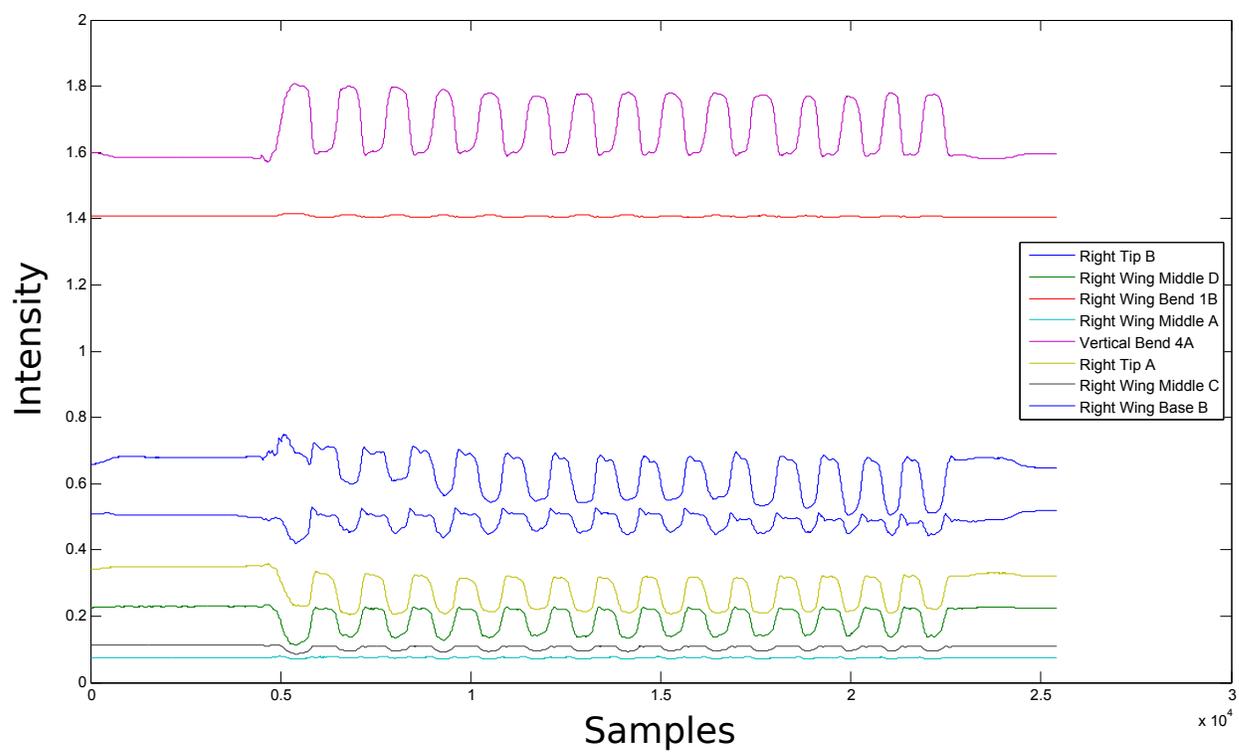


Figure 4.10: Sensor data for repeated “Apex Push” gesture.

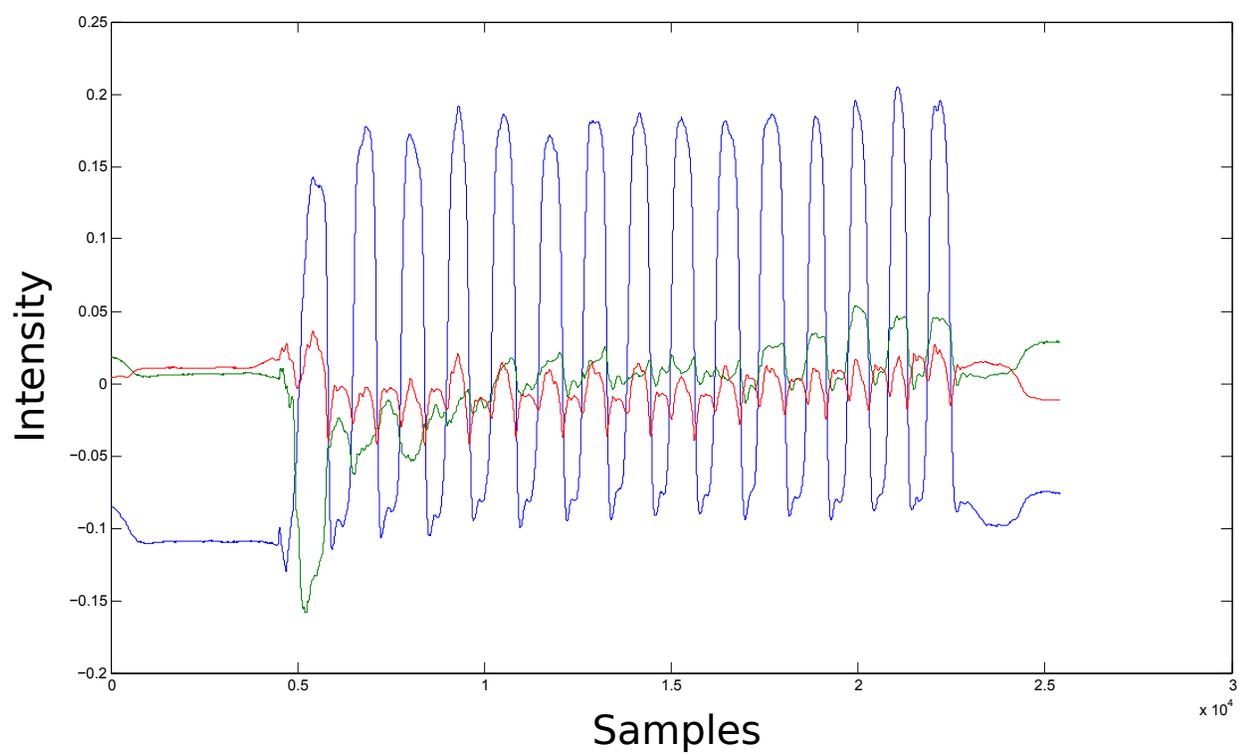


Figure 4.11: The first three principal components of sensor data for “Apex Push” gesture.

of projected data, the rotated dataset without labels was processed by a k-means clustering algorithm into 23 classes (the rest state was, again, excluded due to its high variability and many outliers). The same rotated dataset is shown in Fig. 4.13, with each colour corresponding to classes generated by the k-means clustering algorithm.

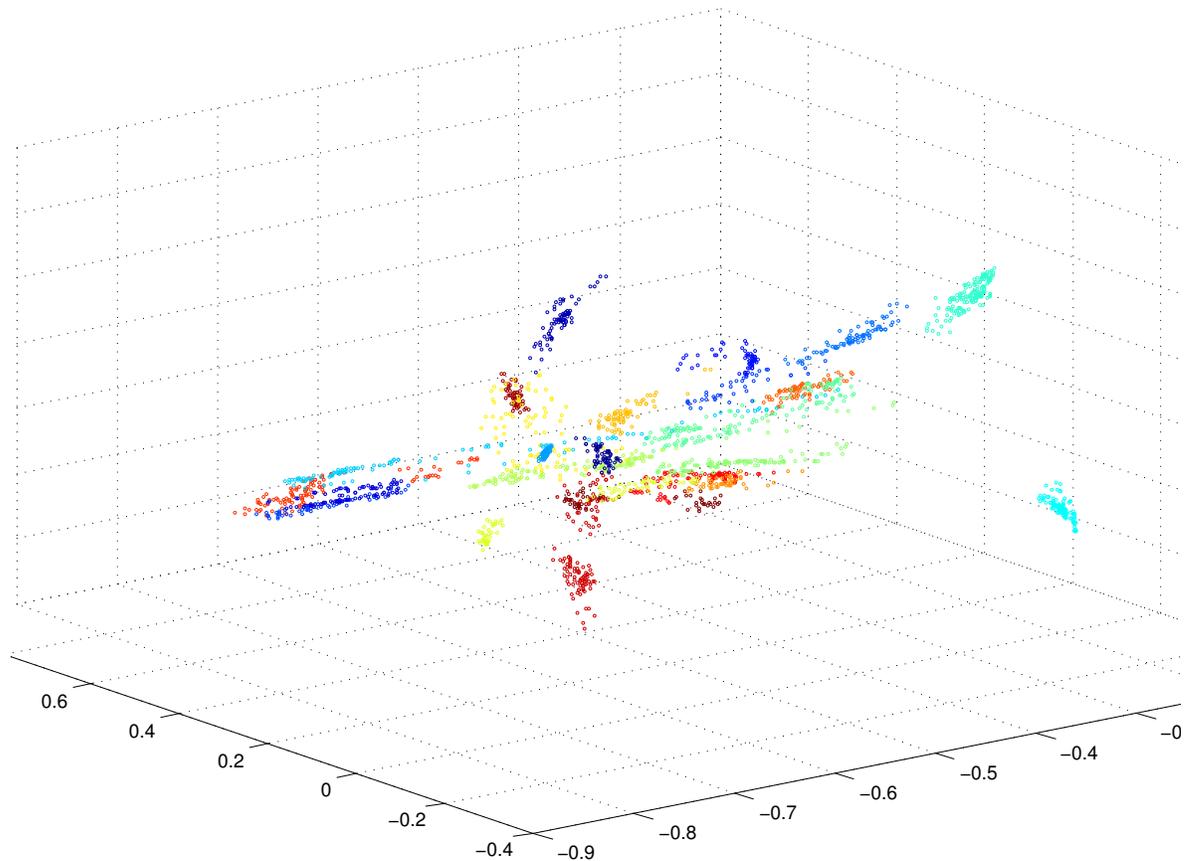


Figure 4.12: Clustering of gestures and their projection on a three-dimensional space using PCA and initial labels.

While the k-means clustering did not classify the *deformation* data with the same boundaries, it did provide very similar results to the hand-labelled gestures, with perhaps even more specificity—as it is possible that certain gestures were not initially performed and reproduced with sufficient care and distinction.

Thus we have shown that dimensionality reduction of redundant sensor data can be used to enhance the computational complexity of mapping, with the use of a small

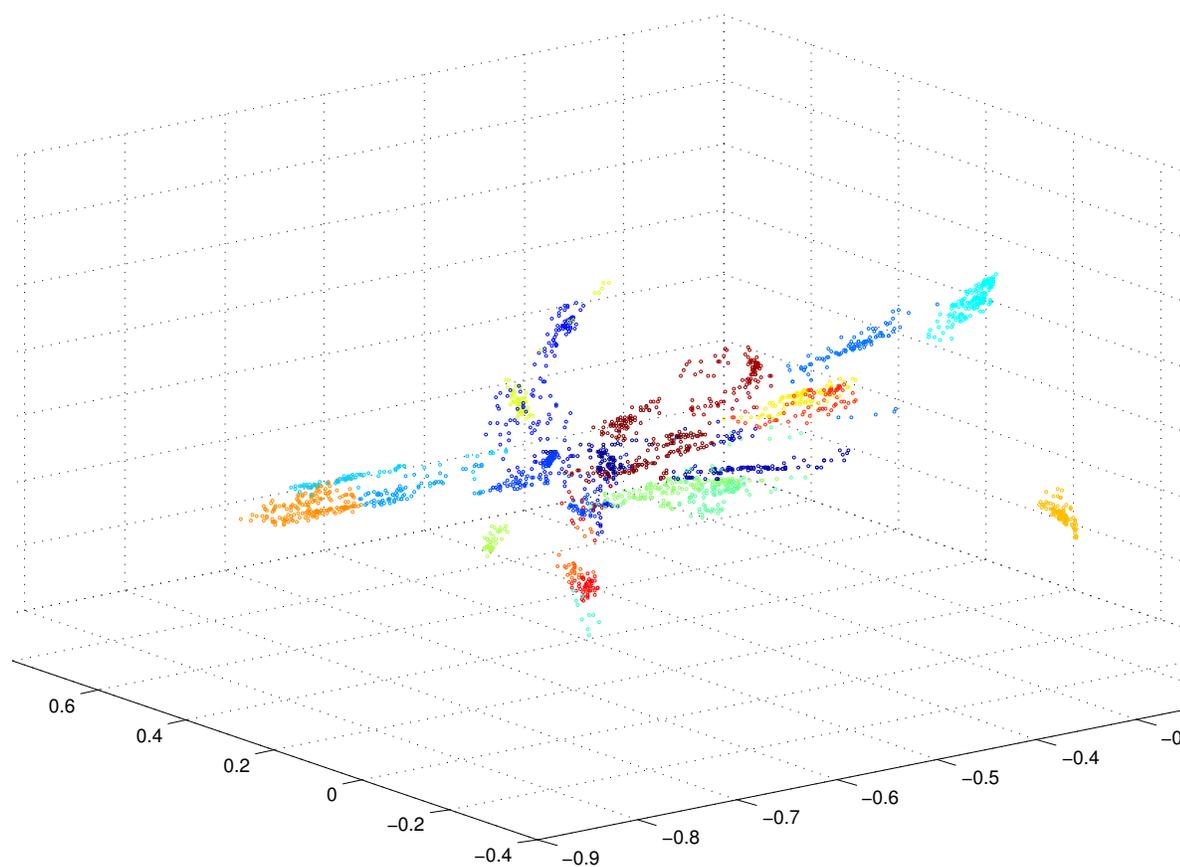


Figure 4.13: Clustering of gestures and their projection on a three-dimensional space using PCA, using labels determined by k-means clustering.

number of salient signals; and that unsupervised gesture classification may be an easier and more effective approach to initialize implicit mapping schemes. Where there exist nonlinearities in the sensor response, it may be advantageous to use more complex and higher-order data reduction techniques, such as nonlinear or kernel PCA [142, 241]. Incorporating temporal information, instead of the static deformations dealt with above, with the use of Hidden Markov Models [131, 148, 286] or recurrent neural networks [108, 109], might help to classify gestures where the data's trajectories provides useful information to unravel the monotonicities inherent in the response of some sensors. As an aside, high-speed acquisition of sensor data could provide differential time-of-flight⁷ measurements, enabling a more precise localization of percussive interaction.

4.3 Discussion

In this chapter we have presented a novel musical interface for real-time control of physical modelling synthesis. The Ballagumi is a flexible silicone interface innervated with fibre optic sensors that can be used to measure its deformation. The sensors are mechanically coupled through the flexible substrate and thus its signals are highly correlated. In order to reduce the computational complexity during mapping, we have suggested that the dimensionality of sensor data be reduced.

The next Ballagumi prototype will have a slightly reduced size and significantly reduced weight, through the use of a flexible low-density foam core. To further reduce player fatigue, increased pliability using a more flexible silicone is planned. To increase sensor robustness, repeatability, and linearity, untreated 0.5 mm closed-loop plastic optical fibre bend sensors, as discussed in Section 6.4.3 and shown in Fig. 6.29, are suggested. Three of these coiled deformation sensors could be combined orthogonally to create a spherical deformation sensor, with each coil being sensitive to deformation a different axis. Spherical deformation sensors, each with three degrees-of-freedom, placed throughout and at different depths within the larger form would provide sufficient data for localizing and classifying deformations, without suffering from problems of reliability and ambient light, and thus reducing the need for re-calibration.

⁷Time-of-flight, in this case, refers to measuring the transit time of an acoustic wave as it travels through the silicone substrate, not to be confused with the time it takes photons to move through the medium.

Along with upgraded acquisition electronics and improved processing power of the *fungible-2* and MR-COSBBE (see Section 3.4), more capable gesture classification and an adaptive control system of physical modelling synthesis are planned to facilitate implicit mapping, where gestures are mapped to sound parameters instead of synthesis parameters, as shown in Fig. 4.14.

For instance, we suggest a “sing-along” mapping system whereby the user deforms the instrument while synchronously vocalizing the sounds that he or she wants or expects to hear. The vocalized audio information is then segmented and analyzed to extract sound features (i.e., amplitude, pitch, and timbral features). A database containing extracted sound features from the physical modelling synthesizer, along with respective synthesis parameters would also be created. The sensor control signals would then be automatically mapped to synthesis parameters that produce sounds with features that most closely match the sound features of the respective vocalizations. Additionally, during performance, an adaptive control system would be used to maintain the synthesized sound features by adjusting the synthesis parameters, as at any one time the output and internal state of a nonlinear dynamical system may not depend only on its current inputs.

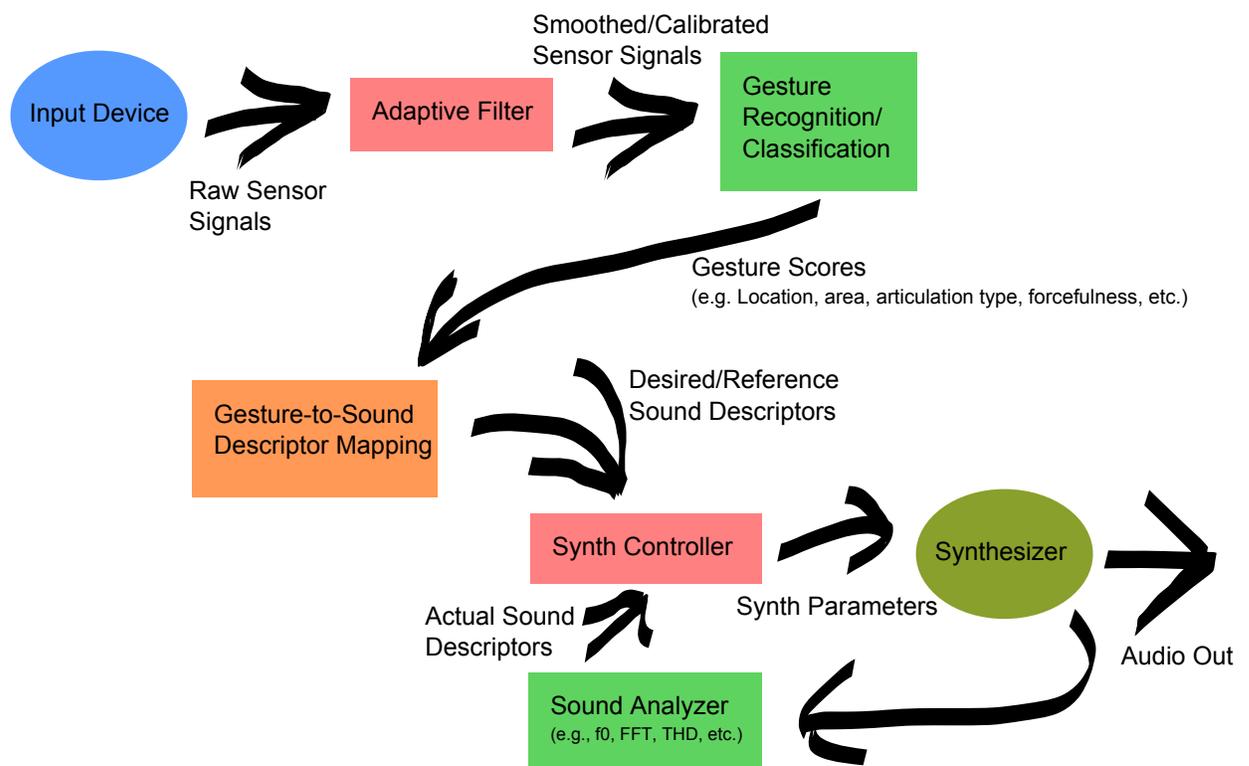


Figure 4.14: Block diagram of an adaptive system capable of mapping gestures onto sound parameters.

Chapter 5

fMRI-Compatible Keyboard

As briefly discussed in Section 2.1.4, we designed an fMRI-compatible electronic piano keyboard used in neuroimaging studies. While the first iteration satisfied the initial criteria of the interface, an improved and more robust clavier was required: one with more keys, better mechanical stability, sturdier optical interconnects, and the ability to capture the keys' trajectories. Presented herein are several design iterations of this device.

5.1 First Generation Keyboard and Limitations

Our first approach employed in the initial keyboard design presented in [113, 116] was to minimize MRI development time and allow the device's use in higher field scanners (i.e., 1.5 T and above) by using only optical sensing and avoiding ferromagnetic materials and electronics within the MRI environment. In aiming to maintain the ecological validity of experiments, weighted keys and the ability to provide real-time auditory feedback were specified. As well, in order to perform more thorough kinematic analysis of piano playing, with applications for motor learning and rehabilitation, a measure of key velocity was sought.

In order to better understand the limitations of the first generation keyboard as it was used in [25], and shown in Fig. 5.1, a brief description follows. As with the initial device design, the eleven key keyboard, modified from an aged upright piano, was de-fanged of all ferromagnetic parts and fitted with pairs of optical fibres: one

pair for each key. A mirror was mounted on the back of each key, and when the key was depressed, the mirror was raised into alignment with the end faces of the optical fibres. With the mirror raised into position, more light was reflected from the emitting fibre into the receiving fibre. The mirrors can be seen on the backs of the keys in Fig. 5.2. Two-level comparators determined the state of each key, whether it was fully-depressed, fully-released, or somewhere in between. The comparator outputs drove interrupts on a microcontroller which then sent messages to the host computer with data including the key identity, transition polarity, velocity, and a time-stamp. Velocity data was extracted from the difference in time between transitions from fully-released to fully-depressed states. The serial communication was changed from the standard MIDI specification for two reasons: the host computer software *Presentation* made by Neurobehavioral Systems, Inc., did not properly recognize MIDI devices; and to increase the baud rate of the communications channel to allow for time-stamp and velocity information, in addition to the usual note-on and note-off messages. The serial data on the microcontroller coming from its UART was converted using an FTDI UART-to-USB chip in order to ensure connectivity on current laptops. Auditory feedback, in the form of sampled piano notes, was presented for each key press.

One major limitation of this implementation was the use of manual trim-potentiometers to set the reference voltages for the comparators. Manual recalibration was occasionally needed when the keyboard was moved between the MRI scanner room and the pilot testing room, due to changes in ambient lighting and poor repeatability of optical coupling with screw-on fibre optic terminals. Manual calibration took up costly scanner time and added an extra variable to the measurement of key transition velocity. As it was, the mechanical differences between keys, as well as the position and alignment of the mirrors, made it difficult to compare velocity measurements across keys and subjects in any absolute sense.

Another issue that arose was the intermittent aliasing of serial communication symbols. That is, the 8-bit field that uniquely identified the key number was being corrupted deterministically and mapped onto an open value that was reserved for future use. Interestingly, this symbol aliasing occurred on only three keys (and much less frequently on only two keys after adding ferrite beads to the USB cable), one of them much more often than the others. It was found that the symbols that were aliasing to other values had more bit-level transitions. This issue was easily overcome in software

by mapping the aliased codes back to their respective symbols.

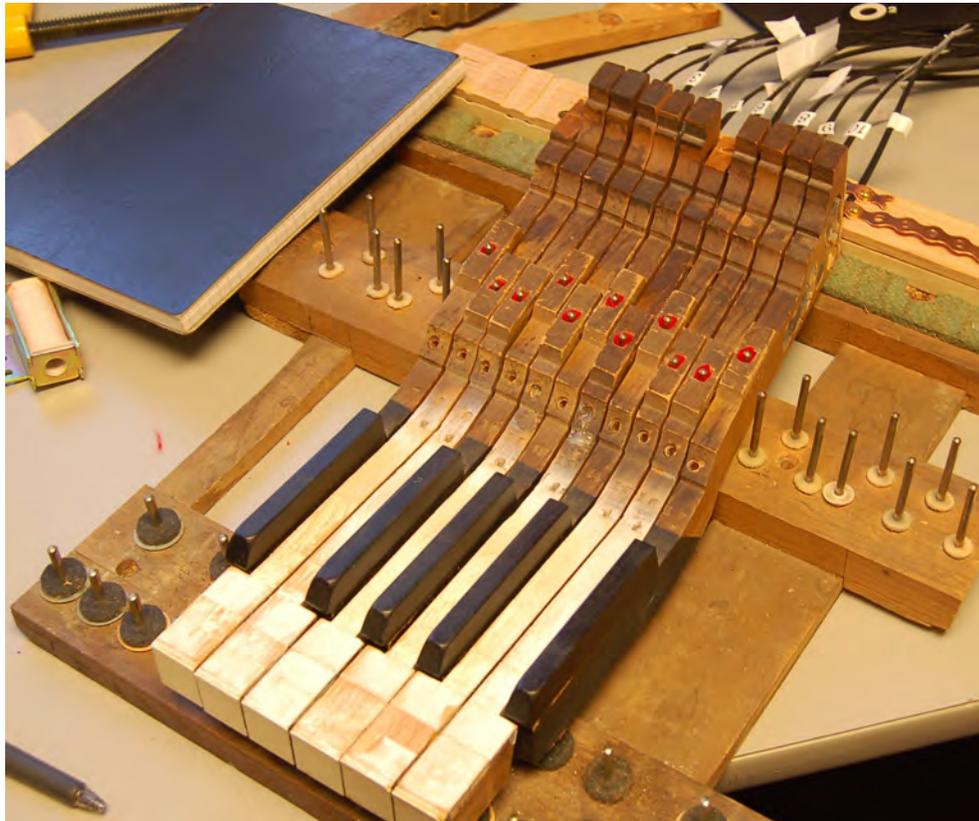


Figure 5.1: First generation piano prototype, presented in [113].



Figure 5.2: First piano prototype with mirrors on backs of keys aligned with emitting and receiving fibres [113].

5.2 First Generation Keyboard with Upgraded Electronics

As a first approach to improve the MRI-compatible piano, the hand-soldered control board was replaced with the *fungible-1* printed circuit board containing the redesigned optoelectronic and control hardware, as discussed in Section 3.3. This board was designed to be used with the piano keyboard and other optically-sensed MRI-compatible electronic interfaces [171]. The board was modular, extensible, re-configurable, and enabled electronic calibration. As well, robust fibre optic connectors allowing repeatable and efficient coupling of light were needed to replace the connector-less screw-terminal housing used previously with the plastic fibre optic cables. Of interest was the ability to measure the position profile of key transitions, which was made more challenging by the very short transition region where the mirror moved into alignment with the optical fibres near the end of the key's travel. The keyboard, in this embodiment, was used for a pilot motor learning study¹ in which only four consecutive white keys of the

¹The multiple finger sequence task (MFST) [228], a variant of the serial reaction time task that had previously been used to investigate learning in non-musicians [229], was used in collaboration with

keyboard were used.

Due to the placement of the mirrors on the backs of the piano keys, a change in amplitude seen by the receiving fibre was only measurable toward the fully-depressed extreme of the key stroke. This transition from almost fully-released to fully-depressed was on the order of 2 *ms*. In order to measure a position profile, a higher sampling rate was required. As the hardware/firmware configuration on the optoelectronic board limited the maximum sampling rate to 1.5 *kHz* per channel (the three-channel ADC running at 3 *kHz* was time-division multiplexed over two acquisition periods to sample the four keys), a 30 *kHz* 8-bit ADC on the PSoC 5 (refer to Section 3.3.2) was used to sequentially sample the position signal of the four keys (i.e., a sample rate of 7.5 *kHz* per channel). The buffered analog output of each key's instrumentation amplifier (with the fibre optic cables of adjacent, unused keys as the reference arms) were wired from the optoelectronic board to the PSoC 5 board. This allowed for an average SNR of 18.5 dB across channels, as the common-mode ambient light was effectively nulled by the instrumentation amp (the SNR of the single-ended amplifier was only about 6 dB), and the programmable gain allowed the dynamic range to be maximized. The key position signals were sent through additional internal PGAs before being multiplexed and digitized. Position data was recorded into a circular buffer of 32 samples per key, along with their respective time-stamps. A state machine driven by a comparator with hysteresis held the state of each key and was used to report only useful (i.e., dynamic) position information to the host PC. As well, the three-threshold comparator scheme properly de-bounced the under-damped piano keys.

Three types of messages were reported by the hub board to the host PC: key onset and release triggers, which included only key ID, a value of 127 or zero, representing an on or off transition, respectively, in the place of the position value, and the time-stamp; a burst of 32 messages sent whenever a key transition was made; and a 1 *kHz* down-sampled signal of the key position whenever a key was not in its fully-released position. As the PSoC 5 USB-UART module had a maximum communication rate of approximately 100 *kBaud*, reporting the full 7.5 *kHz* position information signal to the host PC would have overwhelmed the communication channel and the microcontroller

Christopher Steele and Virginia Penhune to assess learning and performance on different components of learning across multiple days of training. Key press accuracy and timing were measured for a learned sequence interspersed with random sequences.

core. Thus, the burst of 32 samples copied from the position buffer were reported, along with key transition triggers, to give a temporal snapshot of the key's position profile in order to deduce the key transition velocity, while the position signal printed at 1 kHz allowed for analysis of partial key transitions and gestures not obtainable through standard MIDI keyboards.

5.2.1 Calibration and Testing

The calibration procedure for the pilot study was performed in two stages: manual calibration of instrumentation amplifier gains on the *fungible-1* optoelectronic board, and automatic calibration of the state machine comparator thresholds. The gain values were set once and were stored in non-volatile memory and loaded automatically before the start of discontinuous sessions in the same location. Comparator threshold calibration was performed by measuring the minimum and maximum values acquired in both the fully-depressed and the fully-released key positions. In this way, the peak-to-peak noise was implicitly captured and accounted for by setting the fully-released state transition threshold above the noise floor. Calibration of all four keys was done at the same time.

A preliminary development MRI scanning session was performed in co-operation with Christopher Steele and Virginia Penhune at the *Unité de Neuroimagerie Fonctionnelle* (UNF) in Montreal, Canada. During this scan, a measurement of the noise floor indicated no increase in SNR during the scanning.

While the acquired signals do not directly represent the absolute position of the keys, a post-processing normalization and linearization can be applied using the characterization curves. The acquired position signal was compared against the absolute position measured with digital callipers. Fig. 5.3 shows the acquired signal plotted against key displacement. Clearly, there are key-to-key differences in the position of the transition region, as well as differences in their curves. Note that the piano key does not always take the same path and the key can actually lift up at the fulcrum when the key is struck more firmly. Thus the signal and its curve may not be consistent. A mass of around 35 g at its end was needed to fully depress a key.

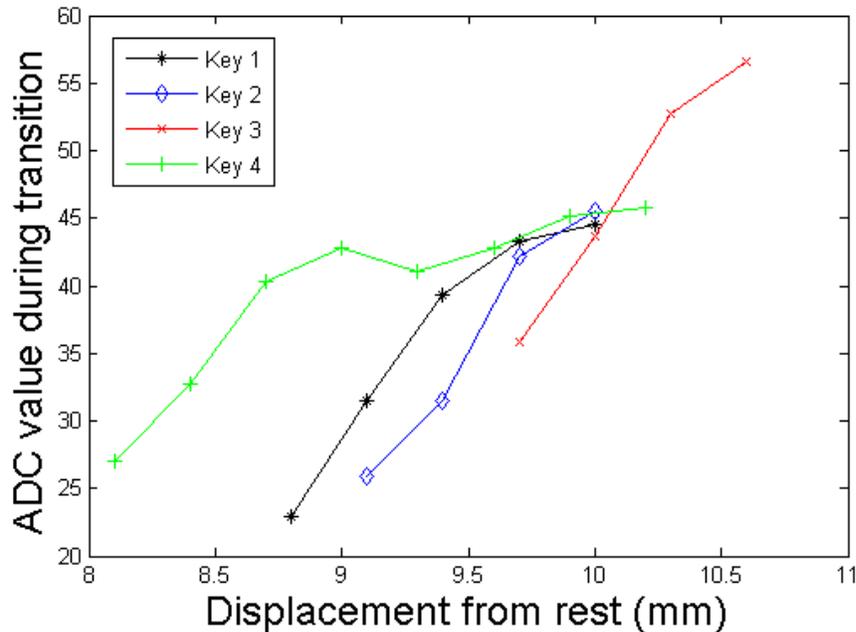


Figure 5.3: Characterization curve of 1st-generation keyboard, key position in transition region.

5.3 Second Generation Keyboard

A redesigned mechanical keyboard was required to overcome the optomechanical key-to-key differences of the first prototype. As with the first prototype, the design constraints included weighted keys to maintain the feel of an acoustic piano keyboard in the absence of a mechanical piano *action*. The new-generation keyboard was to be upgraded with position sensing over the whole range of the key's travel, from resting to fully-depressed, with a linear or easily linearized output for conversion into absolute physical measurements. Proper alignment of the optomechanical system was to be engineered into the design. As well, in order to reduce costs and ensure MRI-compatibility, electronic and ferromagnetic parts on any portion of the device entering the MRI environment were prohibited. Non-ferromagnetic metallic parts were to be kept to a minimum and were to remain outside of the imaging volume.

5.3.1 Mechanical Keyboard

Instead of machining a prototype from the bottom up or modifying an existing acoustic piano, it was felt that modifying an existing electronic weighted-key commercial piano keyboard would provide a compact, robust, and tested mechanical structure on which to base the interface. A Yamaha P85 electronic weighted-key piano was graciously donated by Yamaha Canada Music Ltd., as seen in Fig. 5.4.² The keys and key bed were removed from the case and stripped of all electronics and ferromagnetic components (i.e., screws and steel weighting beams), shown in Fig. 5.5. The 88 key keyboard was cut into three keyboards. Due to issues of mechanical stability and alignment of mounting holes, two keyboards spanned a full two octaves, while the third keyboard had only 22 keys. The three keyboards were identical in all other design respects. The new keyboard prototype consisted of four mechanical sub-assemblies: the key bed, the fibre optic mounting fixture, the augmented keys, and the weighting beams terminated with plastic linkages.



Figure 5.4: Yamaha P85 electronic weighted-key piano.

A fibre optic mounting fixture was designed, as shown in Fig. 5.6, and machined out of Delrin. This fixture fit directly onto the pre-existing mounting holes in the key bed, previously used for securing the key-sensing circuit boards. The fixture had slotted mounting holes in order to adjust the distance between the fixture and the keys. Aligned with each key were three holes for mounting the 1 mm diameter POF cables.

²Many thanks to Yamaha Canada, and specifically Ben Harrison for securing us the piano and for providing detailed support whenever needed.



Figure 5.5: Yamaha electronic piano internals, with ferromagnetic counter-weight beams, screws, and electronic circuit board.

For each key the emitting fibre was placed in the middle hole, with receiving fibres placed on either side. A groove along the top side of the mounting fixture was cut to allow for a bead of adhesive to fix the fibres in place within the fixture. Additionally, the fixture had holes to accommodate a photomask and/or lenses mounted to brass posts which could be moved toward or away from the fibre optic cables. These mounting-posts were made of brass and attached the fibre optic mounting fixture to a transparent polycarbonate beam onto which a photomask could be placed.

The piano keys were augmented with mirrors secured with adhesive. A jig was fashioned in order to properly position and align the mirrors during glue-up. The mirrors on the underside of the keys faced the ends of the fibre optic cables, with the middle fibres aligned with each mirror's centre. The mirrors were cut to dimensions which matched the widths of the inside of the hollow plastic keys to ease mounting. The mirrors were cut to lengths as long as possible so as to reflect light back to the optical fibres throughout the entire travel of the keys, but not so long as to collide with other internal parts, specifically the weighting beams with the keys depressed. Onto

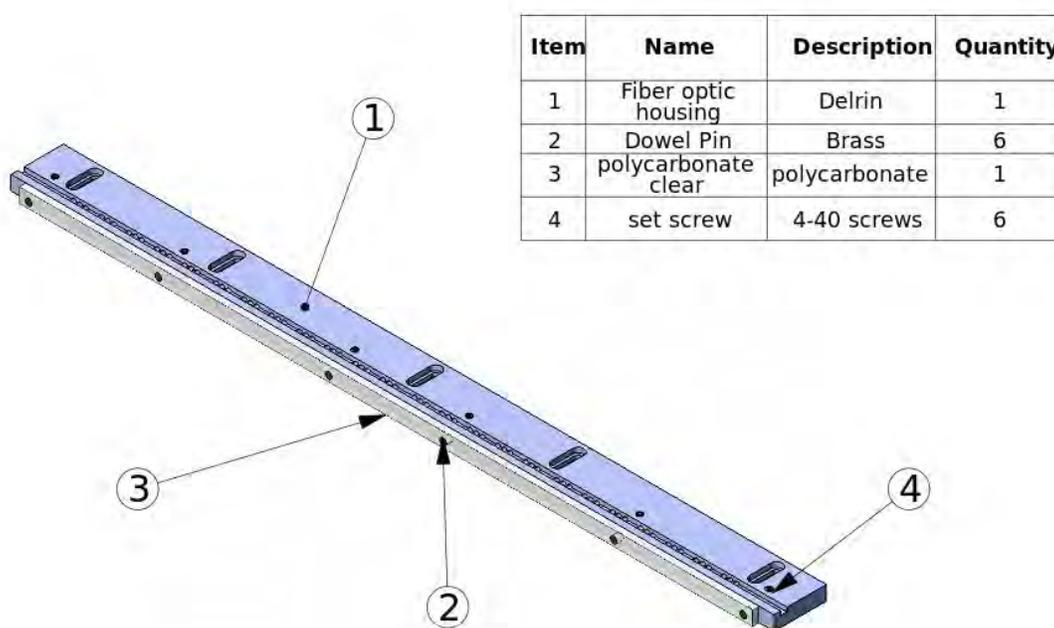


Figure 5.6: CAD drawing of optical fibre mount.

the mirrors were placed photomasks as explained in Section 5.4. The adhesive used to secure the mirrors in place also increased the inertia of the keys, which made for a more realistic feel. A mass of around 70 g at its end was needed to fully depress a key.

The ferromagnetic weighting beams were removed and replaced with brass beams cut to a length which approximated the moment of inertia of the original lighter treble keys.³ The brass beams were secured into the original plastic linkages which attached to the key bed and coupled the weighting beams to the keys. The plastic linkages were filed down slightly to accommodate the mirrors when the keys were fully-depressed. The piano prototype with optical fibre mount, fibres, and mirror is shown in Fig. 5.7.

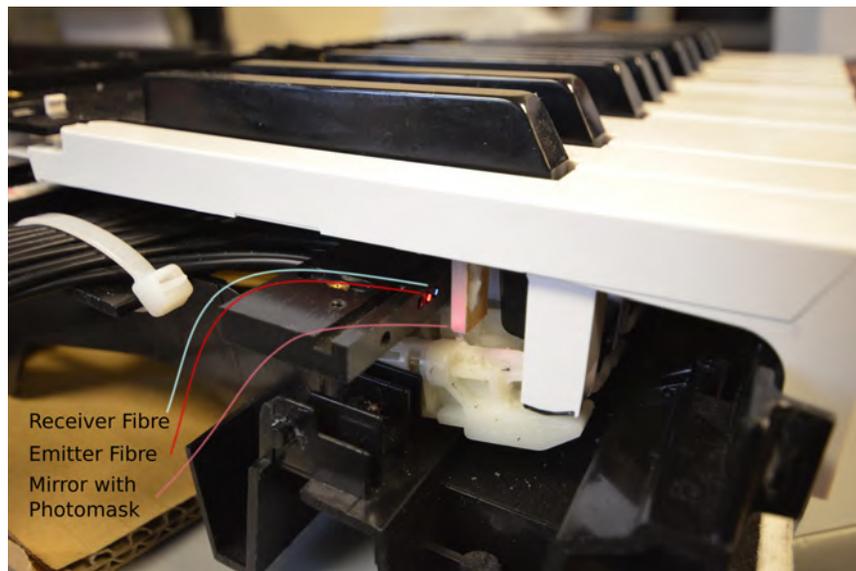


Figure 5.7: Annotated photograph of second keyboard prototype with optical fibres and mirror.

³The initial weighting beams increased in weight as one moved from the treble end of the keyboard down to the lower registers to mimic the feel of heavier hammers of an acoustic piano. This was accomplished by Yamaha by varying the lengths, and thus the effective weight of the steel beams. As the fMRI-compatible keyboard was designed for use in neuroimaging studies, the confounding factor of varying key weights was removed by replacing the steel weighting beams with brass beams of uniform length and weight.

5.4 Sensing Methods

Several optical key position sensing methods were proposed based on simple light intensity measurements: key angle modulation, optical quadrature encoding, shifted parallel gratings producing differential pairs of Moiré patterns, the use of triangular linear photomasks, and continuously-variable neutral density photomasks. There are many other possible sensing methods, including interference gratings and the use of polarization-dependent amplitude modulation, but these were the simplest techniques to approach during the initial design iterations.

As the trajectory of the piano keys was an angular movement about a fulcrum, the distance and angle between the mirror and fibre optic end face varied slightly. Thus as the key is deflected, a monotonic change in light intensity can be measured. A mapping onto a linear position measurement is a simple tangent function. Single-ended measurements such as this are prone to error due to fluctuations in ambient light and photo-emitter power. As well the SNR was higher for the black keys as they had a greater change in angle throughout key travel.

Linear quadrature encoding was implemented using optical gratings consisting of two columns of alternating dark and transparent bands printed on acetate and attached to the mirrors of each key. The dimension of the photomask matched that of the mirror. The two gratings were offset by 90 degrees,⁴ as shown in the first column of Fig. 5.8. The central fibre, aligned in the middle of the two gratings, illuminated the photomask and mirror. The receiving fibres were aligned with their respective grating columns. As the gratings were printed with a known pitch, decoding the quadrature signals using an embedded quadrature decoder block allowed for absolute displacement measurements by counting the number of transitions and using the phase difference between channels to determine directional information. As quadrature encoding of this sort discretizes the positional information, additional sampling of the input signals to provide intra-grating positional estimation was proposed. A possible confound for this method included the angle-dependent amplitude of the optical signal, which would cause transitions to be sensed leading or lagging the true position if the decoder thresholds were static. State-adaptive thresholds could overcome this issue with calibration. Another

⁴This is a linear spatial shift resulting in a phase lead (or lag) of the left-side signal, ahead of (or behind) the right-side signal, depending on the direction of motion.

way to overcome this issue is to use a chirped grating, with the change in grating pitch engineered to match the sigmoidal angle-dependent amplitude shift.

Printed photomasks were used to provide flexibility and minimize costs and time during prototyping. The optical quadrature encoding method was constrained as the opacity and minimum grating pitch using a laser-jet printer was limited. In order to increase the spatial resolution of the sensing system, an approach using two photomasks was tested. One photomask placed on the moving mirror was the same as before, with two columns of shifted gratings, while a second photomask, with another set of gratings of the same pitch, was flipped and placed on the polycarbonate strip between the mirror and the fibre optic cables. The second photomask mounted on the movable polycarbonate strip was positioned and fixed as close to the mirror as possible without colliding. When the key was displaced, an oscillating Moiré pattern was seen at the receiving fibres, with a 90 degree phase shift between the two in order to discern the direction of movement. This phenomenon can be seen in the second column of Fig. 5.8. The same angle-dependent amplitude modulation issues apply to this method as seen in the previous optical quadrature encoding scheme, however, the analog nature of the output signals provide improved intra-grating resolution. Another advantage of this method is that it is less sensitive to the spatial positioning, alignment, and focus of the receiver, in comparison to the optical quadrature method. A similar scheme using gratings with pitch of comparable size to the optical wavelength would create actual interference patterns.

In order to obviate the angle-dependent amplitude modulation, a linear differential sensing scheme which nulled this confound as part of the common-mode signal was developed. For a given key, a photomask was designed which consisted of two vertical sections as with the quadrature encoding scheme, one aligned with each receiving fibre. Each section of the photomask was divided along the diagonal, and half was printed dark while the other half was kept transparent. Note that the photomask is rotation-symmetric with opaque portions at opposite ends of each section. This configuration produces linearly proportional changes in light intensity for each optical receiver, and as the key is moved in a given direction the signal rises in one receiver and falls in the other, as shown in the third column of Fig. 5.8.

In addition to the triangular photomasks, a continuous gradient photomask (also known as a continuously-variable neutral density filter) printed on acetate was tested.

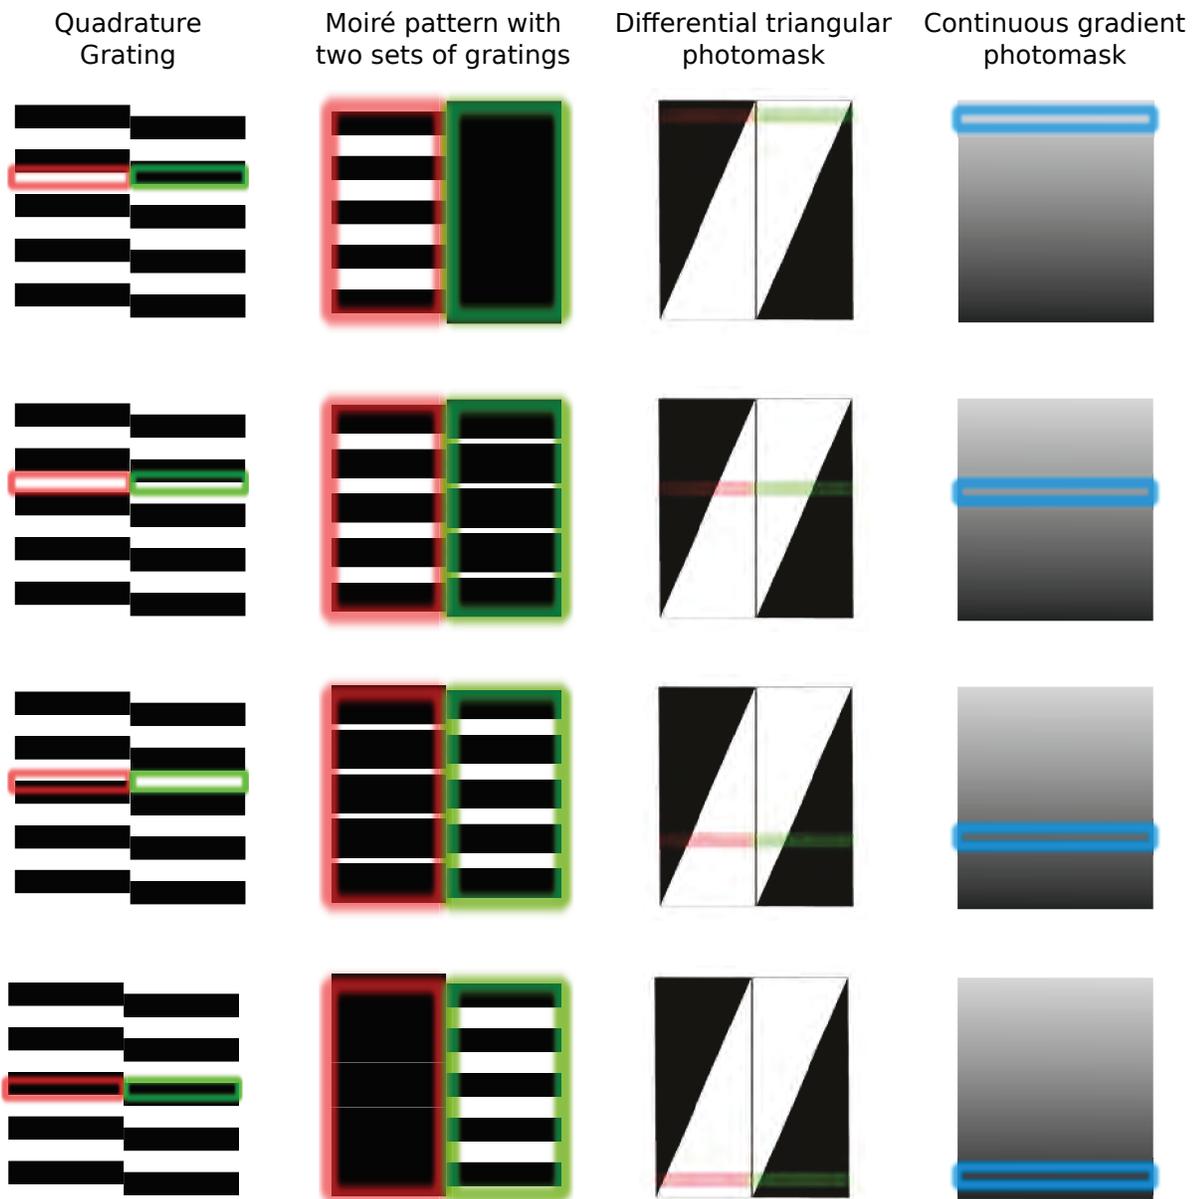


Figure 5.8: Photomasks to measure piano key displacement: quadrature sensing with separate, shifted photomasks; a differential Moiré pattern created by two superimposed sets of quadrature gratings, one static, the other flipped and movable; differential triangular photomask; continuously-variable gradient photomask. Contrast of highlighted areas proportional to light received by photodiodes. Each row shows contrast as the photomask is moved with respect to the sensor (or in the case of the Moiré pattern sensor, one photomask is moved with respect to the other).

A single, as opposed to differential, gradient was used to reduce cross-talk and nonlinearities due to misalignment. A linear increase in opacity was employed, as seen in the fourth column of Fig. 5.8, though other contrast curves were possible.

5.4.1 Testing of Sensing Methods

The various sensing methods described in the previous section were tested in the laboratory using analog measurements made with an oscilloscope at the output of transimpedance and instrumentation amplifiers, as well as data collection using the acquisition system described previously. Position sensing using angle-dependent modulation resulted in measurable shifts in light intensity, especially for the black keys. For the white keys of length of 107 mm (from fulcrum to mirror) the excursion at the mirror was about 6.3 mm, while the black keys with a 60 mm effective length had an excursion of 8.6 mm. Thus the calculated change in key angles of 3.4 and 8.2 degrees, for the white and black keys respectively, resulted in intensity modulations of 17.6 mV and 35.2 mV.

As expected, the linear optical quadrature sensing suffered from angle-dependent signal modulation. For a given white key and a 1 mm grating pitch, the average dynamic range across channels was initially 12.4 dB, but was reduced to 10.4 dB as the key was depressed. For a quick keystroke, the 90 degree phase shift between channels was seen as small as a 380 μ s temporal shift, and the full keystroke lasted around 22.2 ms, giving a little more than six full amplitude cycles. Fig. 5.9 shows a plot of the quadrature signals for a single key during transition. Substituting a red laser diode for the LED, the SNR was increased to 26.6 dB, with an angle-dependent degradation in SNR of -3.1 dB seen at the end of the key's travel.

The Moiré pattern technique suffered from low SNR compared to the quadrature method, with significant losses due to multiple reflections at the air-polycarbonate interface and the two layers of acetate.

The differential triangular photomask method performed better. SNR of the differential signal was measured at 39.5 dB, with average single-ended SNR across channels of 16.3 dB. Photomask misalignment was the one significant issue, causing what should have been a linear function of position to become a three-part piece-wise linear function, due to saturation of single-ended measurements at the extremes of the key's travel

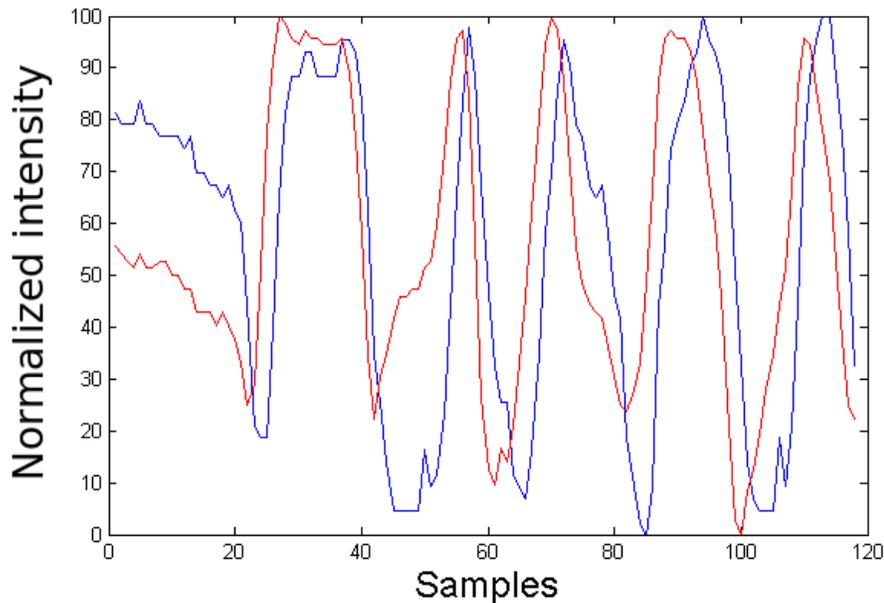


Figure 5.9: Two-channel optical quadrature signal for a single key during transition (1.5 kHz sampling rate).

(see Fig. 5.10).

Better still was the continuously-variable gradient photomask with an average SNR across 8 white keys of 53.8 dB (with a range of 48.8–57.1 dB). The excursion at the end of the key was measured at 11.5 mm. A plot of the normalized signals elicited from the eight key sensors as a function of key displacement is shown in Fig. 5.11.

In order to assess the accuracy and linearity of the key sensors, a Polhemus Liberty electromagnetic motion tracking system [205] was used as a reference measurement system (240 Hz sampling rate, 0.76 mm root-mean-square static position accuracy). A white and a black key were fitted at their tips with Polhemus tracking markers. The keys were depressed, separately and repeatedly, and the Polhemus position measurements were acquired simultaneously with the optical sensor data using the continuously-variable gradient photomasks and a *fungible-2* optoelectronic acquisition board. Figures 5.11 and 5.12 show time traces of the normalized key position and optical sensor data⁵ with white and black keys, respectively. Scatter plots showing the

⁵The optical key sensors, located within the keyboard, sensed a different, though proportional, displacement than that measured with the Polhemus motion tracking system. Thus all key press data is

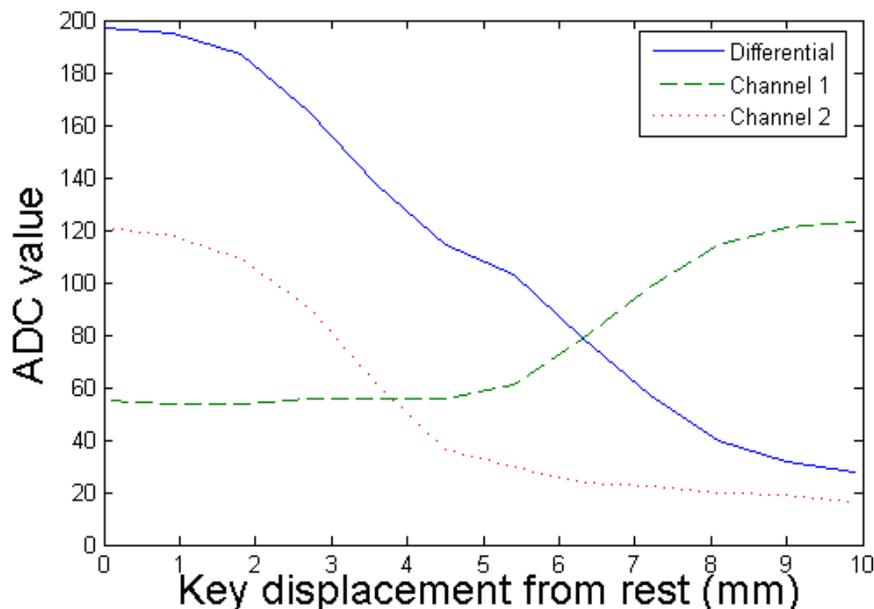


Figure 5.10: Triangular differential photomask sensing of a key moved throughout its full travel, differential and single-ended characterization curves.

optical sensor nonlinearities are shown in Figures 5.14 and 5.15. By using calibration curves for each key as a lookup table, the optical sensor data may be linearized in order to provide a more accurate measurement of position—specifically during key transition, from which velocity and acceleration may be extracted—as shown in Fig. 5.16, where a black key was pressed once and released. The linearized sensor data improved the root-mean-square error for the white key from 5.1% to 1.6%, and for the black key and from 4.1% to 2.0%. The linearized scatter plot for the black key is shown in Fig. 5.17.

The nonlinearities arose from variability in photomask printing and opacity of the adhesive used to adhere to photomasks to the reflectors. Key bouncing at fully-depressed and fully-released positions is responsible for most of the error in the linearized optical sensor data. Additional sources of error, specifically in the reference measurement, can be attributed to movement of the key bed with respect to the Polhemus electromagnetic emitter as well as distortions in the electromagnetic field caused given in a normalized form between fully-released (a value of 0.0) and fully-depressed (a value of 1.0) positions.

by the moving brass counterweights attached to each key. It is important to correct the nonlinearities of the key position values in order to properly estimate the instantaneous velocity of the key as it is depressed.

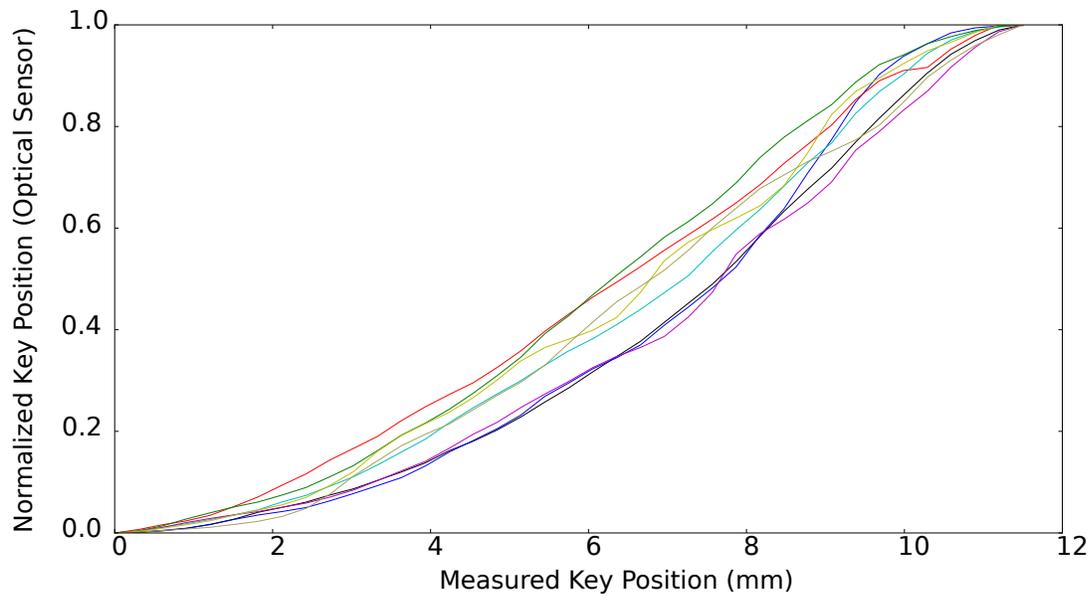


Figure 5.11: Continuously-variable gradient photomask: normalized sensor values with key displacement (in *mm*) for eight white keys.

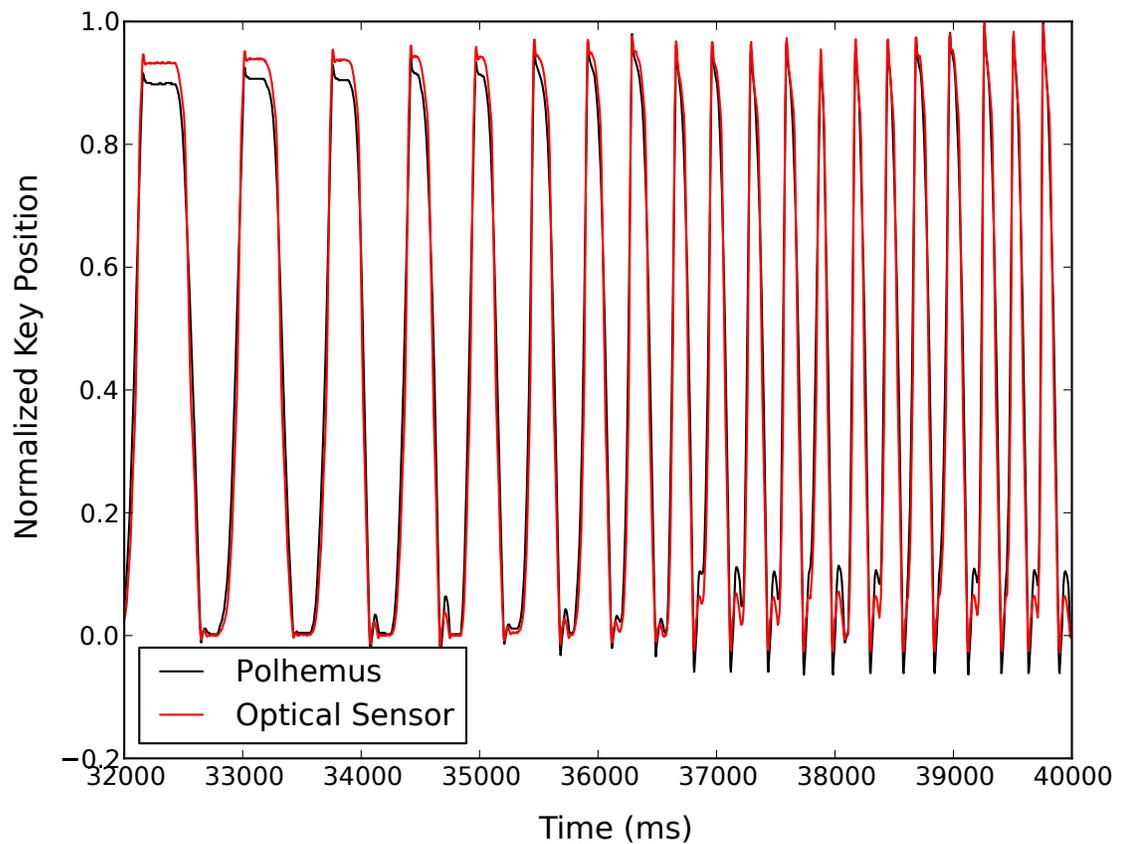


Figure 5.12: Repeated key presses of a white key measured with Polhemus motion tracking system (normalized vertical key displacement) and optical sensor with continuously variable photomask.

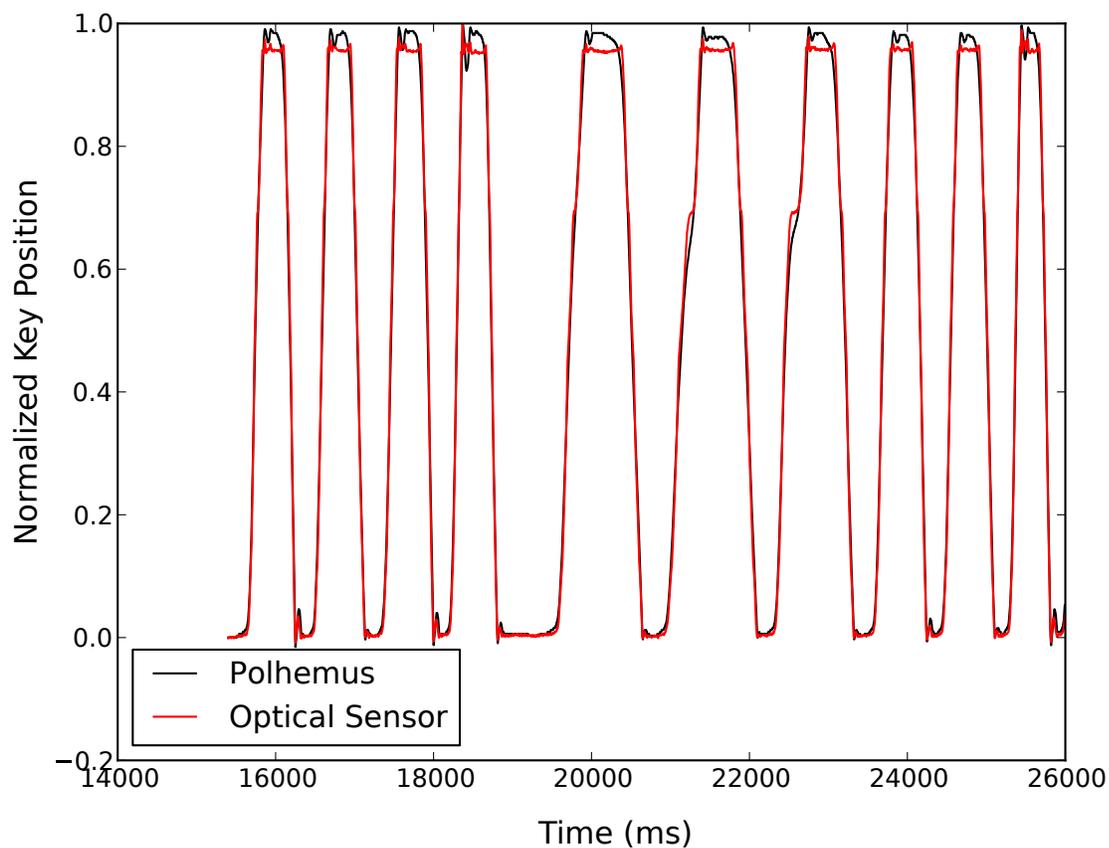


Figure 5.13: Repeated key presses of a black key measured with Polhemus motion tracking system (normalized vertical key displacement) and optical sensor with continuously variable photomask.

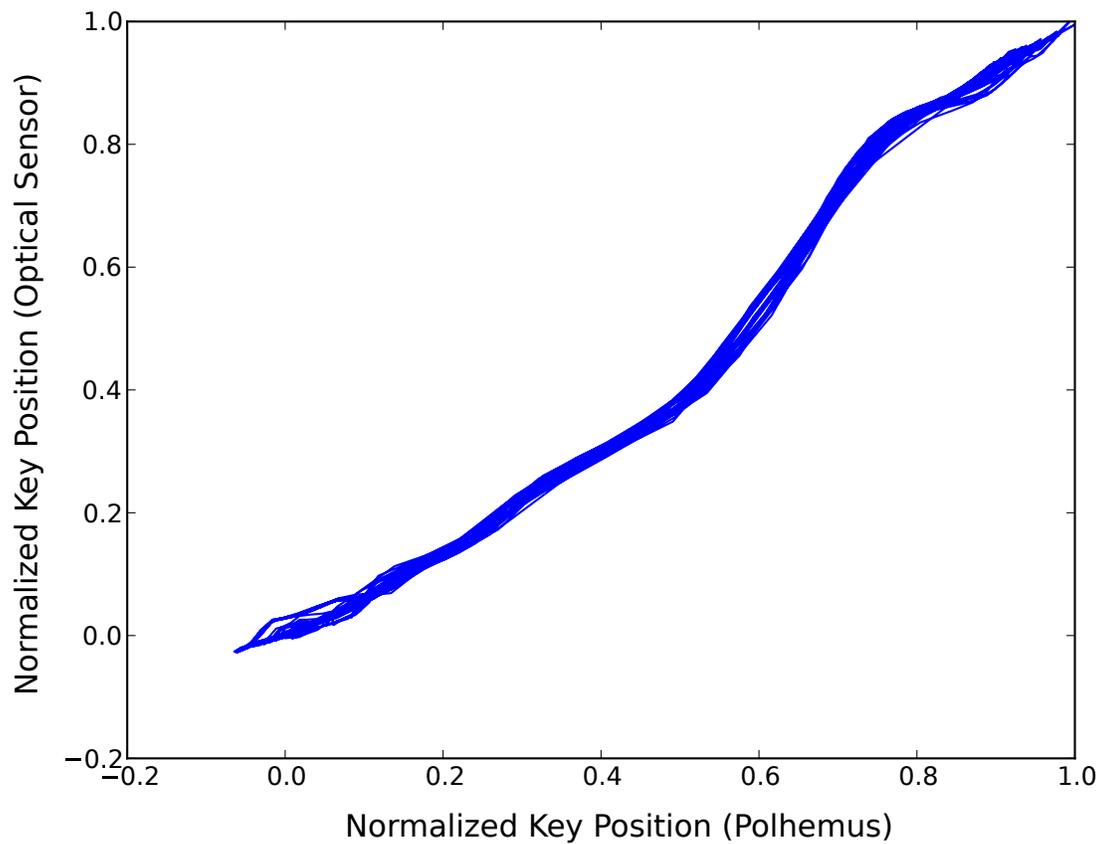


Figure 5.14: Scatter plot of repeated key presses of a white key measured with Polhemus motion tracking system (normalized vertical key displacement) and optical sensor with continuously variable photomask.

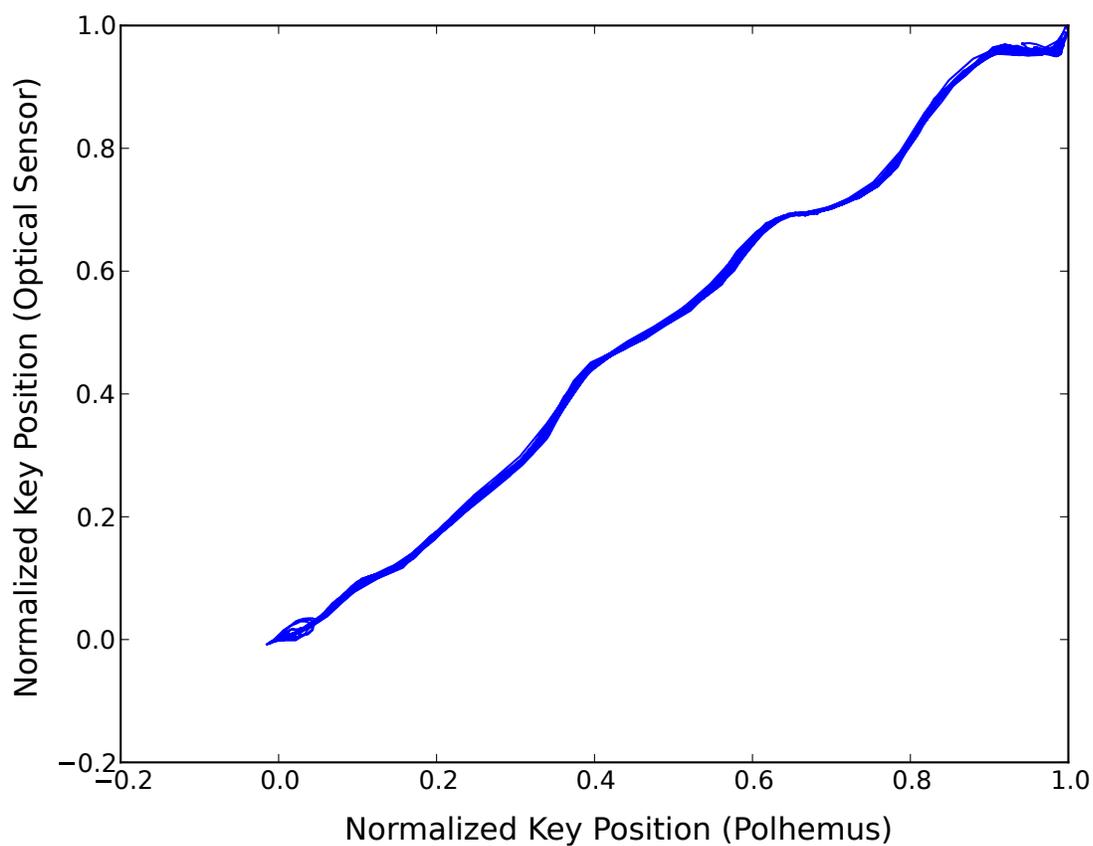


Figure 5.15: Scatter plot of repeated key presses of a black key measured with Polhemus motion tracking system (normalized vertical key displacement) and optical sensor with continuously variable photomask.

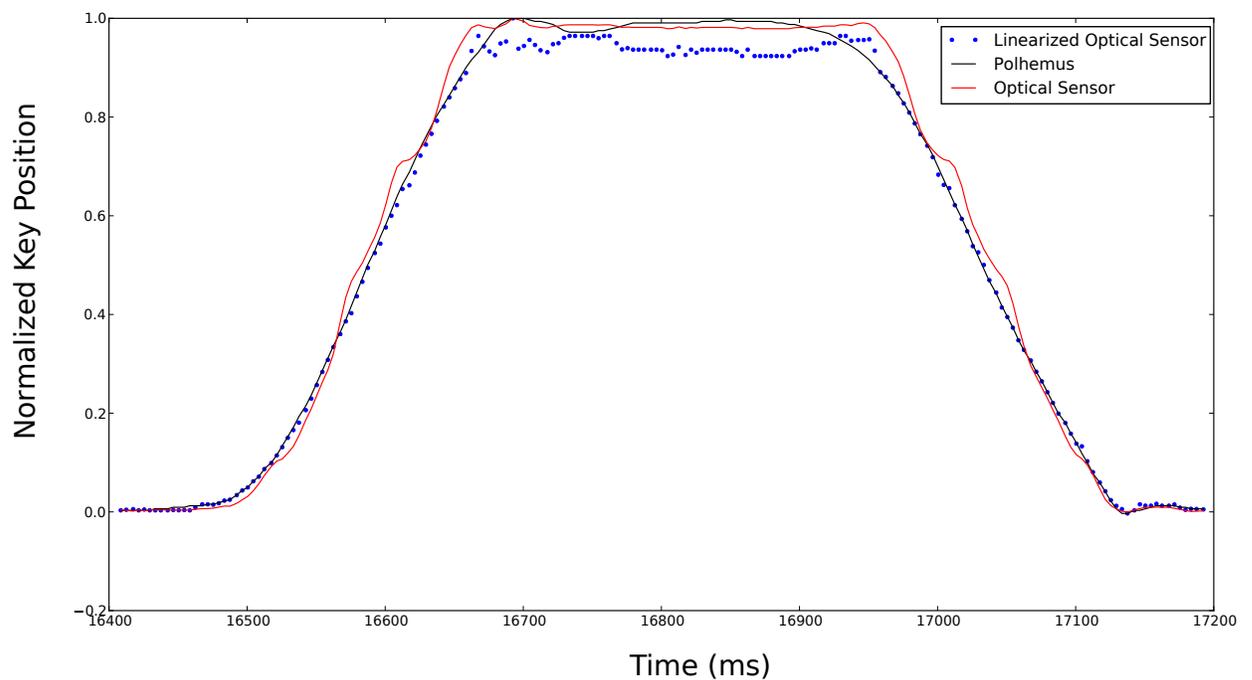


Figure 5.16: Single key press of a black key with release, measured with Polhemus motion tracking system (normalized vertical key displacement), optical sensor with continuously variable photomask, and with linearized optical sensor values.

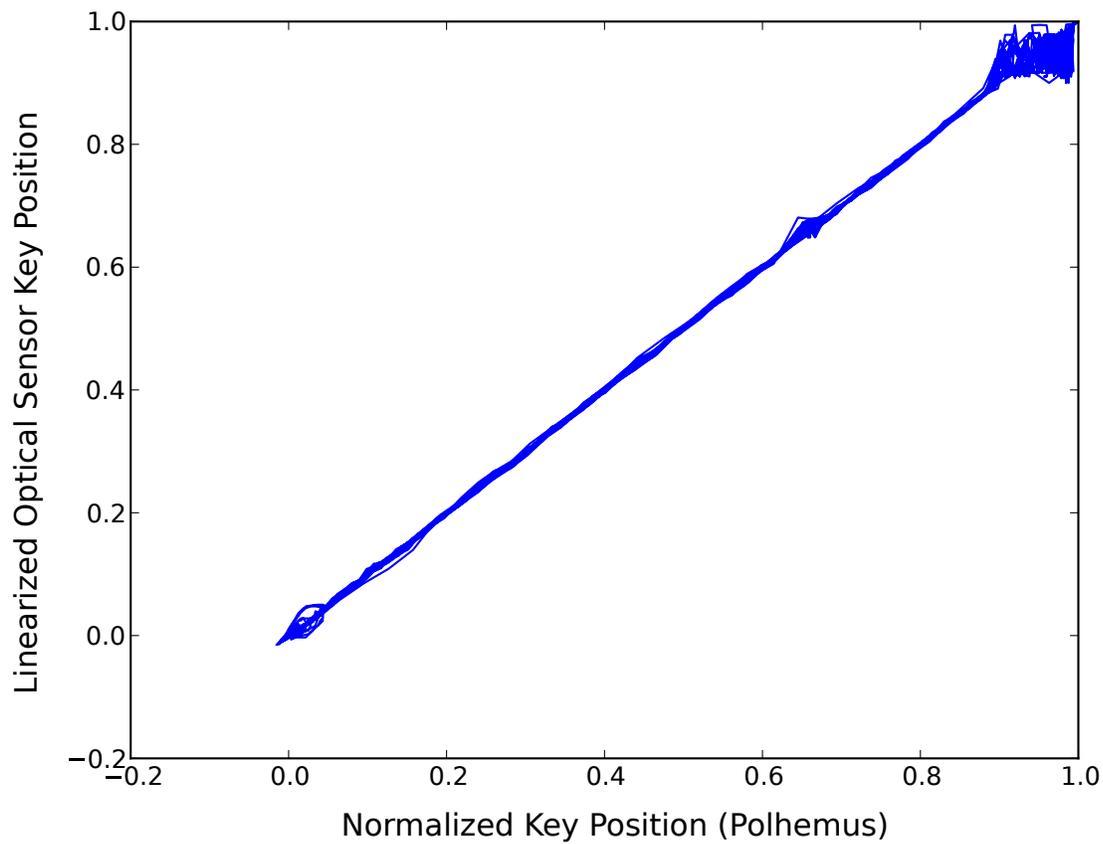


Figure 5.17: Scatter plot of repeated key presses of a black key measured with Polhemus motion tracking system (normalized vertical key displacement) and the linearized optical sensor with continuously variable photomask.

5.5 MRI Testing

Field homogeneity was analysed as a measure of fMRI-compatibility. SNR of fMRI scans is dependent on several factors including the task, scanner acquisition sequence parameters, and field homogeneity [127, 128, 139, 192]. Large distortions in the magnetic field cause spatial and intensity image artifacts, while also affecting the SNR of fMRI. In order to assess field homogeneity, field mapping scans were performed with and without the keyboard in the scanner, positioned at arm's length from the isocentre. An MRI phantom ($NiSO_4 + NaCl + H_2O$) was positioned in the head coil. Each field mapping scan consisted of two gradient-echo pulse sequences ($TR = 1000\text{ ms}$, $TE_1 = 2.99\text{ ms}$, $TE_2 = 5.45\text{ ms}$) using a Siemens Magnetom Trio 3 T scanner, which generated two magnitude images and a phase (or difference) map. The images were masked to include a central region of interest. Static (i.e., B_0) field homogeneity was calculated by taking the difference in phase between the central voxel and the largest outlier which resulted in maximum phase difference within the masked image, $\delta\phi$. The field homogeneity was calculated, as described by Och et al. [192], as

$$\frac{\delta B_0}{B_0} = \frac{\delta\phi(T E_1 - T E_2)}{(\gamma B_0)}$$

where $\gamma = 267.513\text{ rad/s/T}$ is the gyromagnetic ratio for protons (i.e., H^+).

5.5.1 First Generation Keyboard

The first version of the piano keyboard was tested for MRI compatibility by comparing magnetic field maps of a spherical MRI phantom with and without the keyboard. With the keyboard at arm's length from the isocentre, no field distortion or imaging artifacts were visible. Table 5.1 summarizes the results of MRI compatibility testing. Siemens, the MRI scanner manufacturer specifies a field homogeneity of 0.1 ppm [243], which is greater than the change seen with and without the piano keyboard. As well, there are no significant differences in mean image intensity with and without the piano, nor is there a meaningful change in the voxel-by-voxel subtraction of the magnitude images. Thus we can conclude that the piano is MRI-compatible. The MRI scans of the phantom with and without the keyboard are shown in Fig. 5.18. As mentioned previously, this

keyboard was used for an MRI study [25], as well a pilot study of the MFST experiment.

Table 5.1 MRI compatibility of first generation piano keyboard.

| Description | Value |
|--|------------------|
| Maximum change in B_0 field homogeneity ($\frac{\delta B_0}{B_0}$) | 0.038 <i>ppm</i> |
| Larmor (resonant) frequency shift | 1.6 <i>Hz</i> |
| Number of voxels | 11378 voxels |
| Mean image intensity without device | 4006.0 |
| Change in mean intensity with device | 5 |
| Mean voxel-by-voxel percent change | 0.19% |

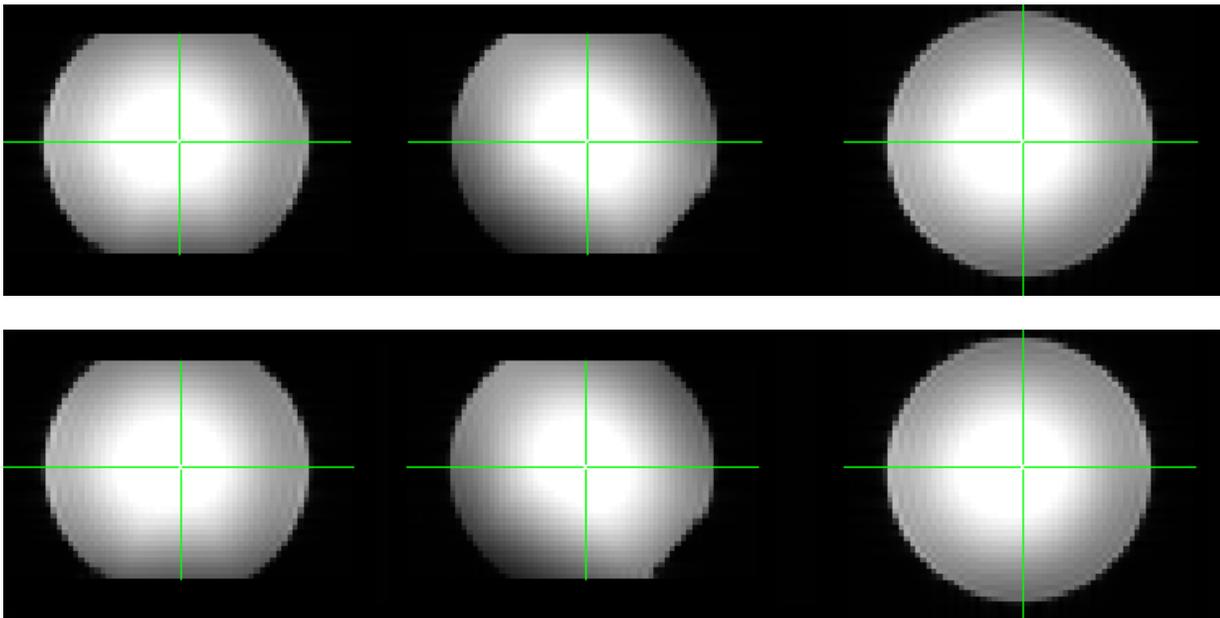


Figure 5.18: Gradient field map of MRI phantom (magnitude images): (*top*) without keyboard; and (*bottom*) 1st-generation keyboard at arm's length from isocentre. No image artifacts are visible.

5.5.2 Second Generation Keyboard

The second revision of the piano keyboard was tested within the MRI as part of a study conducted in collaboration with Christopher Steele and Virginia Penhune using the MFST paradigm. This task included visual stimuli with synchronized motor responses for a learned sequence interspersed with random sequences. No auditory or

visual feedback was presented when any of the four consecutive white keys were depressed. Key press accuracy and timing were analyzed as the behavioural metrics. No image artifacts were found and EMI did not disrupt the proper operation of the electronics of the *fungible-1* optoelectronic acquisition system. The *fungible-1* optoelectronic acquisition system made use of an I2C link between the optoelectronic board and the hub board, which was then connected to a laptop PC as a emulated serial (i.e., virtual COM port) device. MRI synchronization pulses, emanating from the MRI scanner, were captured electronically on the *fungible-1* and sent along with piano key messages. A time-stamp generated on the embedded system, as well as another on the host PC, were used to maintain synchrony within the experiment. A (15,11) Reed-Solomon forward error-correcting code [215] was applied to outgoing serial data. For this study the differential triangular photomasks were employed, as this was thought to provide the most sensitive output. It was found that the optical fibres were not sufficiently glued in place and with movement of the keyboard to and from the scanner room, the fibres were shifted in place. This lead to drastic changes in the output signal, as a small optical misalignment due to a mechanical shift was amplified, in some cases, resulting in saturation of the sensor. This problem was overcome by improved fixation of the optical fibres within the optical mount and strain relief of the fibres as they exited the device. No image artifacts were found.

Another iteration of the piano keyboard, which made use of the continuous gradient photomasks, was used in a neuroimaging study conducted in collaboration with Michael Klein and Robert Zatorre. The keyboard can be seen on its acrylic support in Fig. 5.19. This task involved listening to and performance of major triads. The *play* task was performed with and without auditory feedback, over a range of six consecutive white keys. The *fungible-2* optoelectronic acquisition system was used in HID-keyboard mode, in order to provide key-press messages which triggered feedback in a fashion compatible with the experimental software written in *Python*. High-speed acquisition of key presses was logged on SD cards on each board. No image artifacts were visible and no disruption of the acquisition electronics due to EMI were noted. Fig. 5.20 shows the signals of two key sensors as their respective keys were depressed. A lower signal corresponds to a depressed key. The initial spike and ringing is caused by the mechanically under-damped bounce of the keys.

MRI compatibility was quantified by comparing magnetic field maps of a cylindrical



Figure 5.19: MRI-compatible keyboard in MRI scanner.

MRI phantom with and without the piano (placed at arm's length from the isocentre).⁶ The magnitude and phase images of the phantom, without and with the piano, are shown in Figures 5.21 and 5.22; no differences are visible. The results are summarized in Table 5.2. As expected, the field homogeneity, specified by Siemens as 0.1 *ppm* [243], is not disturbed by the presence of the piano in the scanner. Thus, we can conclude that the piano is MRI compatible.

Table 5.2 MRI compatibility of second generation piano keyboard.

| Description | Value |
|--|------------------|
| Maximum change in B_0 field homogeneity ($\frac{\delta B_0}{B_0}$) | 0.008 <i>ppm</i> |
| Larmor (resonant) frequency shift | 1.2 <i>Hz</i> |
| Number of voxels | 4819 voxels |
| Mean image intensity without device | 1604.4 |
| Change in mean intensity with device | 2.1 |
| Mean voxel-by-voxel percent change | 0.26% |

⁶A spherical phantom would have been preferable due to its symmetry and larger volume, though it was unavailable for this session. Thus a smaller image mask of 4819 voxels was used to exclude the edges of the phantom container, thereby avoiding regions with higher variability in magnetic susceptibility.

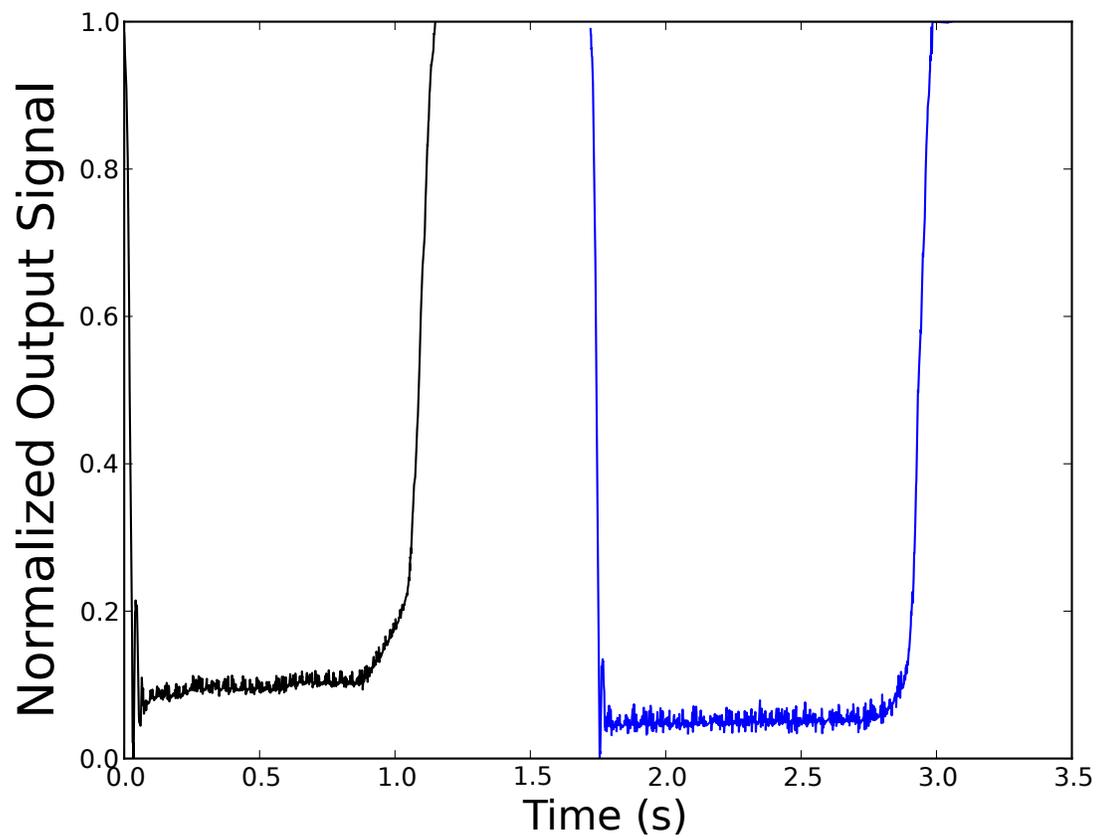


Figure 5.20: Two keys played in succession using continuous gradient photomasks, 1 *kHz* sampling rate.

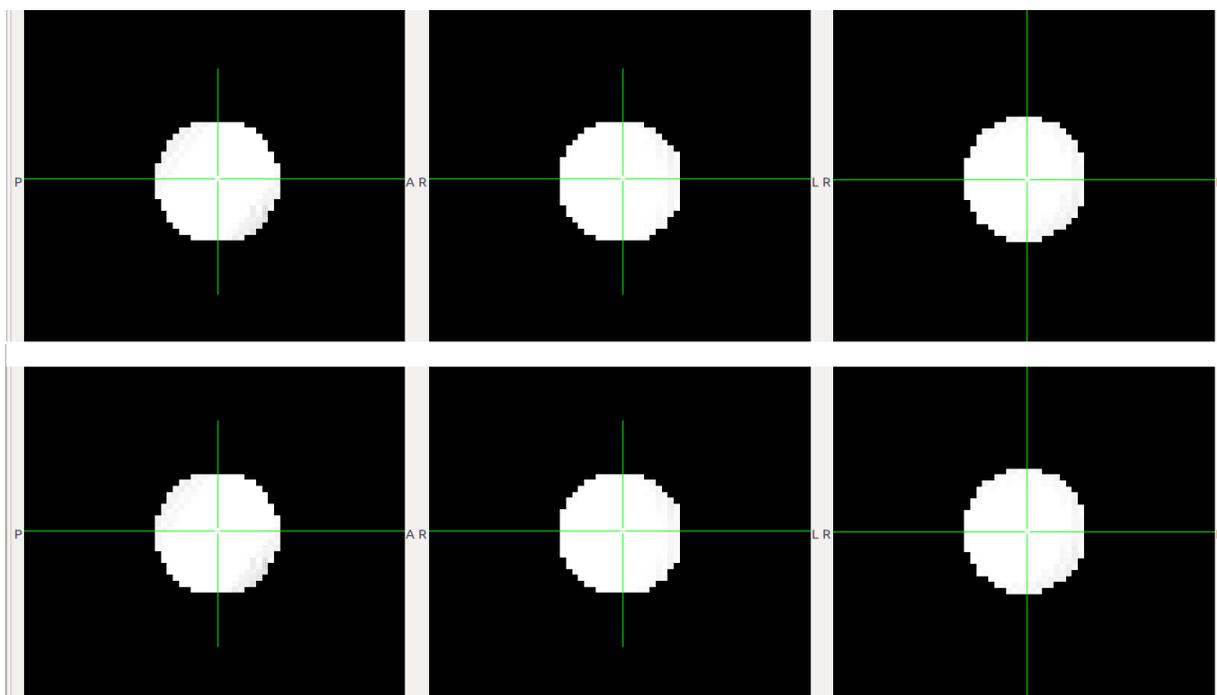


Figure 5.21: Gradient magnetic field maps of MRI phantom (masked magnitude images): (*top row*) without device, and (*bottom row*) with 2nd-generation keyboard at arm's length from isocentre. No differences are visible.

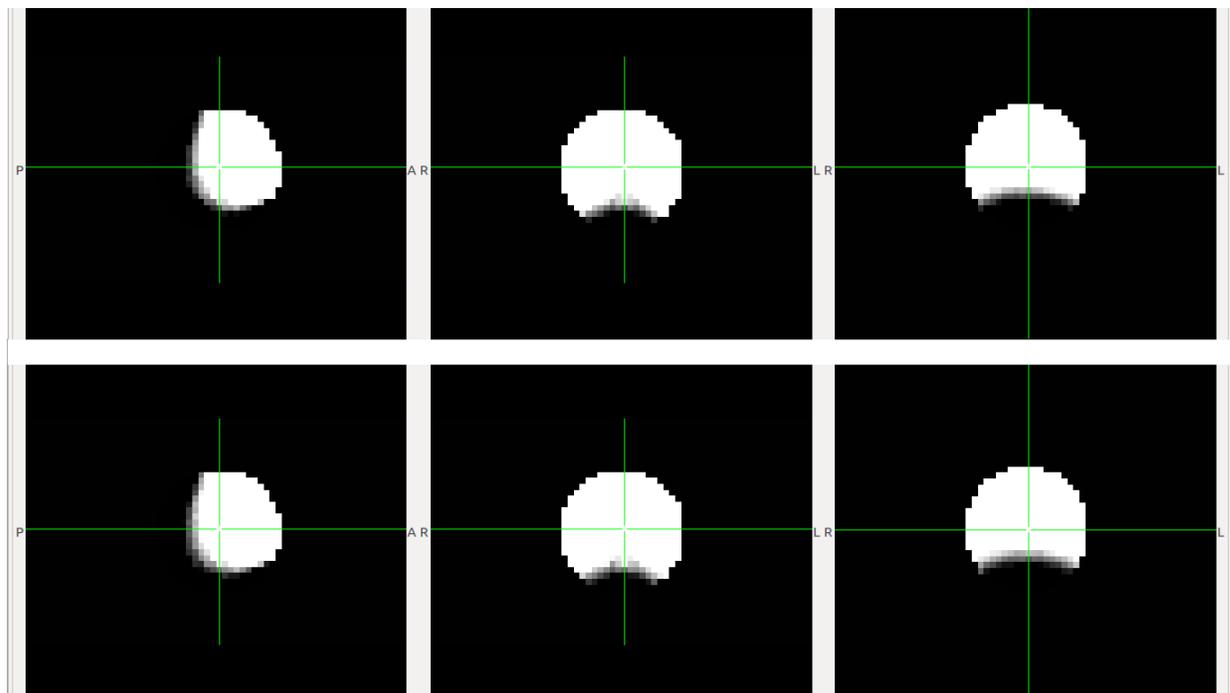


Figure 5.22: Gradient magnetic field maps of MRI phantom (masked phase images): (*top row*) without device, and (*bottom row*) with 2nd-generation keyboard at arm's length from isocentre. No differences are visible.

Magnetic field maps were also taken with a subject, with and without the piano. As seen in Fig. 5.23, no image artifacts are visible with the piano in the scanner bore. Fig. 5.24 shows the phase images with and without the keyboard, as well as a difference image with highly increased viewing contrast. Again, no meaningful differences in field homogeneity were detected.

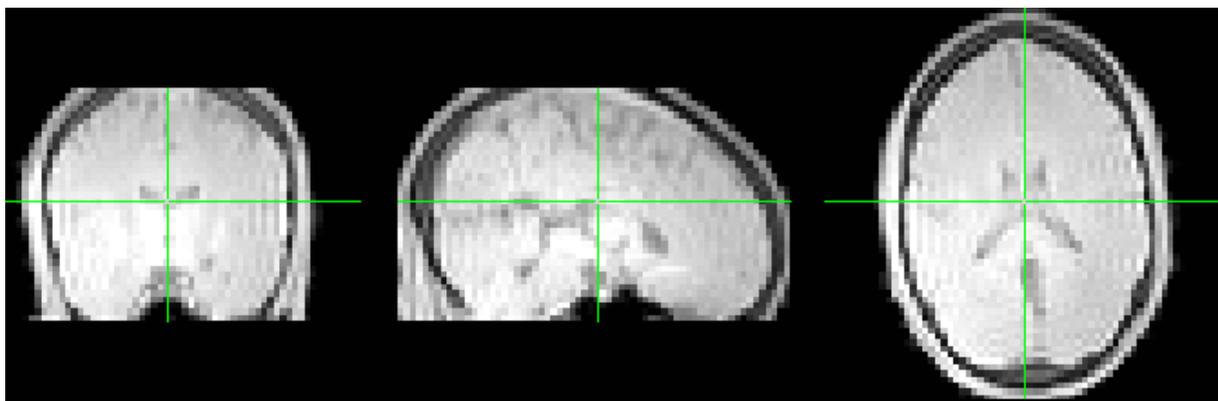


Figure 5.23: Gradient field map of subject with 2nd-generation keyboard at arm's length from isocentre (magnitude images).

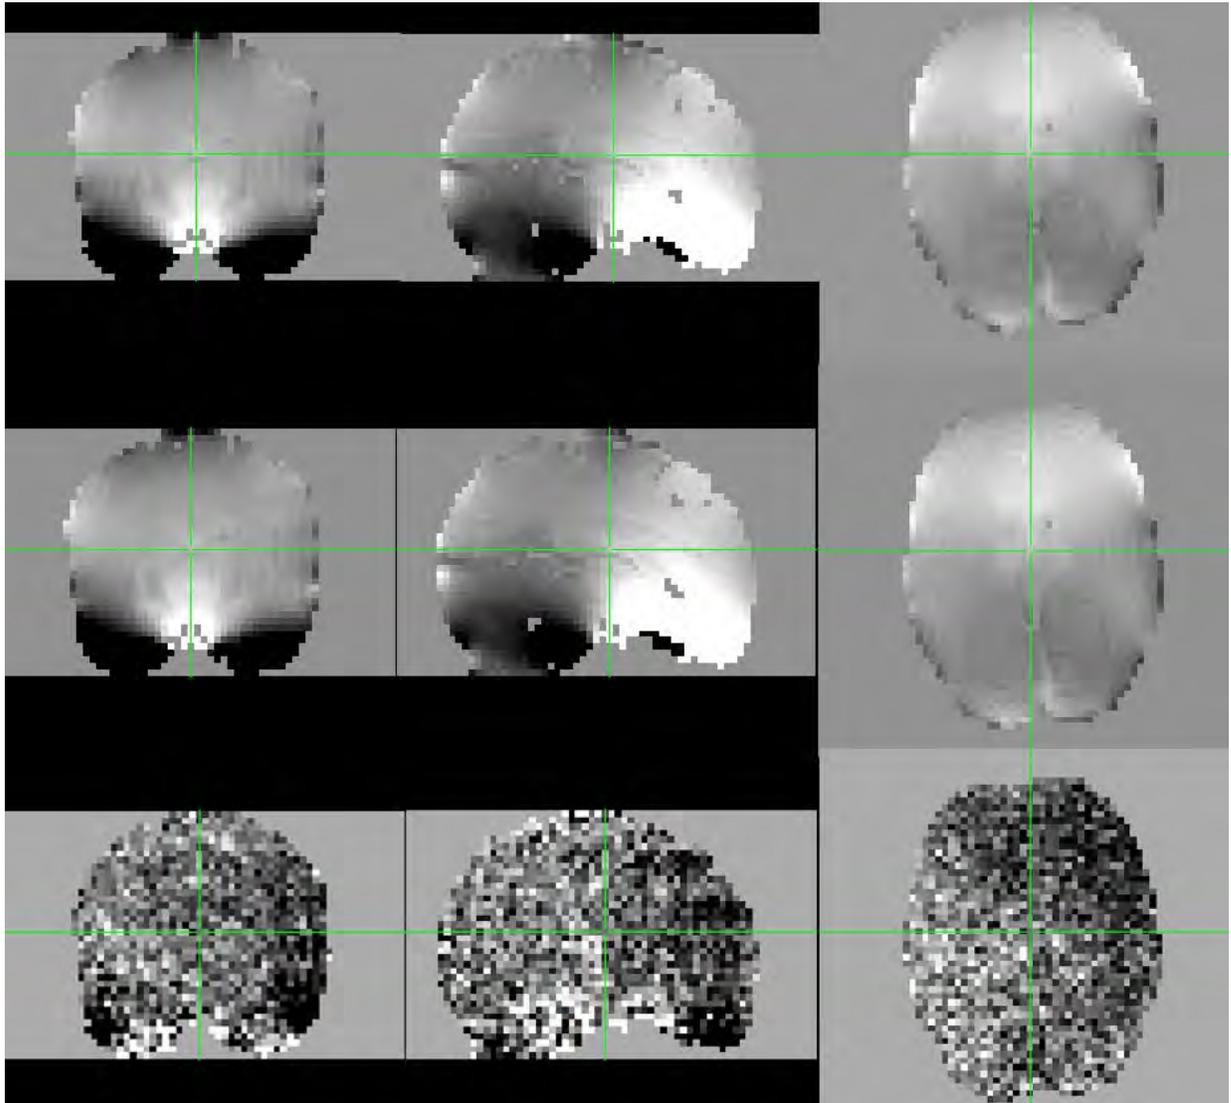


Figure 5.24: Masked phase images from field mapping scan, from top to bottom: without keyboard, with keyboard, and a difference image (with a re-normalized range to show detail).

5.6 Discussion and Future Improvements

While the results of the sensor testing have shown that the continuous gradient photomask provided the best performance from an SNR perspective, further refinements of the Moiré pattern technique could be improved to obtain higher SNR and spatial discrimination. For most musical tasks, however, the SNR and resolution of the continuous gradient photomask are sufficient while its simplicity and robustness are valued even more highly for continued use in the field. Precision-machined/patterned photomasks might also contribute positively to sensor performance, though their increased cost might not exceed the benefit.

Augmentations that would improve the optical sub-system, regardless of the photomask patterning, include lensing and higher output illumination sources. Inexpensive laser diodes housed in 3D-printed ST terminals is one implementation that we have identified and tested with good results. Additionally, plastic compound lenses, similar to those found in CD or DVD players, would improve the SNR as long as mechanical stability was ensured. This would also mitigate any key-to-key cross-talk caused by diffuse illumination or reflections. Lenslet or microlens arrays offer the advantage of lenses without the need to co-align the central axes of the fibres with the respective lenses. Additionally, Fresnel lenses offer good light focusing properties in a flat form factor in applications where imaging is not necessary.

As experiments increasingly make use of auditory feedback, a mapping of key-press velocity to loudness will be the logical next step. The sampling rate is perhaps the most important factor to allow the calculation of velocity and acceleration of key press trajectories. Sensor response can be calibrated to provide a linear displacement output, though without a sufficient number of samples, the derivatives will be obscured by noise. Commercial piano keyboards measure the time difference between the onset of two closely-spaced, though slightly offset, mechanical switches closing.⁷ Modelling the

⁷In Fig. 5.5, each pair of circles on the beige-coloured, flexible rubber strip are pushed into contact with electrodes on the circuit board when the respective key is depressed. Conductive contacts protrude from the underside of the circles, one protrusion being slightly longer than the other, causing the longer contact to close the switch slightly before the shorter one. The time delay between the completion of the electrical circuit is inversely proportional to the loudness of the auditory feedback. This mapping is not consistent across brands, though some models allow the user to tune this curve, with a parameter described as *touch sensitivity*.

trajectory as a smooth function, and taking the derivatives of said function will provide the best results, though at a computational cost that may be prohibitively expensive on an embedded system.

It is interesting to note, that during testing of the triangular differential photomask, where alignment of the photomasks with respect to the optical fibres was crucial, a slight side-to-side movement of the keys could be sensed. This sensing technique has further applications in extending performance of the usual piano keyboard, for example, to allow individual key vibrato or pitch-bend. While the use of side-to-side movement of piano keys is not novel, we were able to detect these slight gestures without altering the mechanical action of the keys.

5.7 Conclusion

The study of the neural correlates of piano performance using fMRI requires electronic tools to measure and analyze the fine motor movements of subjects within the MRI environment. With the goal of providing an ecologically-sound experimental setup for neuroimaging piano performance studies, an improved fMRI-compatible electronic piano keyboard that emulated the feel of an acoustic piano was developed. The optically-sensed keyboard allowed the collection of key presses and timing, as well as key position profiles, in order to quantify music performance and motor learning. The design methodology, implementation, testing, and experimental use of multiple iterations of an fMRI-compatible piano keyboard were presented, along with limitations and further improvements.

The current design includes a modified 24-key commercially-available electronic piano keyboard, devoid of electronics and ferromagnetic parts. Key bed mounts were fabricated to support and align fibre optic cables used to sense the position of each key individually. The keys are augmented with mirrors and photomasks, and weighted with brass beams. Various photomask designs were discussed, with the continuous gradient photomask being the best suited to the application. Modular optoelectronic control and acquisition boards were designed to calibrate, interrogate, and digitize the key position signals. The key sensors were tested and found to provide sufficient kinematic information. MRI-compatibility testing found that the piano keyboards did not create image artifacts and the acquisition electronics were not susceptible to EMI

emanating from the MRI scanner.

An fMRI study was conducted using the first generation keyboard, and was presented in [25]. A behavioural pilot study was performed with a transitional version of the fMRI-compatible keyboard. The MFST task included the synchronization of key presses with visually-presented sequential stimuli. Behavioural measures included key accuracy and timing. Key transition velocity data was also collected. An fMRI study using the MFST task was conducted, using the second generation piano and *fungible-1* acquisition system. The second generation piano was also employed with the *fungible-2* acquisition system in an fMRI study of pianists where auditory feedback was presented. Development of yet another study using this device is under way.

Chapter 6

Optoacoustic Cello

6.1 Why a Cello?

Few instruments rival the mellifluous tones of the cello. In the right hands, a cello can put forth melodies that are deep and haunting or light and soaring. Many scholars have sought to understand the complexities involved in playing the cello: from bow-string interaction [96, 235, 236, 289, 290], to the ability to discern and correct for micro-tonal pitch errors [36, 38], and on to the critical synchronization of left-hand and right-hand gestures [6].¹

Like other members of the violin family, the cello is played with two hands, each performing a vastly different function and gesture. Additionally, the cello allows for continuous control of pitch, amplitude, and timbre, throughout the duration of the sounded note. These two properties of cello performance gestures—disparate bimanual trajectories and continuously-variable on-line control—differentiate it and other non-fretted bowed string instruments from most other instrument groups. For example, keyboard instruments, such as the piano, use both hands to perform the same fixed-pitch selection gestures with no on-line control, other than terminating the sounded note.² From an experimental neuroscience standpoint, perturbations in auditory feedback requiring the cellist to perform some on-line motor correction are of great interest,

¹Henceforth, left-hand fingering gestures will simply be referred to as fingering and right-hand gestures as bowing, as is the case for most cellists.

²The clavichord is distinct within the clavier-based instruments with its nominal pitch selection, continuously-variable on-line pitch adjustment, and an ability to modify timbre, especially the attack.

and something that cannot be studied on fixed-pitch instruments. For the cellist, for example, this means that if the pitch that is heard is different from the one that he or she is intending to play, a shift in finger position is needed to play in tune. Cellists are not limited to correcting for pitch shifts, of course; perceived loudness, timbre, and a mixture of these sound parameters are what make up cello performance, with expert players having the most refined ability to perceive and react, and moreover to anticipate.

Another way of probing auditory-motor integration is to manipulate the perceived physical gestures, as opposed to the auditory trajectories, through the use of altered haptic feedback using mechanical actuators [100], peripheral nerve stimulation [220], or direct cortical stimulation [107, 151]. The latter methods are much more invasive, especially in the case of surgically implanted electrodes, thus modifying the design of musical instruments, instead of performers, is perhaps the best place to start.

In this chapter, we present the design of an MRI-compatible optoacoustic cello. This composite fiberglass instrument, for use in neuroimaging studies of cellists, optoelectronically captures the vibration of the gut cello strings as well as the performer's fingering and bowing. Through the design and application of novel fibre optic sensors, we were able to measure left-hand finger position, create string pickups, and quantify bow force, position, and velocity. Thus providing behavioural measures to the neuroscientist and, where intended, auditory feedback to the subject.

6.2 Design Goals

As it is our goal to enable the study of music performance through functional neuroimaging, specifically fMRI, we require musical instruments that can be brought into the scanner and played, to some extent, as one would outside the scanner. While the Baroque violin seems like a reasonable choice due to its small size and all-wood construction, standard violin technique has dispatched with a posture and holding of the violin toward the front with the tailpiece abutting the collarbone.³ This posture, as Leopold Mozart notes, "...looks well but is insecure" [22] and is an impediment to moving with agility over the fingerboard, and has been superseded with a firmer grip

³There is some debate whether this posture, seen in Baroque art, may have been more a pose struck for the *camera obscura* than actually used by violinists [183].

over the left shoulder. As such, there is not room for this posture within the MRI scanner, especially with the head coils used with fMRI. A standard cello, with its large arched resonating body, would not fit within the scanner either. If, however, the instrument were stripped of its resonating body down to the performance essentials, what would remain? We would have strings with a vibrating (or scale) length of 695 *mm* [255] tensioned between a nut and bridge. Additionally, a fingerboard supported by a neck of the proper shape and size are essential to cello performance, as is a bow. All other elements provide mechanical support or acoustic reinforcement.⁴ With these essential components in mind, we can focus on tailoring the instrument for the confines of the MRI scanner.

Our primary objective is to design a playable cello, ensuring that all components with which the cellist comes into contact, remain true to form. While cello dimensions are fairly standard [254], there is still some variability with regard to certain dimensions and curvatures, due to the craftsmanship and decisions of individual luthiers. Thus a design that aims at the middle, and not the extremes, should allow for a comfortable playing experience for all but the most discerning cellists. With the knowledge that metal strings are not MRI-compatible and that Baroque performance practice is still common practice for early music specialists, gut strings are still available and an ecologically-valid choice for our cello. Similarly, Baroque instrument-making techniques can be applied to the bow to ensure MRI-compatibility: replacing the metal worm-drive, used to tighten the bow hair, with a fixed or discrete tension arrangement.

A secondary goal is to directly acquire cello performance gestures. Depending on the behavioural task, the experimenter may require a recording of the sound produced by the cello or co-ordination of bowing and fingering gestures. Therefore we strive to capture several features of cello performance: the acoustic signal, finger position, bow velocity, and bow pressure. Through the acquisition of the strings' vibration, the sound of the cello could be amplified (and even processed) electronically before being fed back to the performer through headphones as auditory feedback. The acoustic signal, and especially its frequency content, can also be used as a behavioural measure, along with finger position sensors embedded in the fingerboard, to quantify left-hand gestures and timing, as a representation of the performer's pitch tuning acuity. Bow force in

⁴The upper bout may be used as a reference point by cellists to augment their proprioceptive memory when shifting to fourth position, so we consider this too for musical excerpts in fourth position.

combination with bow speed and its position are important performance parameters, which contribute to the sound intensity and timbre. Bow force can be measured in the bridge, as a low-frequency component of string displacement, or in the bow, as a compressive force modifying the space between the bow hair and the stick. In the fringe field of the scanner, the magnetic field varies with position, which can be used to determine the location and orientation of the bow.

6.3 Implementation

6.3.1 Parts of a Cello

An acoustic cello is made up of several parts, as shown in Fig. 6.1.⁵ The neck, fingerboard, and strings are parts with which the cellist comes into contact. The bridge, nut, tailpiece, peg box, and tuning pegs are used to tension the string, keep them at the right height, and properly spaced. The top and bottom plate, the ribs, bass bar, sound post, and end pin are mechanical supports to keep the instrument together under the string tension. Other than the end pin, these parts also greatly affect the tone of an acoustic cello.

6.3.2 Modification or Fabrication?

Two approaches to designing an MRI-compatible cello that seem inviting are piecing together acoustic cello blanks or modifying a commercial electric cello. Both methods have the advantage of using proven forms, thus cutting down on design time and costs, but adding time to customize and modify the instrument for the scanner. With the former, acoustic cello parts can be purchased and refined. The *blanks*—for the neck, fingerboard, bridge, and nut—that come from the luthier supply store, are roughly shaped and need to be refined as a luthier would. A mechanically-sound method to support the strings, and the underlying bridge, fingerboard, and nut, by means of a reinforcing beam would be required.

The modification of a commercial electric cello would require removing all ferromagnetic parts and electronics. While replacing steel screws with brass or nylon is fea-

⁵Ninat Friedland. *Cello Study*. Ink and graphite.

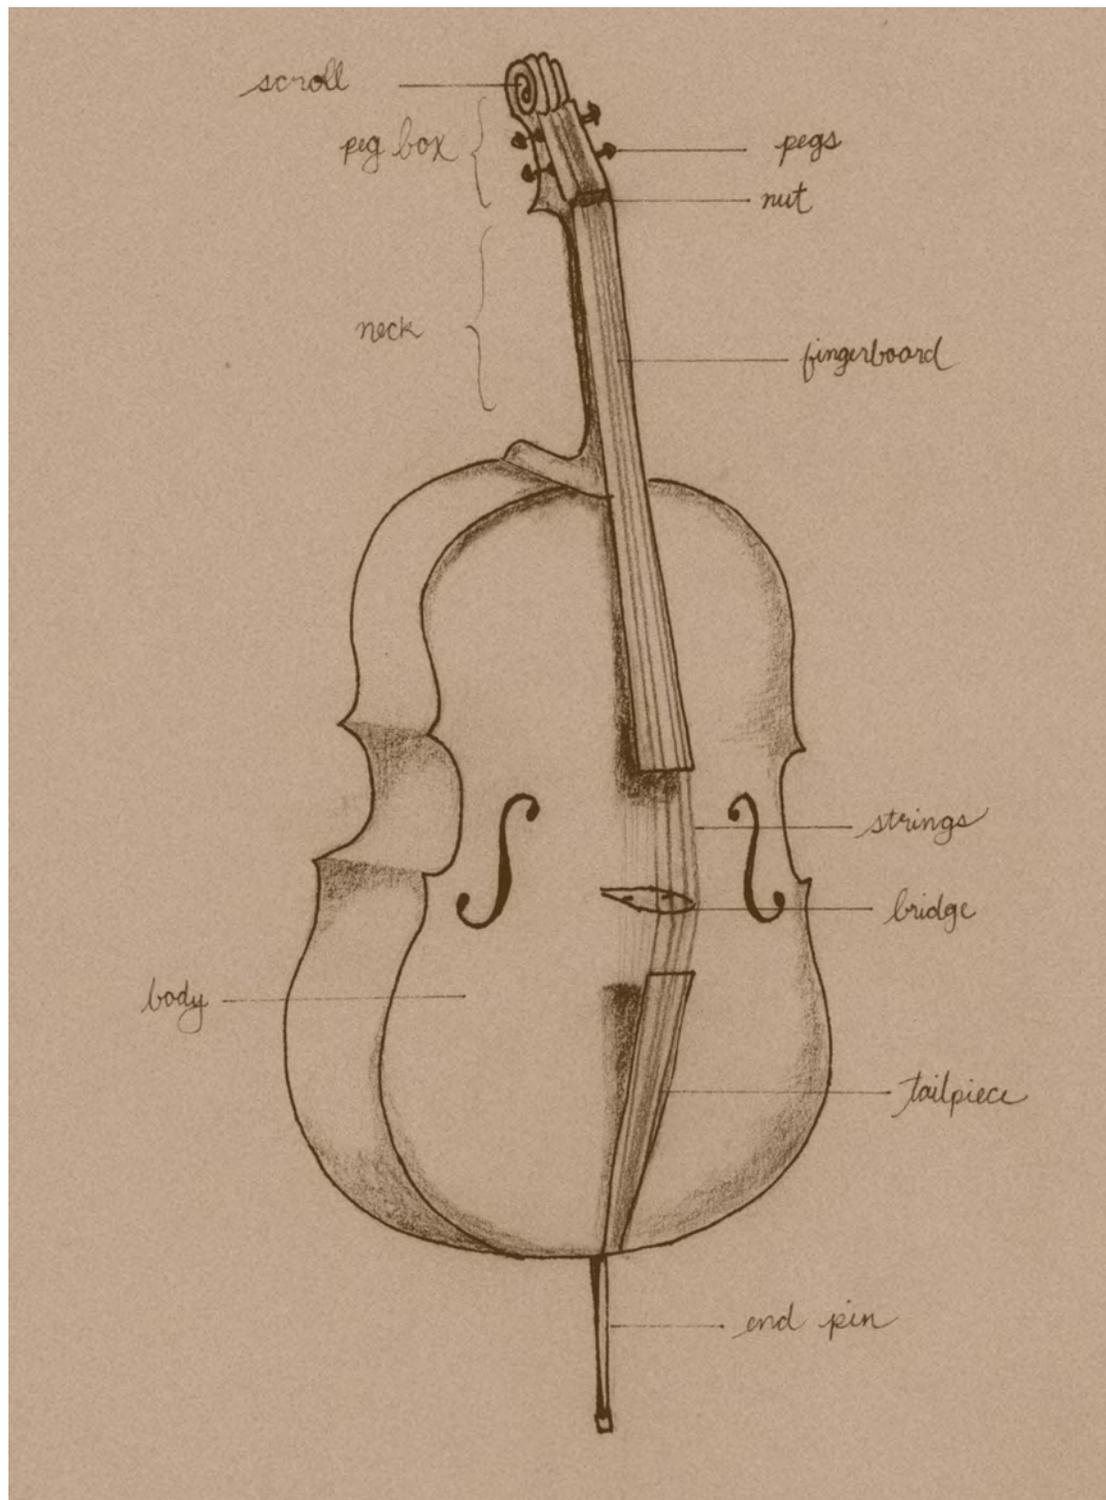


Figure 6.1: Annotated diagram of cello body.

sible, certain electric cellos contain other metallic hardware and housings (e.g., tuning pegs and fine tuners, bridge and pickup, enclosure for electronics and battery compartment). It is important to note that gut strings are of larger diameter than metal strings, making them incompatible with the bridge, nut, and tuners of pre-fabricated electric cellos. The electronic pickups that are usually built into the bridge would likely not function properly with gut strings or would not be MRI-compatible, and quite likely, both of these issues would arise.

Ultimately, the incorporation of sensors to allow the acquisition of performance gestures would require the modification of at least the fingerboard and bridge. There are limitations to modifying these and other parts such that they constrain the design process and make it less flexible. Specifically, the material of pre-fabricated parts is fixed and has already been determined, and it is usually easier to take away material than to add to it.

What appears to be a cost- or time-saving decision at the start may encumber the design so much that a work-around needs to be devised and implemented at every turn. It was thus determined that where feasible, the cello parts would be fabricated based on new or pre-existing designs using less-restrictive materials than traditionally used and more repeatable manufacturing processes. As well, traditional techniques should be used where their application does not compromise the goal of augmenting the instrument with sensors. There is also a trade-off of form versus function and impact on usability, as applying traditional techniques and materials will create an instrument that looks and feels more similar to what the musician expects, putting them at ease and enabling them to perform comfortably. Overall, a mixture of traditional techniques and materials with non-standard ones will allow the creation of an instrument on which a cellist will emulate cello playing.

6.3.3 Cello Body

The cello body, which usually performs the job of a resonating cavity, was designed as a single piece consisting of the peg box, neck, and supporting structure, as seen in Fig. 6.2. The cello body had to be strong enough to support the string tension and thin-walled to allow fibres to pass under the fingerboard while maintaining the outer shape of the neck. Without using metal, a composite material was selected because of its

high strength to weight ratio and previous work being done with resin-fibre composite musical instruments [202]. Carbon fibre is becoming more common in musical instrument design as it can be made to mimic the anisotropic properties of wood [176, 193]. Carbon fibre possesses a lower magnetic susceptibility than metals [82], such as those used in MRI-compatible electrodes, and thus generates smaller artifacts due to distortions of the static and gradient magnetic fields [216]. The conductivity of carbon-fibre materials combined with the RF emissions of the MRI scanner, however, may cause RF heating [160, 189], or artifacts due to RF shielding or eddy currents [82, 163]. While the image artifacts of carbon-fibre biopsy needles are relatively small, artifacts vary greatly with changes in placement and orientation of these foreign objects within the scanner. A carbon-fibre cello would have a relatively large volume, increasing the possibility of image artifacts. Additionally, the electrical properties of carbon-fibre composites are difficult to predict and perhaps impossible without well-characterized materials, as they depend on the geometry, interaction of component materials, and variability due to hand layup [81, 154]. As the MRI-compatibility of a carbon-fibre composite cello could not be ensured without testing the actual instrument (i.e., after design and fabrication), it was decided that the cello body would be manufactured as a non-conductive fiberglass composite instrument.

Fiberglass is often laid into a mould, sorbed with resin, and heat-cured. Due to the liquid epoxy resin, this *wet layup* technique is often messy, and releases potentially-hazardous fumes, thus a *dry layup* technique was employed. For this, thin sheets of fiberglass pre-impregnated with resin (prepreg) are cut to shape and laid into a mould, then heat-cured. During the curing process, the heat first causes the resin to re-flow in order to take the shape of the mould, and then as the cure cycle progresses and the temperature increases, the resin-infused fiberglass hardens. The cure schedule was 12 hours at 80 °C with a 2.5 °C/*min* ramp, up and down.

The first part of the body design was to conceive the general shape and layout of the cello. Within the MRI scanner, there is very limited room around the head, as the MRI head coil surrounds the head, thus the cello cannot be laid against the head or the neck. The peg box and tuning pegs, which are usually placed near the head while playing, would make playing difficult in the scanner, and so they were moved to the bridge-side, where the tail-piece would normally be. The back and sides (or ribs) of the cello were diminished in width and depth, while retaining the heel of the neck so that shifting to

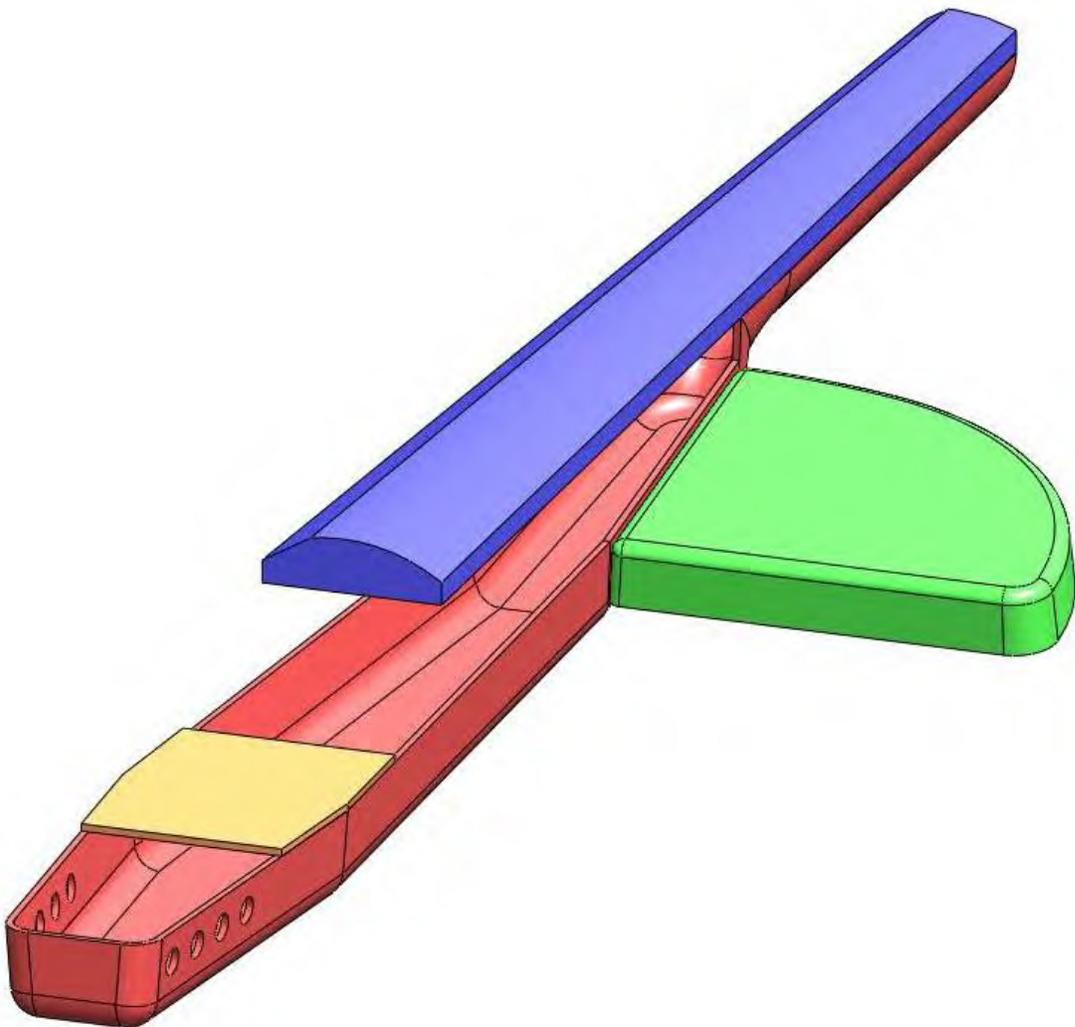


Figure 6.2: Early assembly drawing of MRI-compatible cello interface.

higher positions felt natural. The length of the underside of the neck, measured from the nut to the heel, was maintained at 240 mm [255]. The neck, back, and peg box were merged together to create the cello body, thus reducing to one the number of moulds that were required for the cello body, and simplifying post-processing and assembly. The body was created as a hollow *U*-shaped channel, which allowed optical fibres to pass within the neck, under the fingerboard. The hollow shell design also simplified the fabrication process, enabling the use of a single-part mould. The rounded *U*-channel allowed the fiberglass to better conform to the mould and eased de-moulding, after curing. The cello body design was prepared as a 3D model using computer-aided design (CAD) software. From the cello model shown in Fig. 6.3, a mould model was created from its negative image, shown in Fig. 6.4. The mould was machined out of a block of aluminium, with a computer-controlled mill, as shown in Fig. 6.5. The mould was then hand-sanded and polished to smooth away the tool paths and ensured the end-product would be just as smooth without the need for post-processing of the fiberglass. After applying a sealant and release agent to the mould, it was ready for layup. Two iterations of the fiberglass body were layed up.

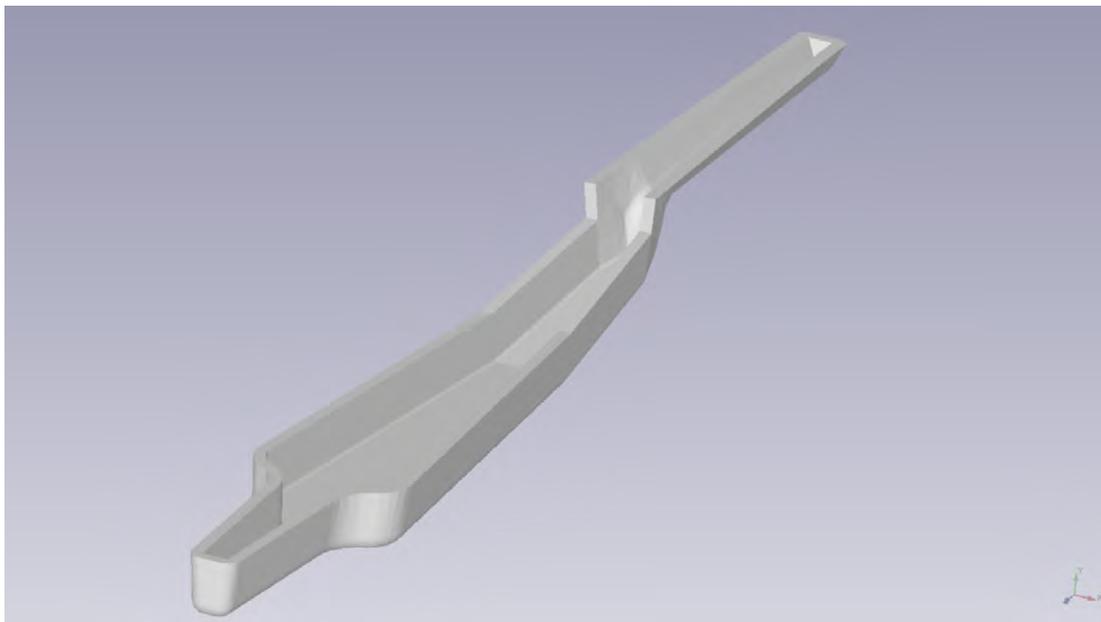


Figure 6.3: CAD model of MRI-compatible cello body.

Layup commenced by printing out paper stencils from the orthogonal-views of the

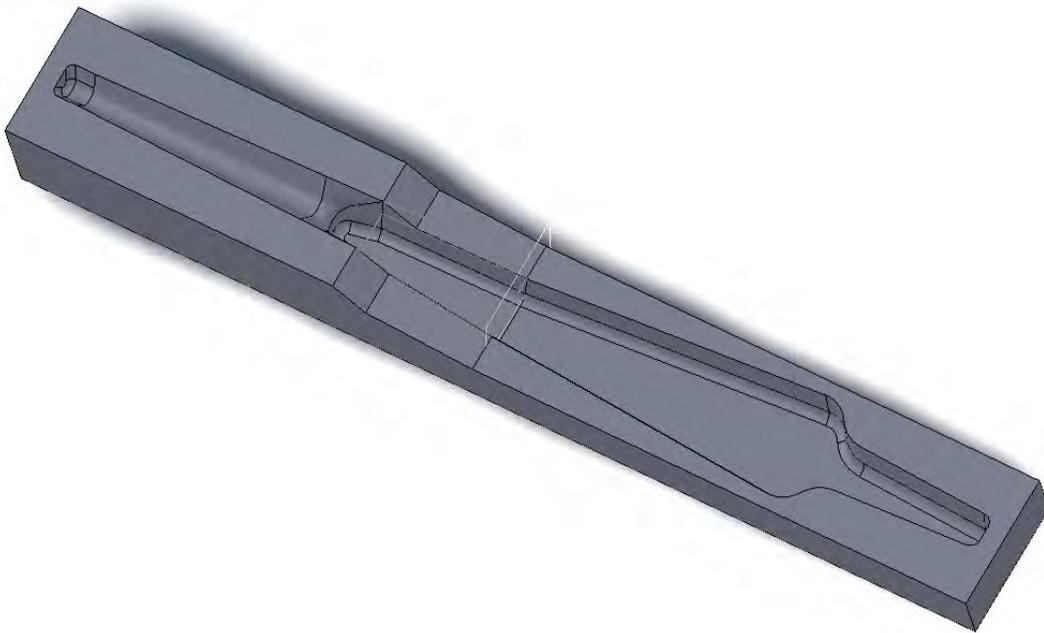


Figure 6.4: CAD model of MRI-compatible cello body mould.

CAD model, which were cut to shape (though slightly over-sized to account for the curvature) as shown in Fig. 6.6, temporarily laid in the mould, and then trimmed to the appropriate dimensions. This process aimed to create a set of planar templates that matched the curved internal surface of the mould, which were then used to cut the prepreg fiberglass to shape. The number of templates was minimized, while the size and fit of each template was maximized. Special attention was paid to areas of higher stress⁶ (e.g., the heel), as seen in Fig. 6.7, in an attempt to reduce the number of seams and reinforce these areas. While computerized modelling software exists to generate and project templates onto the composite layup workspace [146, 245], the time taken to perform this step iteratively and by hand, outweighed the cost and modelling time for this limited-run prototype. Once the paper stencils were suitably sized, the fiberglass sheets were marked and cut into shapes to match. The body was made of

⁶The simulation shown in Fig. 6.7 was performed using finite element analysis in SolidWorks by fixing the peg box and applying a force at the nut-end of the neck, simulating the forces exerted on the body due to string tension. The simulation was not meant to accurately model the material properties, the actual finished geometry (i.e., variable wall thickness), nor the forces exerted on the cello body. It is proffered only to show areas of relative stress and to confirm the location of expected weak spots.

three distinctively-shaped templates: the sides, the back, and the neck. Each shape, consisting of four layers of fiberglass, was laid into the mould, slightly overlapping. Additional pieces of fiberglass were added to patch any voids and to strengthen seams. Shapes cut from fiberglass-backed balsa wood, ranging in thickness between 3 and 10 *mm*, were placed atop the first set of fiberglass shapes, as shown in Fig. 6.8. This inner layer, known as the *core*, gives the composite material its strength in the form of high bending stiffness⁷ without a weight penalty that would arise from a denser solid material [11]. Several more layers of fiberglass, cut to shape—however slightly smaller by this point, as the internal dimensions were reduced with each additional layer—were laid on top of the core, seen in Fig. 6.9.

With the sandwich of fiberglass prepreg and wood laid in the aluminium mould, the assembly was ready for heat-curing in the oven, under vacuum, as shown in Fig. 6.10. The low pressure of the vacuum forced the composite material against the single-sided mould, ensuring no air gaps or *voids* were formed, which might cause de-lamination of the fiberglass layers. The assembly was first wrapped in a synthetic wool in order to protect the vacuum bag from being pierced by the sharp edges of the mould. A vacuum bag was created to encompass the assembly using a plastic film, a valve, and specialized vacuum tape. Once the bag was formed and sealed, the air was evacuated from the bag with a vacuum pump, and the bag was checked for leaks and further sealed as necessary. The pressure was then released, slightly, to repositioned the bag inside the mould, allowing it to better conform to the internal structure, further reducing the possibility of voids during curing. With the vacuum applied, the assembly was placed in an oven for 12 hours at 80 °C. After curing and cooling, the assembly was de-moulded and post-processed; excess fibres were cut away and rough edges were sanded down. Holes were drilled through the peg box for the tuning pegs and other holes were drilled in the nut-side end for the strings to pass through, as shown in Fig. 6.11.

In order to increase the stiffness of the cello body and improve the surface finish, a second iteration of the process employed 16 total fiberglass layers, instead of the initial eight. As well, a surfacing film (Surface Master 905C) was placed on both inner and outer layers during layup to provide a smoother and more regular finish. This film

⁷The bending stiffness is proportional to the square of the core thickness.



Figure 6.5: Aluminium cello body mould during initial stages of layup.



Figure 6.6: Paper template preparation for cello body.

flowed easily and was better able to take on the contours of the mould, eliminating voids in the surface and on a smaller scale, eliminated the surface roughness of the fiberglass weave. Additionally, adhesive layers (MeltBond 1113), were placed on both sides of the balsa wood core during layup to provide a stronger bond between the fiberglass and core, reducing internal voids and de-lamination.

6.3.4 Fingerboard

The role of the fingerboard is to provide a surface against which the fingers depress the strings in order to shorten the sounded string length, essentially fretting (or stopping) the strings. During cello performance, multiple fingers may depress the same string in order to smoothly play the repeated notes of a trill, in anticipation of a downward run, or to share the string tension and ease the use of weaker fingers, especially in thumb position. Thus the fingerboard, integral to pitch selection and vibrato, presents an interface that may be innervated with sensors to directly measure either gross or nuanced, dynamic or anticipatory, left-hand gestures not available through indirect, acoustic means.

Modern cello fingerboards are usually made of ebony, a hard wood that can stand up to the wear of the metal strings. Softer gut strings were paired with a maple fin-

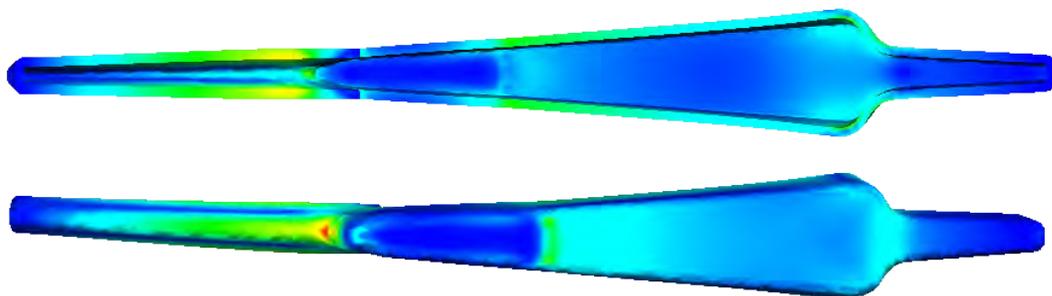


Figure 6.7: Cello body simulation of relative stress distribution with views of top and underside, showing higher-stress area at the heel of the neck.

gerboard, often veneered with ebony, on Baroque instruments, reducing weight and cost [144]. Baroque instruments also had shorter necks for a combination of practical reasons. A flatter bridge and curved bow allowed two or more strings to be bowed simultaneously, often heard in the music of Bach and Telemann [200, 206]. The flatter bridge would have also made it more difficult for strings to be sounded individually, especially at higher fingering positions. Lowering a middle string (i.e., the D or A string) on to the fingerboard, far toward the bridge end, would eliminate the clearance with respect to adjacent strings, required to properly and cleanly bow the string.

Augmenting the fingerboard to sense finger-string-fingerboard interaction can be accomplished in a variety of ways, even when constrained to only optical sensing. What exactly one seeks to measure is a fundamental question. If the contact position is of interest, then is a distributed and continuous measure needed, or would a discrete and binary measurement suffice? Is the contact area or force of interest, too? While we ideally wanted the most comprehensive measures possible, the feasibility of a distributed fibre optic strain sensor (using Brillouin scattering, for example, as discussed in Section 2.2.5) embedded within the fingerboard was questionable, and thus posited only as a long-term goal. A rigid fingerboard would undergo very little localized strain during cello performance, while back-scattering techniques relying on minute spectral shifts, slight amplitude modulations, and highly sensitive optoelectronic components require highly stable, laboratory-like conditions. A prototype instrument used for musical performance, with the goal of minimizing set-up time in the MRI scanner, does not provide these ideal conditions even if the fingerboard strain were measurable in

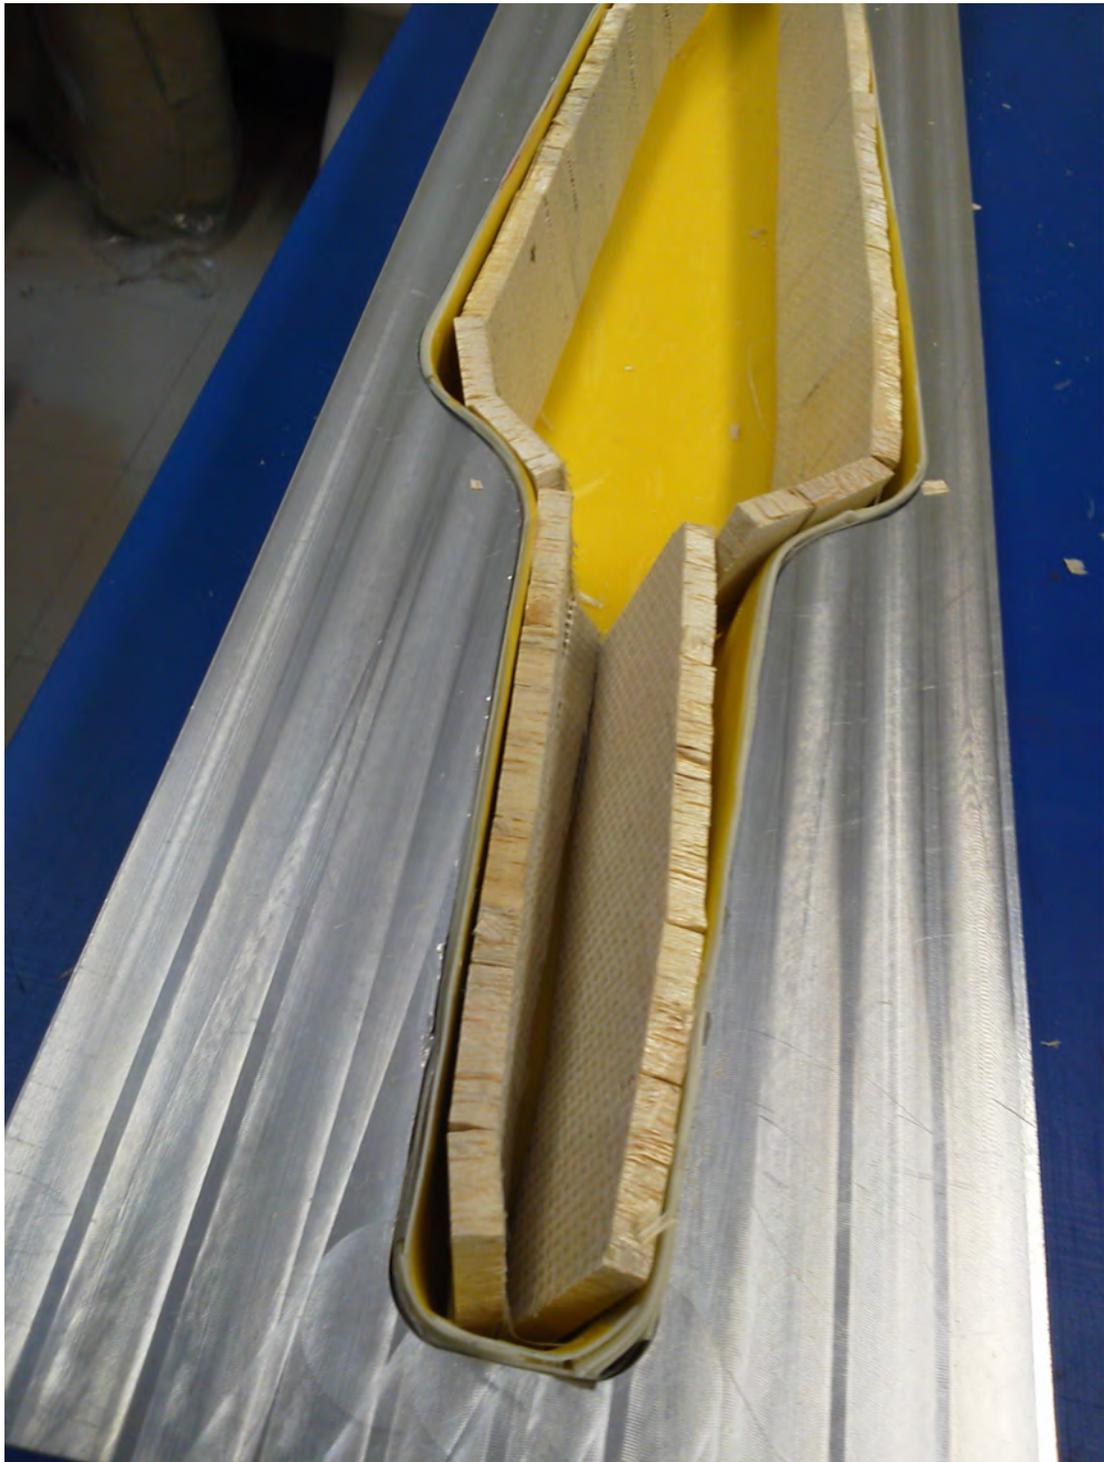


Figure 6.8: Composite cello layup with first layers of fiberglass prepreg and balsa wood core.



Figure 6.9: Composite cello layup with second set of fiberglass prepreg.



Figure 6.10: Composite cello in vacuum bag and conformal protective envelope, ready for heat-curing.

this way. Similarly, quasi-distributed strain sensing using in-fibre gratings requires a reasonable amount of strain to be used to sense fingerboard compression. While it was possible to modify the fingerboard to incorporate a conformable top layer to increase the strain on the sensing element, the low-pass filtering effect of the added layer would undoubtedly impair the playability of the instrument.

So while strain sensing within the fingerboard was not feasible due to the rigidity of the material and the relatively small forces applied by the fingers of the cellist, proximity sensing through the acquisition of light reflected off of the string and finger back into the fingerboard, however, is feasible. Several sensorized fingerboard design iterations are presented here.

Planar Optical Fibre Layup in Fiberglass

A first attempt was made to incorporate fibre optics within a composite fiberglass fingerboard, with the optical fibres laid parallel to the surface of the fingerboard, cut at an angle. Initial tests with a fiberglass rectangle (pictured in Fig. 6.12), using only two layers between the optical fibres and the sensing surface, enabled the acquisition of finger presses on the test fixture, as shown in Fig. 6.13, providing a signal-to-noise ratio (SNR) of 26.1 dB. The limitations of the planar embedding of optical fibres within the composite material were repeatability of placement and sensitivity, heat-induced damage due to the cure schedule of the fiberglass, sensor density and the absolute number of discrete sensorized areas, and allowing the fibres to exit the fingerboard in a manner in which the fingerboard could be joined to the body while minimizing de-lamination. In order to overcome the issue of repeatability, several methods were proposed: the use of a jig to cut the fibres and a machined foam core embedded within the composite, which would be used to maintain fibre alignment. As well, modifying the fiberglass cure schedule to a lower peak temperature, within the recommended limits of the optical fibre (i.e., less than 70 °C [65]), would ensure that the POF was not damaged during curing. While sensor density could be increased slightly with the use of smaller diameter POF, reduced sensitivity would follow. The planar POF layout was found to be infeasible, as the goal of a quarter-tone resolution over a span of an octave for each string was not reachable using this geometry. As well, at the desired density and sensitivity, there would be no space for all of the fibres to exit the end of

fingerboard.

Fiberglass Fingerboard

A fiberglass fingerboard was produced using perpendicular POF alignment. First a CAD model of a cello fingerboard was designed, then a two-part aluminium mould was fabricated from a negative image of the model. The main mould, shown in Fig. 6.14, included the contoured portion of the fingerboard, while the second part was a flat bottom plate.⁸ The mould also included two ejection pins to help de-moulding of the part. These holes were plugged with slightly-recessed bolts during layup. The mould was hand-sanded and polished, sealed and coated with a release agent. Three layers of prepreg fiberglass were laid in the main mould. A foam core, shaped to match the mould, was then inserted into it. Finally, three more layers of fiberglass were laid on top of the foam core and the mould was closed with the top plate and fastened with bolts. After the assembly was cured, the fiberglass fingerboard was de-moulded and cut to size, as shown in Fig. 6.15.

Structurally, the fiberglass fingerboard performed adequately on the cello. Augmenting the fiberglass fingerboard with optical fibres placed perpendicular to the length of the playing surface proved more challenging. Holes were marked, according to the placement of a diatonic scale, and drilled through the fingerboard. Triples of POF were inserted into the holes and the sensitivity of the sensors was tested, with one fibre serving as the emitter and the other two as receivers. As there was no space between the fibre end-face and the finger-string interface, light was occluded instead of being reflected down the receiving fibres. Therefore, channels were cut in the fiberglass between the fibre ends and the fibres were recessed, allowing light emanating from the emitter to bounce off the target (i.e., finger and string) onto the receiving fibre. While the finger-to-fingerboard interaction was measurable, the surface finish was not sufficiently flat, such that it impeded performance.

Transparent Urethane Fingerboard

The strings and fingers, while translucent, did not internally convey a sufficient amount of light from the emitter to the receiver of the fiberglass fingerboard when there was

⁸The fingerboard mould did not include the carved underside nor the subtle *scoop* [86].

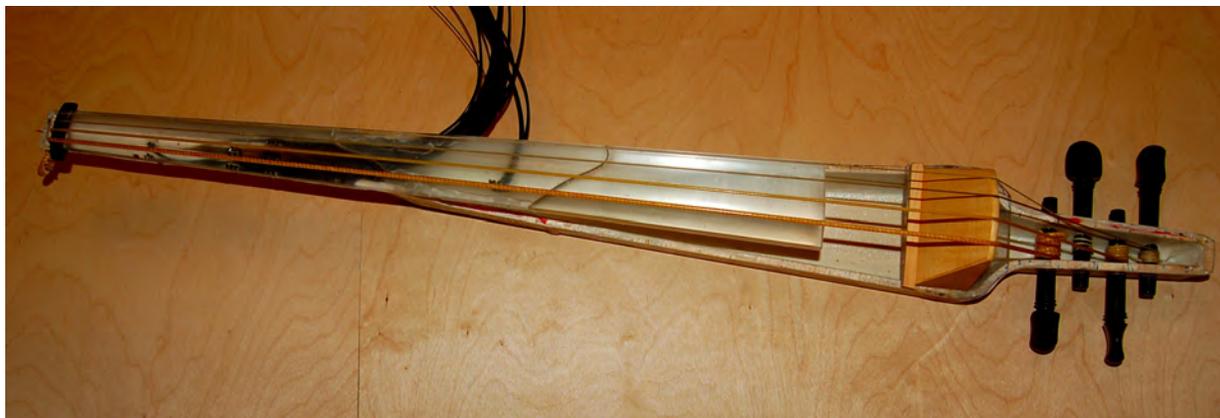


Figure 6.11: First prototype of composite fiberglass cello body, with gut strings and transparent fingerboard embedded with optical fibres.



Figure 6.12: Fingerboard test fixtures with embedded optical fibres: composite fiberglass touch sensor).

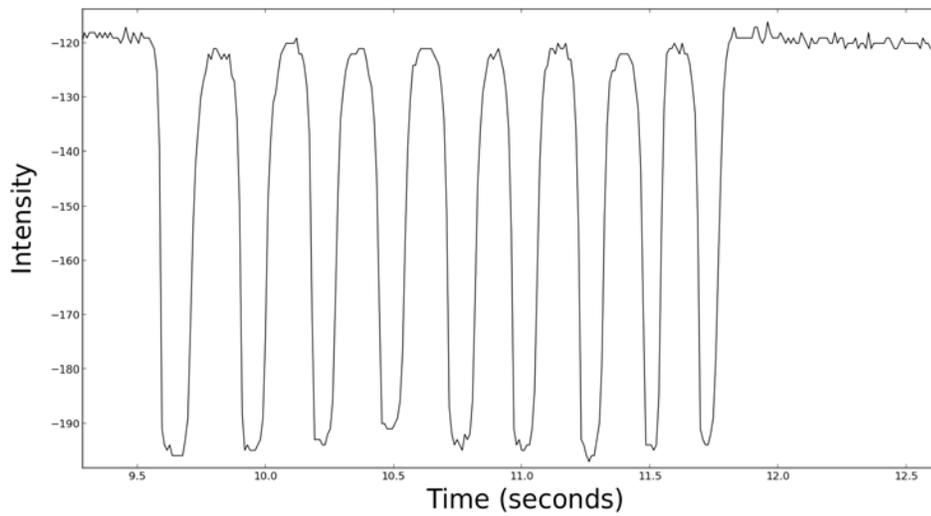


Figure 6.13: Repeated touches on fiberglass touch sensor (100 Hz sampling rate).

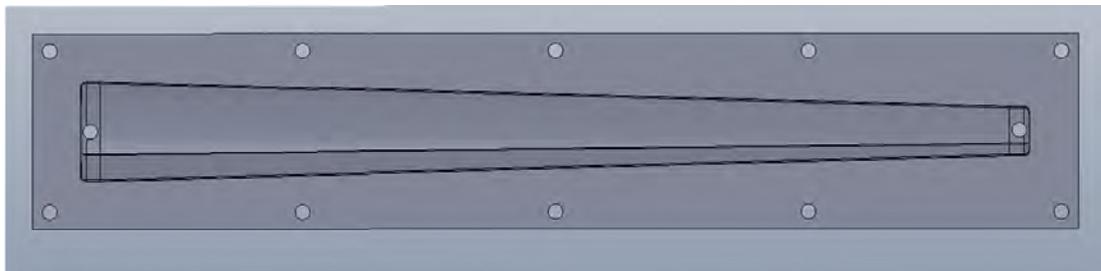


Figure 6.14: 3D CAD model of cello fingerboard mould.



Figure 6.15: Wooden (top), fiberglass (middle), and cast transparent urethane (bottom) cello fingerboards.

no gap between the interfaces. Recessing the fibres allowed a path for light to reflect off of the target surface, however, to ensure the fingerboard surface finish was not affected, the fingerboard needed to be designed with this in mind, preferably using a transparent material. A transparent, rigid urethane casting resin, Smooth-On Crystal Clear 200, was selected for its optical qualities, specifically its refractive index of 1.499 [249], closely matching the index of refraction of the POF, of 1.49 [66]. A transparent test piece was cast with 0.5 mm POF, as shown in Fig. 6.16 and was used to measure repeated finger touches with an SNR of 15.0 dB, as shown in Fig. 6.17.



Figure 6.16: Transparent urethane test piece cast with 0.5 mm POF.

The two-part resin was mixed and de-gassed in a vacuum chamber until the large bubbles were removed. It was then poured into the aluminium fingerboard mould, where it set and cured at room-temperature for 48 hours. After de-moulding, the part was trimmed (see Fig. 6.15) and holes were marked and drilled. Fibres were potted with Epo-Tek 301 [64], a transparent epoxy with an index of refraction of 1.539. Using this fingerboard, we were able to measure finger-string-fingerboard interaction, though there were some limitations. The process of drilling holes in the urethane caused the inner surface of the holes to take on a cloudy, less transparent finish, and thus reducing

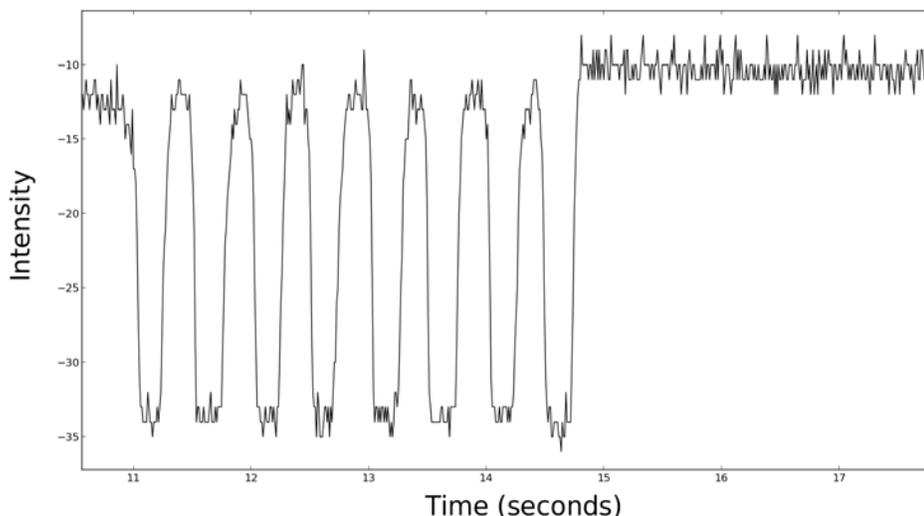


Figure 6.17: Repeated touches on cast transparent urethane proximity sensor, measured with 0.5 mm POF.

the sensitivity. As well, using a simple drill press, without a computer-aided mill, the hole angle was not normal to the surface of the fingerboard. Additionally, optical fibres tended to shift slightly, out of the holes, during curing as fixing them in place was challenging.

A second fingerboard was cast with the POF embedded in the fingerboard during casting. The POF were cut and polished beforehand. In order to mount the fibres in the right position and to a usable depth, another fingerboard was used as a template. It was marked and drilled and the POF threaded through, and secured above the mould. The resin was mixed and de-gassed, as before, and poured into the mould with the ends of the optical fibres submerged in the resin. The resin was allowed to cure for 48 hours. As shown in Fig. 6.18, this process created a transparent fingerboard without any clouding, provided a sensor with high sensitivity, and enabled the acquisition of finger-string-fingerboard interaction. As the optical fibres were not fixed within the mould, shifting during potting and internal stress during curing caused the fibres to move slightly. This affected both the depth to which the fibres were recessed, with respect to the curved surface of the fingerboard, as well as the angle and placement. Appendix A gives the equations governing the proper fibre placement within the fingerboard.



Figure 6.18: Transparent urethane fingerboard cast with POF mounted on cello body.

Both versions of the transparent urethane fingerboard allowed the measurement of finger contact position and timing, as well as vibrato, however, an issue was anticipated with extending the prototypes, which had only a few sensor triples. Due to the curved face of the fingerboard, the fibres were not aligned normal to the surface of the fingerboard. The depth to which the fibres were recessed could not be finely controlled. A deeper recess projects the light onto a larger surface area, allowing one sensor to cover a larger range, at the expense of decreasing the signal-to-noise ratio. Scaling up the number of sensitive regions from just a few to several dozen, or equivalently from nine fibres to over 100 fibres is no simple task. Curing that many fibres, either one at a time or all at once would require a properly machined jig, more expense, and additional time.

3D-Printed Fingerboard

While the cast resin material made for a highly transparent fingerboard, we questioned the feasibility of accurately placing optical fibres in a scalable and repeatable fashion. Rapid prototyping, specifically the use of 3D printing, allowed for relatively inexpensive manufacturing with fast turn-around time of iterative prototypes. Thus models could be designed, fabricated, tested, and re-designed fairly quickly. Starting with the fingerboard CAD model developed for the mould, a test piece of reduced dimensions extending only to second position was created (see Fig. 6.19). The underside was hol-

lowed out to allow for more room for fibres to bend with a larger bend radius, thus reducing bending losses. In the first design simple recesses, 3 mm deep and 1 mm in diameter, with a 3 mm spacing were included for inserting 1 mm diameter POF. This design was printed in a translucent plastic. With this truncated prototype we were able to see that the 3D-printed design was feasible but needed further revision. The iterations of these prototypes is shown in Fig. 6.20.

The small recess, alone, was not sufficient to support the optical fibre within the fingerboard, and so the recess was extended to form a collar around the fibre. Additionally, the recess geometry was changed to form two concentric recesses, so that 1 mm and 0.5 mm diameter fibres could be embedded within the fingerboard. As the 1 mm and 0.5 mm fibres had 2.2 mm and 1 mm outside jacket diameters, respectively, the insertion depth was controlled by stripping away the jacket by the prescribed length from the fibre. With holes distributed along and across the fingerboard, the hole geometry provided the most versatile prototyping fingerboard, so that different diameter fibres could be placed with a variable, yet precise depth, wherever they were required. The hole angles were designed so that the fibres within the holes were aligned normal to the fingerboard surface. A small test piece was printed in a transparent material, and after initial tests, a shortened fingerboard allowing for just over an octave on each string, was printed. Due to the 3D printing process, these layers, when printed off-axis tend to cloud slightly. As well, the process of removing support material during post-processing left a slightly cloudy residue. With the addition of an optical adhesive, the fibres acted reliably to reflect light off of the fingerboard-string-finger interface, back to the receiving fibres. The illuminated 3D-printed fingerboard mounted on the cello body is shown in Fig. 6.21.

6.3.5 Bridge

In order to reduce the size and complexity of the cello, the instrument was designed so that the strings passed over the bridge, directly to the tuning pegs within the peg box. The cello was abridged, without a tail piece, and thus the string angle passing over the bridge was significantly greater than that found on an acoustic instrument. A traditional cello bridge would have collapsed under the uneven tension of the strings. The bridge in an acoustic cello is designed to transfer energy from the vibrating strings

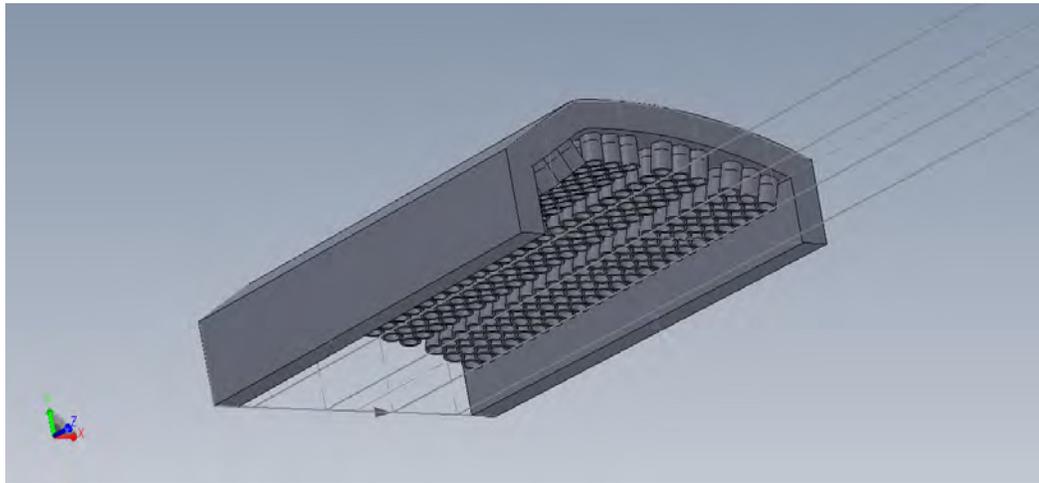


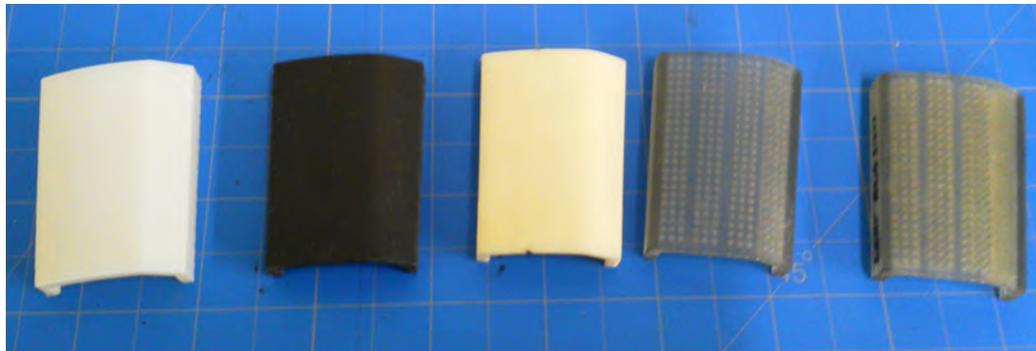
Figure 6.19: Truncated portion of CAD model of 3D-printable cello fingerboard.

to the body of the instrument, projecting a tonally balanced sound, loud enough to be heard on stage. The MRI-compatible cello, however, was not meant for stage performance, and thus a solid bridge without feet was designed to rest atop a flat acrylic bridge plate. It was expected that the acoustic sound would thus lack the bass response and “bridge-hill” peak at 2–3 *kHz* [126]. Furthermore, in experiments where feedback was presented through headphones and altered electronically, the actual acoustic emissions of the cello were masked with noise. Thus it was beneficial to minimize the acoustic resonance of the cello.

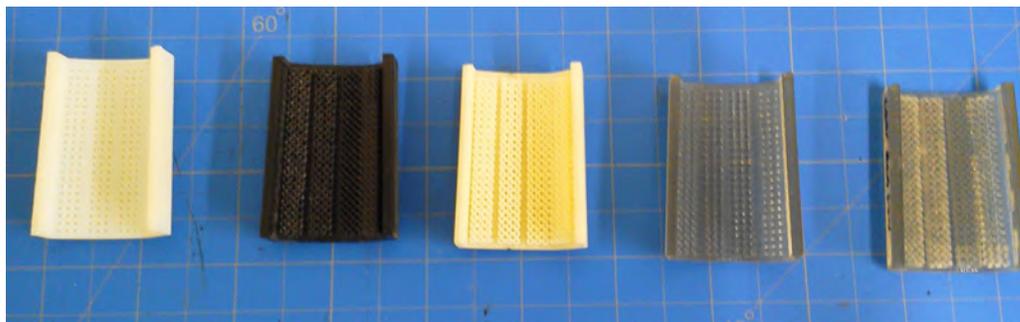
A bridge, made from a wedge of wood and shown in Fig. 6.22, was constructed and shaped by a luthier so that the top curve of the bridge conformed to the curve of the fingerboard. The front face of the wedge was also contoured so that the bridge thickness, where it met the strings, was uniform; thus maintaining the same, correct string length and intonation. The development of a bridge augmented with sensors to measure bow-string pressure and string vibration was carried out through a number of design iterations, presented here.

Contact Bridge Pickup

The bridge functioned as it was designed, though there was an opportunity to augment it with an optical pickup. Initial tests with a tightly-coiled 0.5 *mm* POF under compression suggested a bridge pickup could be made to sense string vibration. A coil



(a)



(b)

Figure 6.20: Truncated 3D-printed fingerboard prototypes: (a) top-side and (b) bottom-side views. From left to right: initial prototype with fibre recesses only, printed in translucent white plastic; prototype with fibre collars and concentric recesses, printed in black opaque plastic—the thickness between the inside of the hole and the top surface was sufficiently small to transmit some light; test piece made on lower-cost printer—dimension tolerances were not sufficient for off-axis holes, white plastic was opaque; clear plastic with fibre recesses only; clear plastic with fibre collars and concentric recesses—similar to full-size fingerboard.



Figure 6.21: Transparent 3D-printed fingerboard mounted on cello body with gut strings.

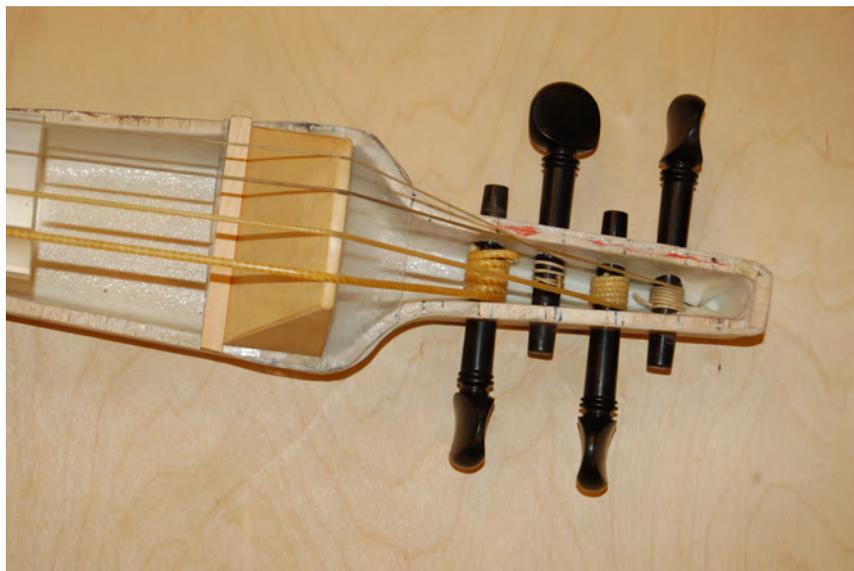


Figure 6.22: Wooden wedge-style cello bridge on acrylic bridge plate, and peg box.

of fibre placed between the bridge and the string was sufficient to sense the acoustic vibration, and due to the slight asymmetric placement, up-bows were distinguishable from down-bows. Additionally, the low-frequency component of the signal reflected a combination of bow pressure and string tension modulated by left-hand finger position along the string. With proper placement, two matched coils of POF could form a differential pair, thus increasing SNR. The main issues with this sensor included the lack of robustness in positioning, the difficulty in obtaining repeatable measures, and damage to the POF. As the sensor coil prototype was simply sandwiched between the string and the bridge, the string was not well seated, allowing the string to drift out of the notch in the bridge. As the gut strings needed to be re-tuned between and within some sessions, the fibres tended to be displaced slightly. The losses caused by bending was nonlinear for large bends [66], yet the small losses caused by the modulation of compression could be approximated as a linear function. Thus, the fibre needed to be biased into an optimally-sensitive regime by pre-stressing the fibre. Three parameters accounted for the pickup's sensitivity: the coil diameter, the number of turns, and the pressure on the fibre. A smaller coil diameter translated into more light lost per turn, due to bending losses. The number of turns, in addition to biasing the sensor to provide a certain static loss, controlled the area in the fibre under compression, thus increasing sensitivity. The static compression of the string pinching the optical fibre accounted for another large-signal loss. With use, the POF degraded as the sharp angles within the notches of the bridge and the large compressive force caused the fibres to deform catastrophically.

A CAD model of the MRI-compatible cello bridge was prepared in order to mount the fibres in a repeatable way, while minimizing the wear on the fibres. The model, shown in Fig. 6.23, was 3D-printed and mounted with fibres. The 3D-printed bridge was designed with a secondary notch for each string, containing a disk, used as a bobbin, around which the POF was coiled. It was found that micro-bending, rather than compression of the fibre, accounted for the greatest loss and sensitivity to vibration. A single loop of fibre, bent at almost 180° and wedged within the groove in the 3D-printed bridge, under the string, was used to measure string vibration. Though, again, this pickup method suffered from wear.

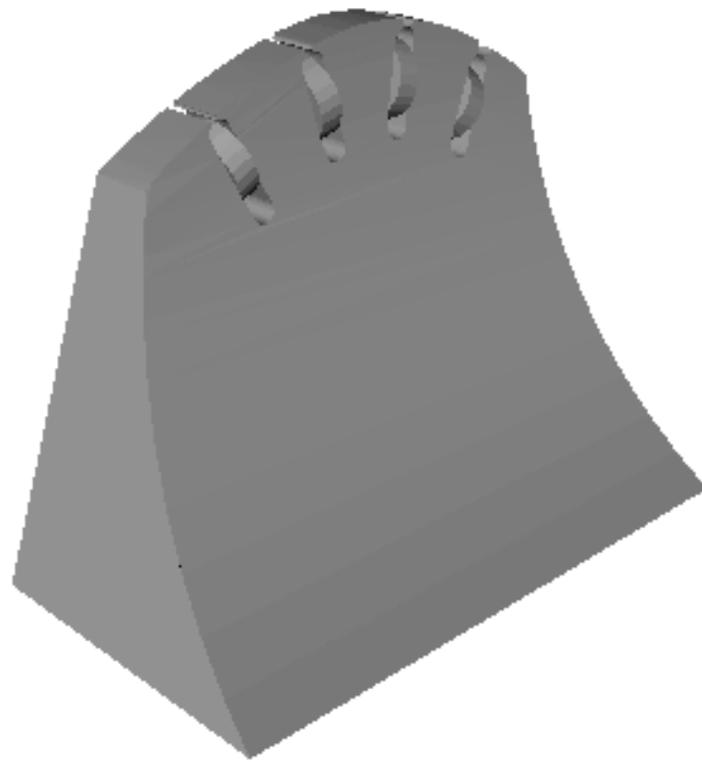


Figure 6.23: CAD model of contact sensorized bridge.

Non-Contact Bridge Pickup

As the lifetime of contact bridge pickup described above was limited by use, a non-contact pickup was designed to eliminate this issue. The idea was to illuminate the string from below and measure the amount of light transmitted to the top side of the string, modulated by the movement of the shadow cast by the vibrating string. By offsetting the alignment, the displacement of the string from its rest position could produce a monotonic variation in transmitted light intensity, given a displacement within the limits of normal bowing. A CAD model was designed to incorporate a pair of emitter and receiver fibres for each of the four cello strings, along with respective lenses, as shown in Fig. 6.24. The bridge and fibre mounts consisted of three distinct parts: the solid wedge-shaped bridge with mounting points and fibre exit channels, bottom lens and fibre mount, and top lens and fibre mount. With this modular design, the lens-string distance could be tuned for each string individually, in order to accommodate the differences in maximal string displacement and string diameter. Fig. 6.25 shows the bridge mounted on the cello body with a sensor head on the D-string.

6.3.6 Bow

Cello bows are usually around 72 *cm* in length [144] while the bore of a Siemens Magnetom Trio MRI scanner (the 3 *T* research at the MNI) is only 60 *cm* [244]. A shorter bow than normal was clearly required. The MRI-compatible cello bow was constructed from a 1/16th scale cello bow. An annotated diagram of a cello bow is given in Fig. 6.26⁹. The variable tension screw, which contained ferromagnetic material, was removed from the frog and notches were carved into the stick. A piece of gut rope was threaded through a hole drilled in the frog, looped over a notch in the stick, and tied off, thus tensioning the bow hair. The stick was re-cambered, in the modern style, to match the length of the bow, and the tip of the bow was truncated to reduce the bow length even further, giving it a snub-nosed appearance. The bow length, from tip to frog was 30 *cm*, with 22 *cm* of usable bow hair. Fig. 6.27 shows the MRI-compatible miniature cello bow.

The co-ordination of bowing gestures with fingering, as well as the effects of bowing

⁹Ninat Friedland. *Cello Bow Study*. Ink and graphite.

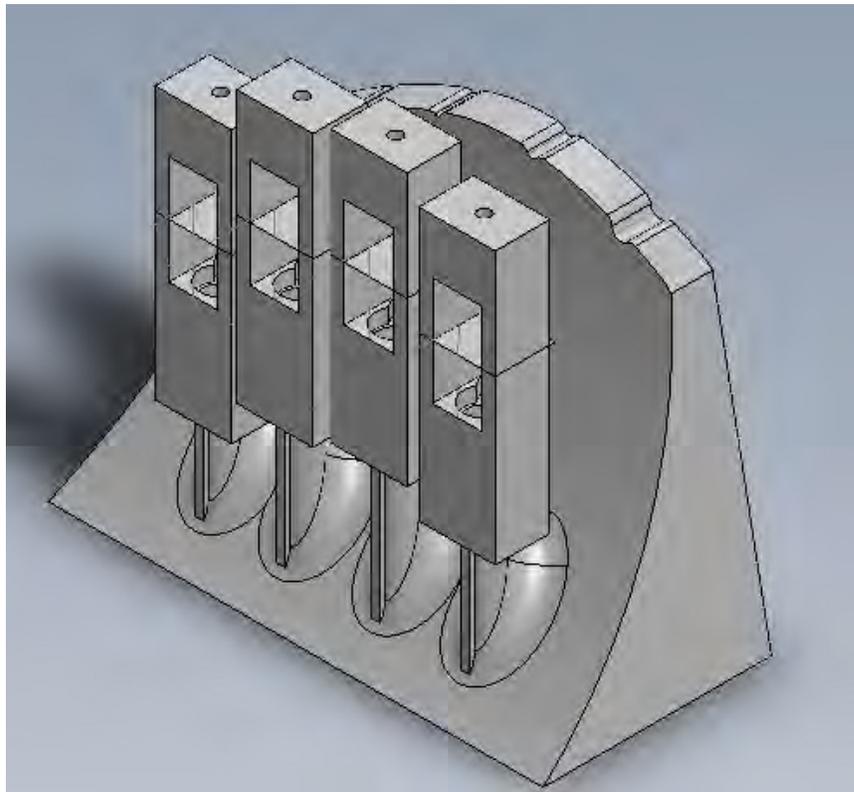


Figure 6.24: CAD model of non-contact sensorized bridge.

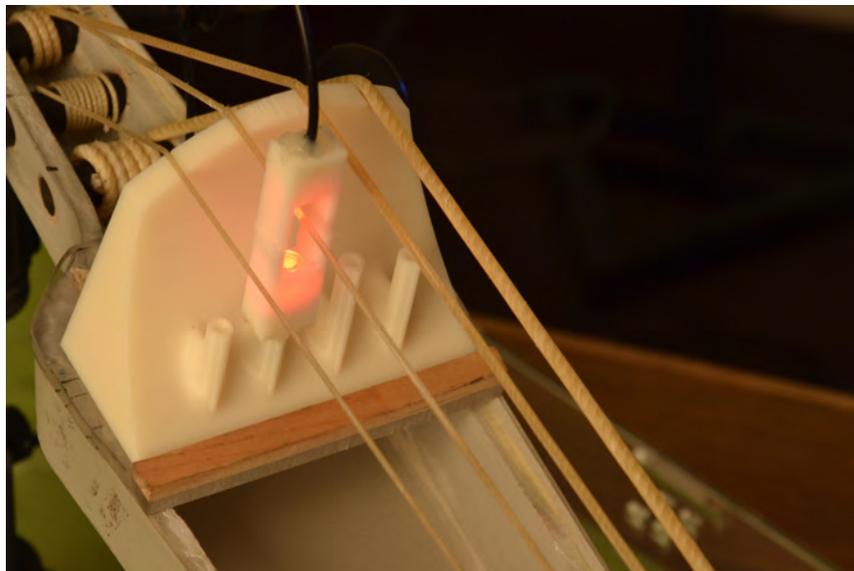


Figure 6.25: Non-contact sensorized bridge with sensor head on D-string.

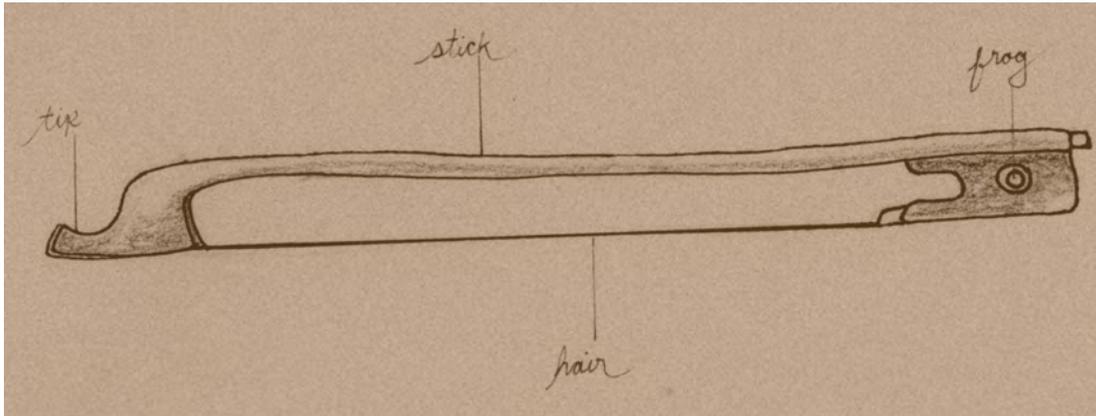


Figure 6.26: Annotated diagram of cello bow.



Figure 6.27: 30 *cm* long MRI-compatible cello bow.

gestures on the acoustic signal are of considerable interest to neuroscientists, instructors, and those interested in musical acoustics, kinematics, and simulation [6, 36, 37, 133, 199, 213, 224, 233, 235, 236, 285, 290]. Several technologies have been applied to bow (and finger) sensing of augmented cello instruments, including resistive, electromagnetic, and inertial [87, 289] techniques, though none of these is appropriate for our application. The MRI scanner presents many constraints (e.g., limited space, movement, and materials) to measuring bowing gestures, however, it allows for some novel sensing techniques as well. The designs of three different types of sensors for augmented cello bows are presented here.

Bow Force Sensor

An optical force sensor was designed and mounted between the bow stick and hair at the frog of the bow. A second sensor was mounted between the hair and the stick at the tip-end. Windings of optical fibres were embedded in a silicone toroid, with a 1 cm inner diameter and 2.5 cm outer diameter. A two-part toroidal mould was designed in CAD and 3D-printed, one half of which is shown in Fig. 6.28. The mould contained three ports, two for POF exit points and one for pouring the silicone. The number of turns of the fibre and the turn radius allowed the sensitivity to be tuned. Sensor iterations consisted of between five and ten turns, with a nominal bend radius of 8.75 mm (with a maximum range of 0.5–1.25 cm). Plastic optical fibre with 0.5 mm diameter was employed, which has a manufacturer-specified minimum bend radius of 15 mm [181]—thus, a compressive force on the sensors was expected to yield macrobending losses.

The bow force sensor casting and de-moulding process can be seen in Fig. 6.29. This sensor effectively measured the displacement of the bow hair. The force applied by the bow on the string can be calculated, as long as the position of contact along the bow is known. By characterizing the bow hair displacement-force curve at multiple points along the bow, the bow force can then be estimated, as was done by Demoucron and Schoonderwaldt [56, 236]. A detailed calibration strategy for the acquisition of violin bowing parameters was also given by Maestre et al. [164]. The force sensors can be seen at the tip and frog of the bow in Fig. 6.30, between the hair and the stick.

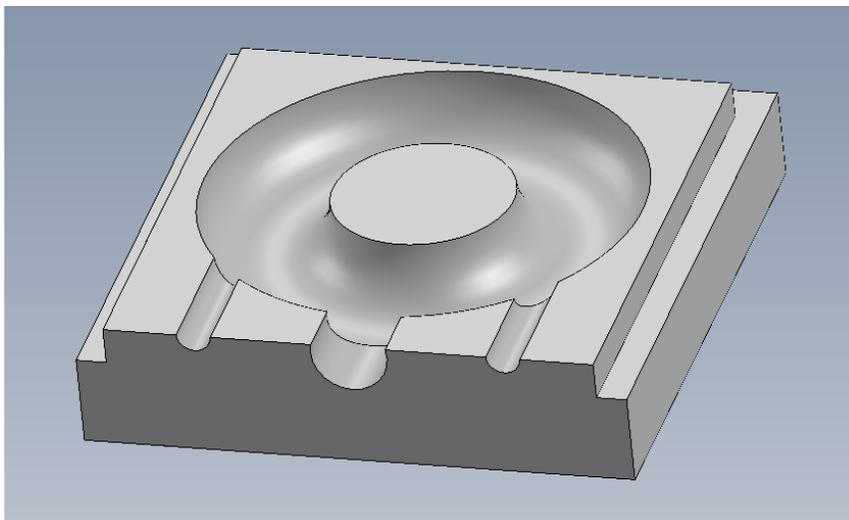


Figure 6.28: CAD model of the bottom half of the silicone force sensor mould.

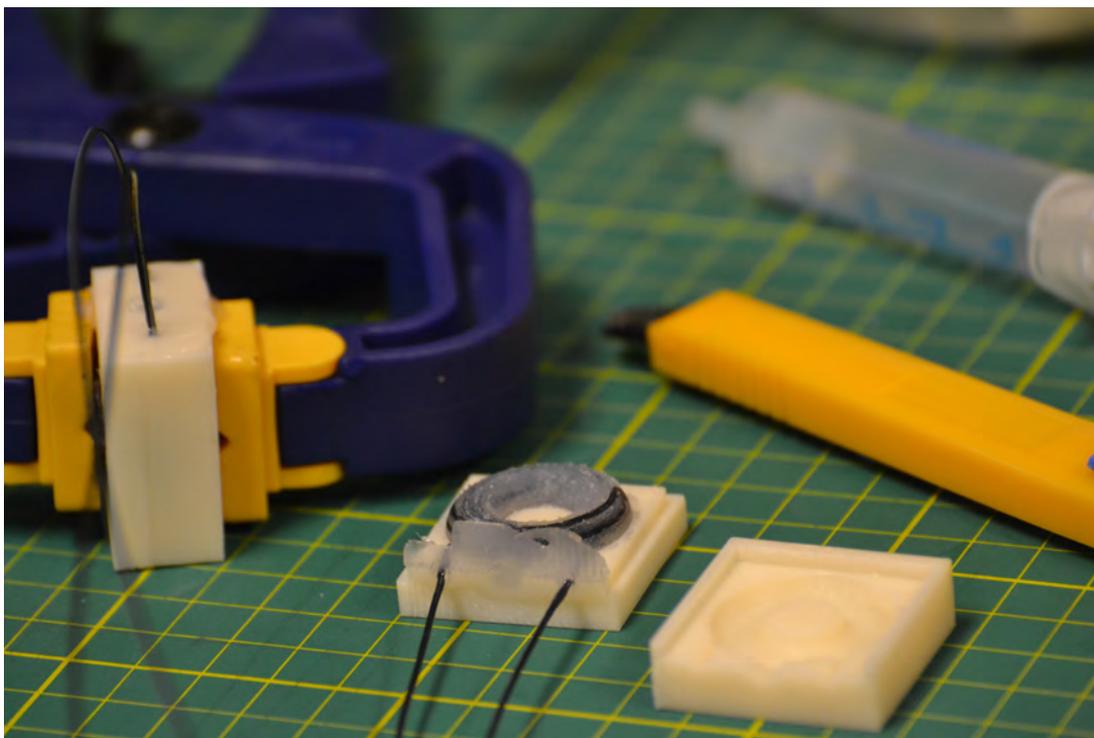


Figure 6.29: Fibre optic force sensor: curing in mould (left), de-moulding of silicone form with optical fibres (centre), and top half of 3D-printed mould (right).



Figure 6.30: MRI-compatible cello bow with bow force sensors mounted at tip and frog, and 3D magnetic field sensor.

Bow Position Sensor Using Optical Flow

Bow position sensing, in addition to enabling the calculation of absolute bow-string force, allowed the measurement of dynamic bowing gestures and their derivatives. Precisely controlled bow speed, for instance, is integral to producing *good* cello tone [224]. Articulation styles, such as *legato*, *spiccato*, and *détaché* are used to convey disparate musical ideas [290]. Each technique has very specific gestural components and measures of bow position over time can be used to provide pedagogical feedback.

A position sensor, capable of measuring bow displacement and distance from the bridge, was designed using a commercial optical flow sensor and a plastic multi-core fibre optic cable. The optical flow sensor, Avago Technologies ADNS-9500, is an optical mouse sensor meant for high-precision gaming applications. The integrated package contains a 30×30 pixel image sensor, a laser¹⁰, and a configurable chip-on-board [4]. The run-time configurable chip-on-board, itself contains an image acquisition system, a digital signal processor (DSP), and an SPI communications port. The ADNS-9500 differs from similar Avago parts, as it also contains a programmable read-only memory,

¹⁰The laser was not employed as its near-infrared spectral peak ($\lambda = [832\text{--}865] \text{ nm}$) is outside the transmissive region of the plastic optical fibre.

which doubles as a temporary storage space for full-frame image acquisition.

The sensor measures displacement of salient image features in order to compute the direction and speed of movement, known as *optical flow*. The DSP, which implements the optical flow algorithm, essentially takes successive image frames, looking for patterns of pixels of high contrast that are shared between frames, and through a correlation analysis of successive frames, measures the direction and magnitude of motion [156, 250, 265]. Normally, the sensor and target are co-located, with the illumination source almost parallel to the target plane, in order to illuminate the surface texture from the side. In the embodiment that was meant to measure bow movement inside the MRI scanner, the sensor and target would be located in different rooms, thus fibre optic cables were employed to convey light between the two. A simple, single-core fibre cannot transmit an image, as light launched into the fibre does not maintain its geometry, phase, nor coherence. Using a multi-core fibre (Mitsubishi Rayon Co. Ltd., part number LH4001-N [59]), where 151 cores share a common cladding, each core carries the equivalent of a pixel of information from one end to the other.

A lens was needed to project a focused image onto the fibre end and another lens was needed to project the image from the fibre onto the image sensor. Avago produces a moulded plastic lens that fits the ADNS-9500, which focuses at 2 mm from the lens plane. Tests were conducted with a second lens that projected light from the fibre onto the Avago moulded plastic lens. In order to mechanically stabilize the test assembly, a 3D-printed fixture was designed, as shown in Fig. 6.31 and 6.32. The fixture secured the multi-core fibre over the lens assembly and aligned it with the image sensor.

Instead of illuminating a diffusely reflective target, as is the case for optical mice, a laser source illuminated an emitting fibre that was placed on the bow while the receiving fibre and lens were placed on the bridge, as shown in Fig. 6.33.

The optical flow sensor was connected to a *fungible-2* acquisition and control board (refer to Section 3.4), using SPI, and firmware routines were created to communicate with the optical flow chip, load settings, read movement data, as well as capture full frames. The usual mouse settings were over-written with custom settings to increase the sensitivity and decrease the threshold of the image acquisition and processing algorithm. Frame rate, resolution, and power (i.e., sleep and power conservation modes) setting were maintained at their maximal settings so that the chip performed as consistently and effectively as possible. The frame capture feature was used extensively to

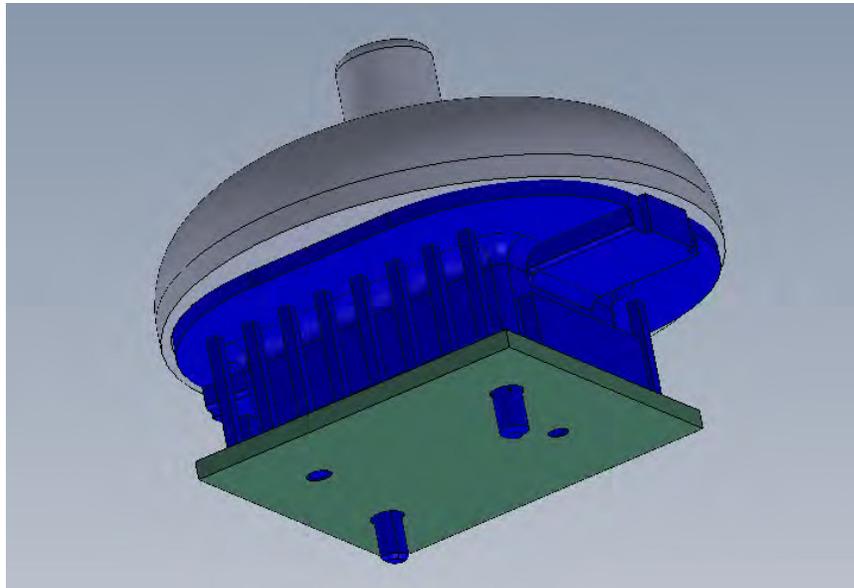


Figure 6.31: CAD assembly drawing of optical flow sensor: 3D-printed fixture (top), moulded lens (middle), and board-mounted flow sensor (bottom).

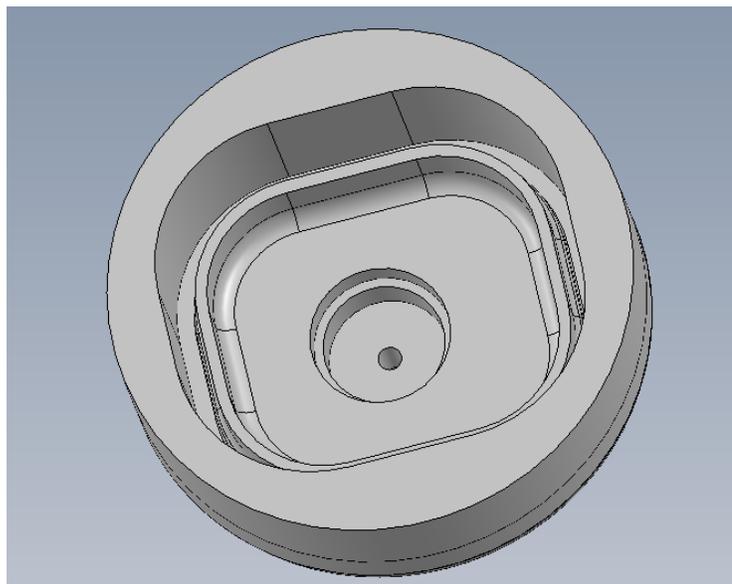


Figure 6.32: CAD model (underside) of fixture used to secure and align POF with the optical flow sensor.

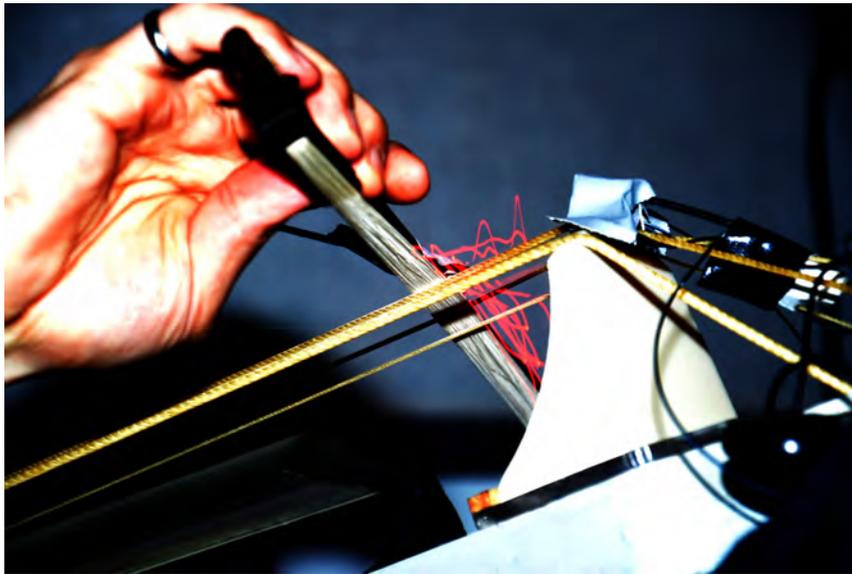


Figure 6.33: Long exposure photograph of cello bowing trajectory showing red trail produced by emitting fibre attached to bow. Receiving fibre and lens fixture mounted on bridge with distal end focused on optical flow sensor.

test alignment and to visualize the apparent image movement. The chip featured not only x - and y -axis displacement values, but also the salient feature quality and sum of pixel intensities which were useful in adjusting the optical and on-chip parameters. By placing the optical flow sensor on the bridge, with emitters on the bow, the lateral position, as well as the distance from the bridge were captured.

The optical flow sensor accrued movement data internally between microcontroller polling. The microcontroller was programmed to poll the flow sensor at regular intervals to determine if there had been any movement since the last request. If the movement flag was asserted, the accrued movement data was transferred to the microcontroller, where the movement or delta position data was added to the current position. The absolute position was reset through a command sent to the microcontroller.

Bow Position Sensor Using Faraday Rotator

In contrast with the previous set of sensors, where we have gone to great lengths to avoid interaction between the sensor and the MRI scanner environment, for position

sensing we might make use of the scanner's magnetic field to determine where a sensor is in space. The static magnetic field, while homogeneous near the isocentre, drops off, axially, at a high rate of change, in Tesla per metre, at the edges of the bore. This fringe field, while not linearly related to position, decreases maximally with a rate of 7 T/m at arm's length (i.e., 0.8 m) from the isocentre [243]. This is ideal for measuring cello bowing gestures within the MRI scanner using magneto-optic means.

The Faraday effect is a magneto-optic effect wherein the polarization of a beam of light travelling through a birefringent medium¹¹ is rotated as a consequence of the application of an external magnetic field in the direction of propagation [67]. Thus, by analyzing the degree of rotation of a linearly polarized beam of light, one can measure the local magnetic field intensity. The relation that governs Faraday rotation is given by:

$$\theta = \mathcal{V}Bl \quad (6.1)$$

where θ is the angle of rotation, \mathcal{V} is the wavelength-dependent Verdet constant of the medium, B is the longitudinal component of the local magnetic field, and l is the path length through the medium [112]. The Verdet constant is relatively low in most substances: in air, $\mathcal{V} = 0.0019 \text{ rad/Tm}_{\lambda=633 \text{ nm}}$ [123]; in water $\mathcal{V} = 3.35 \text{ rad/Tm}_{\lambda=632 \text{ nm}}$ [278], and in silicate glass $\mathcal{V} = 3.78 \text{ rad/Tm}_{\lambda=633 \text{ nm}}$ [210]; however, in specialized optical crystals such as terbium gallium garnet (TGG) the Faraday effect is more than an order of magnitude greater, with a Verdet constant of $\mathcal{V} = -133.9 \text{ rad/Tm}_{\lambda=632 \text{ nm}}$ [279].

A Faraday rotator was designed using a 1 cm long, 5 mm diameter TGG crystal sandwiched between two linear polarizers. A plastic lens was used to focus a 660 nm beam launched into a 1 mm POF into the optical assembly. A second lens and fibre were used to receive the light and convey it to a *fungible-2* optoelectronic acquisition board (see Section 3.4) within the control room. Fig. 6.34 shows a schematic view of the Faraday rotator. An initial prototype shown in 6.35, using commercial non-ferromagnetic (aluminium) lens tubes and a 3D-printed ST mating connector, was used to test the feasibility of the sensor. The sensor was initially tested in the lab by passing

¹¹A birefringent medium is one in which left- and right-hand circularly polarized beams of light propagate at slightly different speeds. A beam formed by left- and right-hand circularly polarized components of equal magnitude is, through the property of superposition, a linearly polarized beam of light.

it over a rare earth magnet with a bowing-like gesture, the results of which are shown in Fig. 6.36. The sensor was taken into the MRI scanner and bowing-like gestures were captured, as shown in Fig. 6.37.

Clearly, the sensor held promise, but the eddy currents created by moving the metallic lens tube across the magnetic field lines, caused significant dampening forces, anticipated according to Lenz's law. Thus a CAD model was designed for a lens tube which integrated the TGG crystal, linear polarizers, lenses, and alignment slots for the POF. Fig. 6.38 shows the 3D-printed Faraday rotator housing with optical elements exposed. The lens tube was also designed in a modular way, such that three such sensors were snapped together with orthogonal alignment, allowing for measurement of the magnetic field in three dimensions, as seen in Fig. 6.39. The cello bow with 3D-printed magnetic field sensor is shown in Fig. 6.30.

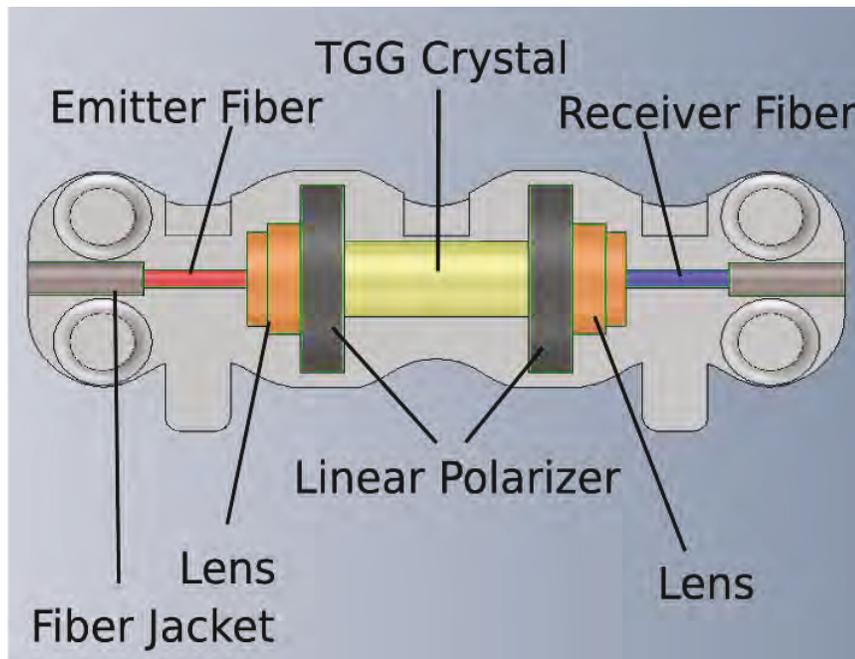


Figure 6.34: Faraday rotator schematic showing optical fibres, lenses, linear polarizers, and high Verdet constant terbium gallium garnet [115].



Figure 6.35: Prototype Faraday rotator in aluminium lens tube.

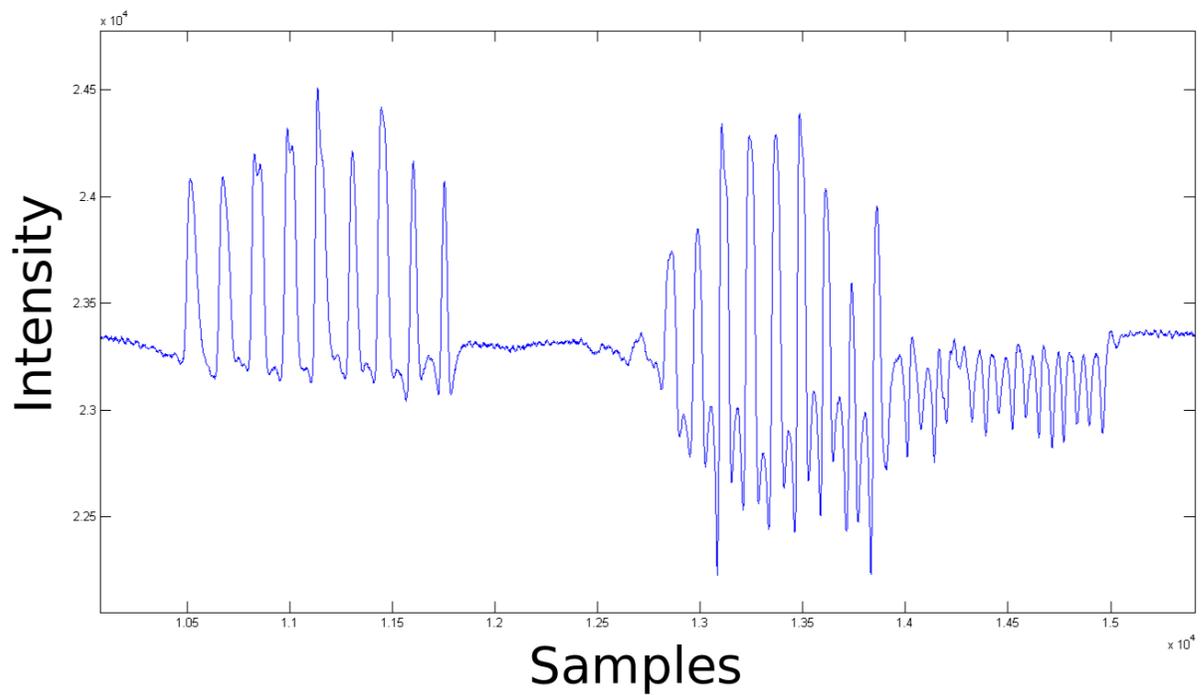


Figure 6.36: Proof-of-concept Faraday rotator moved past a niobium rare earth magnet with a bowing-like gesture (500 Hz sampling rate).

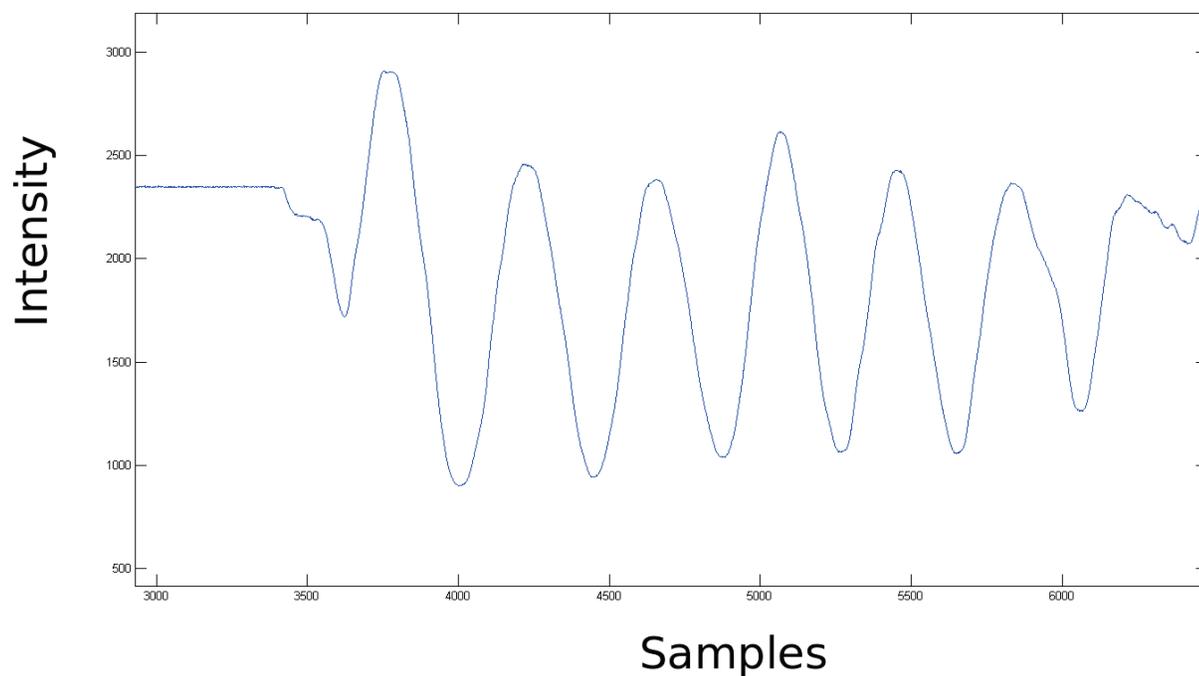


Figure 6.37: Bowing-like gestures performed in 3 T MRI scanner measured with proof-of-concept Faraday rotator in aluminium lens tube (500 Hz sampling rate).

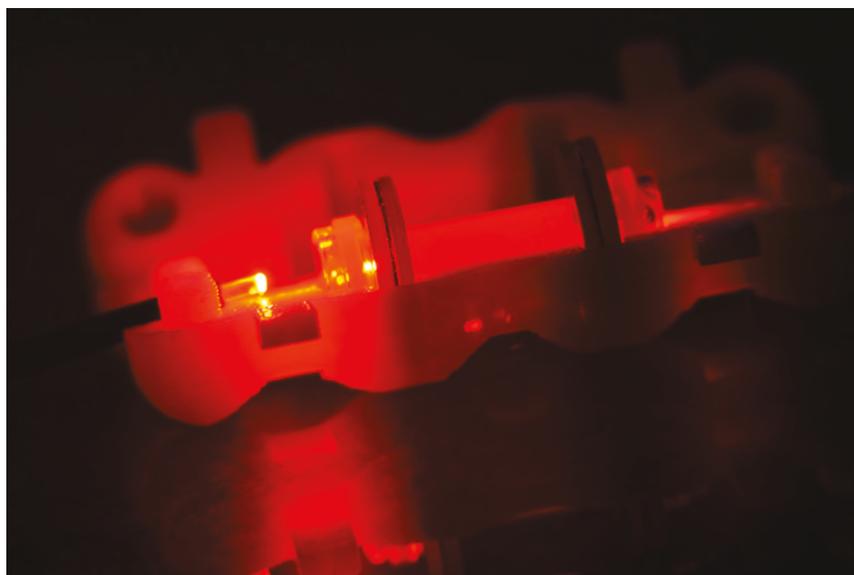


Figure 6.38: Magnetic field sensor with optical elements and 3D-printed housing.

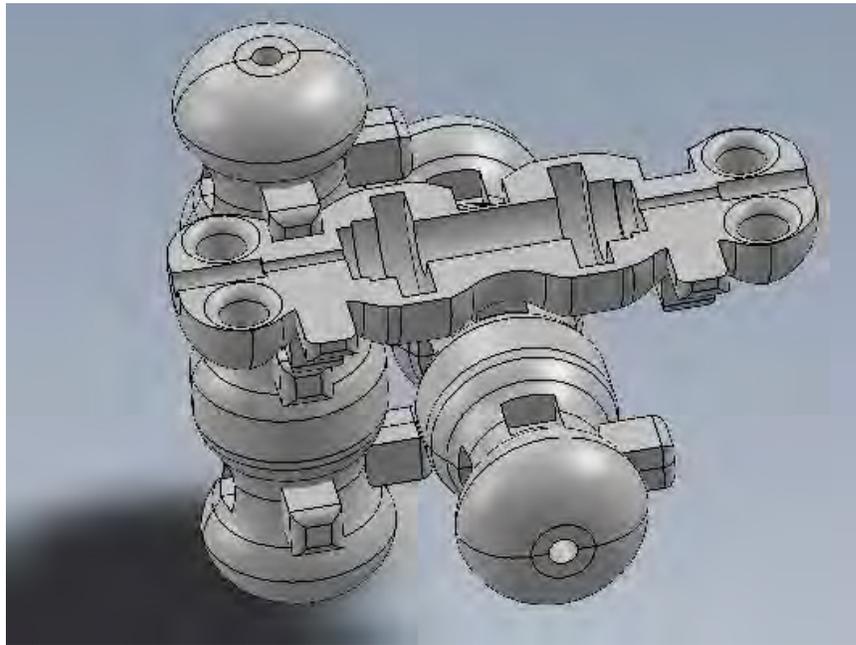


Figure 6.39: CAD assembly drawing of 3D magnetic field sensor.

6.3.7 MRI-Compatible Cello Support Fixture

A specially designed fixture was built to support the cello inside the MRI scanner to reduce strain on the cellist. Whereas a cello would normally be supported by the end-pin and the cellist's legs, this was not feasible in the scanner. The cello, lacking a proper body, needed to remain steady in response to the torques applied with the bow. As well, by keeping the cello elevated and at a comfortable playing position, we were able to reduce fatigue, exacerbated by the cellist's supine playing position wherein the cellist is working against gravity in ways that are quite different than the normal, upright, cello posture.

The fixture, shown in Fig. 6.40, was designed with a base that lay flat against the scanner bed below the cellist's legs. The peg box of the cello was inserted into the fixture head and secured by the over-hanging lip. The fixture head was adjustable in height, using a telescopic arm. Atop the arm was a ball joint socket fixed with a screw-collar. The base was made of acrylic, while the other parts were all, initially, 3D-printed. Due to repeated material failures of the ball joint while it was being tightened, as seen in Fig. 6.41, the ball joint was milled out of Delrin. The fixture, supporting the

cello, is shown in Fig. 6.42.

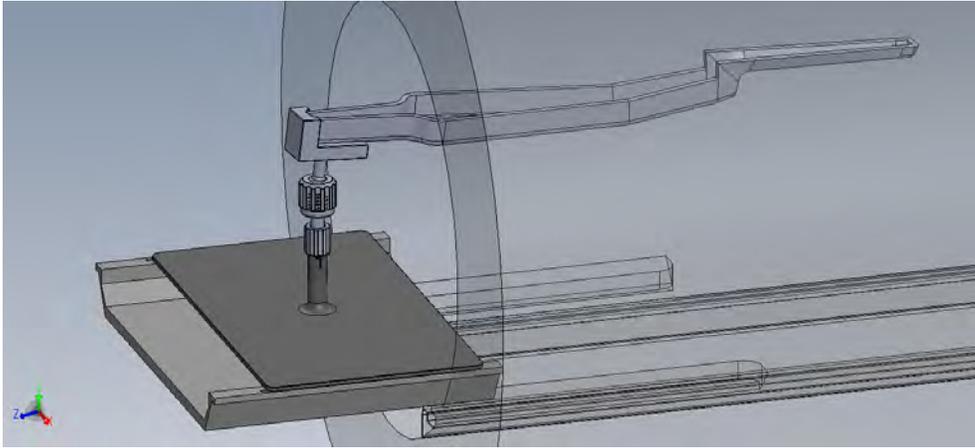


Figure 6.40: CAD assembly drawing of cello support fixture with cello body in scanner.



Figure 6.41: Material failure of 3D-printed ball joint.



Figure 6.42: Cello supported in MRI-compatible fixture.

6.4 Testing and Results

The cello and its sensors were tested at various phases of development outside and inside the MRI scanner. It is important to note that no image artifacts were encountered during MRI testing of the cello and its components in a 3 T scanner. The proof-of-concept magnetic field sensor having an aluminium lens tube, which was meant only as a demonstration of the sensing technique and never meant for actual behavioural use, was the only metallic component used and demonstrated the counter electromotive force of a moving conductor within a magnetic field. Otherwise, through the use of non-ferromagnetic and non-conductive materials, a completely MRI-compatible interface with optical sensing was designed, built, and tested. This section presents the results of tests conducted with the cello sensors using cello-like performance gestures or during actual cello performance where possible.

6.4.1 Fingerboard Sensors

Cast Transparent Urethane Fingerboard

The transparent urethane fingerboard, with cast POF, was tested outside of the scanner using the *fungible-2* acquisition system (refer to Section 3.4) by repeatedly pressing the finger against the fingerboard. The finger-to-fingerboard interaction was captured in low ambient light conditions to reduce noise caused by environmental sources (e.g., fluorescent lights, computer screens, etc.). A plot of the results are shown in Fig. 6.43, with an SNR of 43.8 dB.

3D-Printed Fingerboard

The 3D-printed fingerboard was tested under similar conditions as the cast fingerboard with the same equipment and test procedure. The results are shown in Fig. 6.44, with the lower values corresponding to string-fingerboard contact. SNR for this fingerboard was measured between 12.7 dB and 16.9 dB, depending on the specific sensor, with the variation in performance owing to the optical quality of the POF-fingerboard adhesive bond.

The 3D-printed fingerboard was also tested inside the scanner during a cello shifting task. The traces of Fig. 6.45 show data from pairs of sensors, each pair was centred

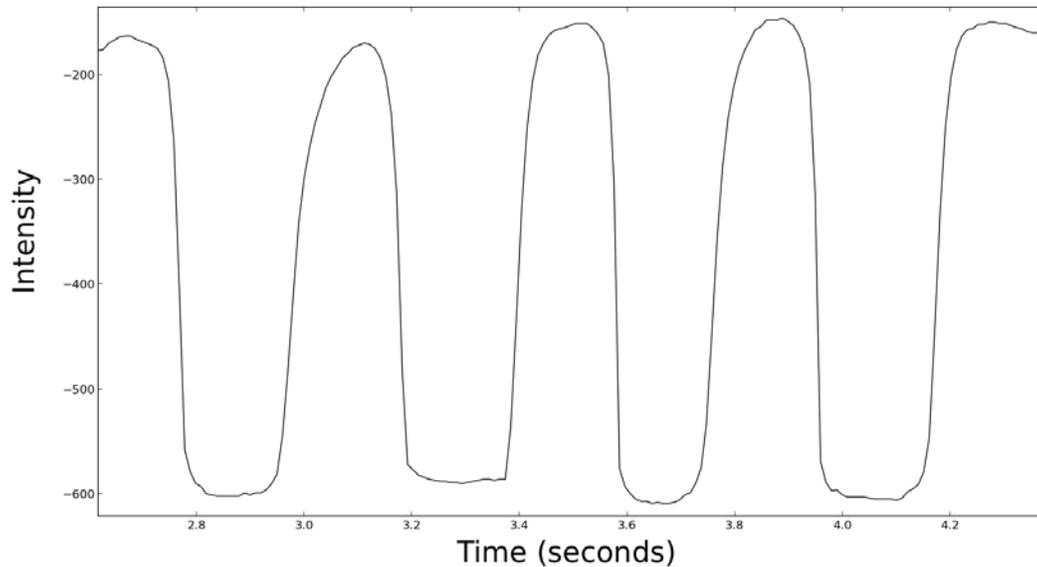


Figure 6.43: Results of repeated finger-string-fingerboard contact with cast transparent urethane fingerboard (100 Hz sampling rate).

around the position on the fingerboard corresponding to in-tune and neighbouring semitones. In these plots the finger started centred on the D^\sharp sensor pair and moved over and past the E and F sensor pairs, to land on the F^\sharp sensor, resulting in a shift up of approximately one and a half semitones.

6.4.2 Bridge Pickup

Contact Bridge Pickup

The contact bridge pickup was tested outside the scanner, using the transimpedance amplifier of the acquisition system described in Section 3.4 as a preamp. The signal was then buffered on-chip and recorded using a 744T professional-quality portable four-track recorder, made by Sound Devices, LLC at a sampling rate of 48 kHz. Fig. 6.46 shows a plot of both up-bow and down-bow strokes on the open D-string. The sensor had to be adjusted into a bending loss regime by creating a tight loop with the 0.5 m POF, placed between the string and the bridge, and tuning the cello to provide sufficient pressure on the sensor. A decreased bend radius (i.e., one smaller than the

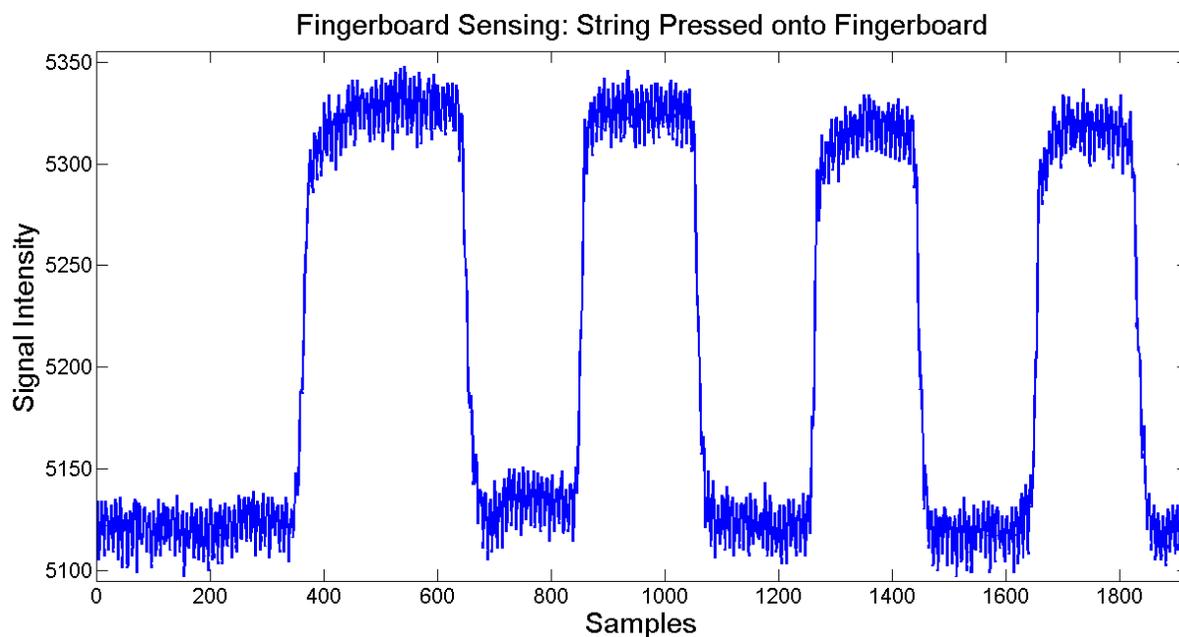


Figure 6.44: Results of repeated finger-to-fingerboard contact with 3D-printed fingerboard (500 Hz sampling rate) [115].

absolute minimum given by the manufacturer) enabled a larger variation in signal, yet the sensor itself began to fail under wear within minutes. Bow pressure, which was measurable using this method, also slightly affected the sensitivity of the sensor.

Non-Contact Bridge Pickup

The non-contact bridge pickup was tested outside and inside the scanner, using the buffered transimpedance amplifier setup, as described in Section 6.4.2, with the addition of digitizing the signal on-board the *fungible-2* and relaying that to a laptop computer over USB, in order to collect the low-frequency component of the string displacement in addition to the vibration. Fig. 6.47 shows a plot of a down-bow on the open D-string. Up- and down-bows could be distinguished, though the asymmetry showed some dependence on sensor-head alignment with the string. A perfectly symmetric placement of the string would cause non-monotonic perturbations in sensor amplitude, if the string's vibration were small and centred around the medial axis. The bowing action, however, caused the bow to pull the string off-centre and vibrate

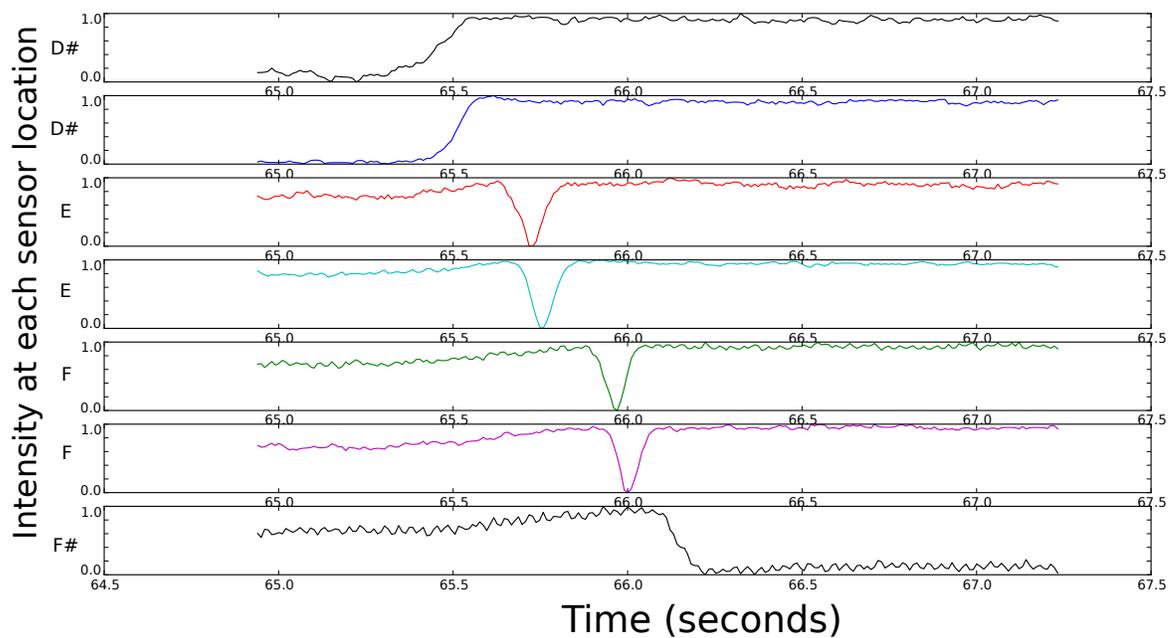


Figure 6.45: Finger-string-fingerboard interaction, showing a shift from D^\sharp to F^\sharp on the D-string, as sensed with the 3D-printed fingerboard [239].

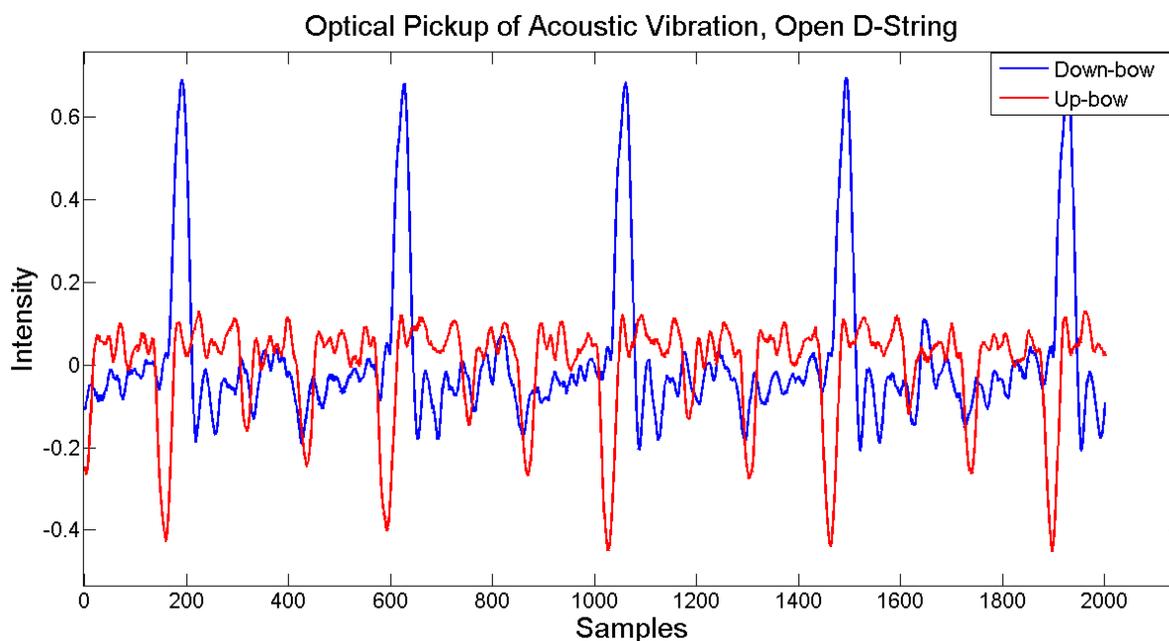


Figure 6.46: Contact optical bridge pickup signal, up- and down-bowing on open D-string (48 kHz sampling rate) [115].

around a shifted bias point. This is demonstrated at the top of Fig. 6.48, where the string displacement can be seen to shift with alternating down- and up-bows.

6.4.3 Bow Sensors

Bow Force Sensor

A 7-turn toroidal force sensor was characterized by compressing the sensor in 1.59 mm ($\frac{1}{16}$ inch) increments while measuring the load using a precision balance scale as well as noting the average sensor values acquired with a *fungible-2* optoelectronic acquisition board. The process was performed twice—once with load increments and once with load decrements, to account for hysteresis—and the results were averaged. Fig. 6.49 shows the force measurements and average sensor data plotted against the compressive strain. Fig. 6.50 shows the sensor's force response with a characteristic nonlinearity at very low and high forces. The sensor shows good linearity, however, in the 1–5 Newton range.

The toroidal silicone fibre optic force sensor was tested using a *fungible-2* optoelec-

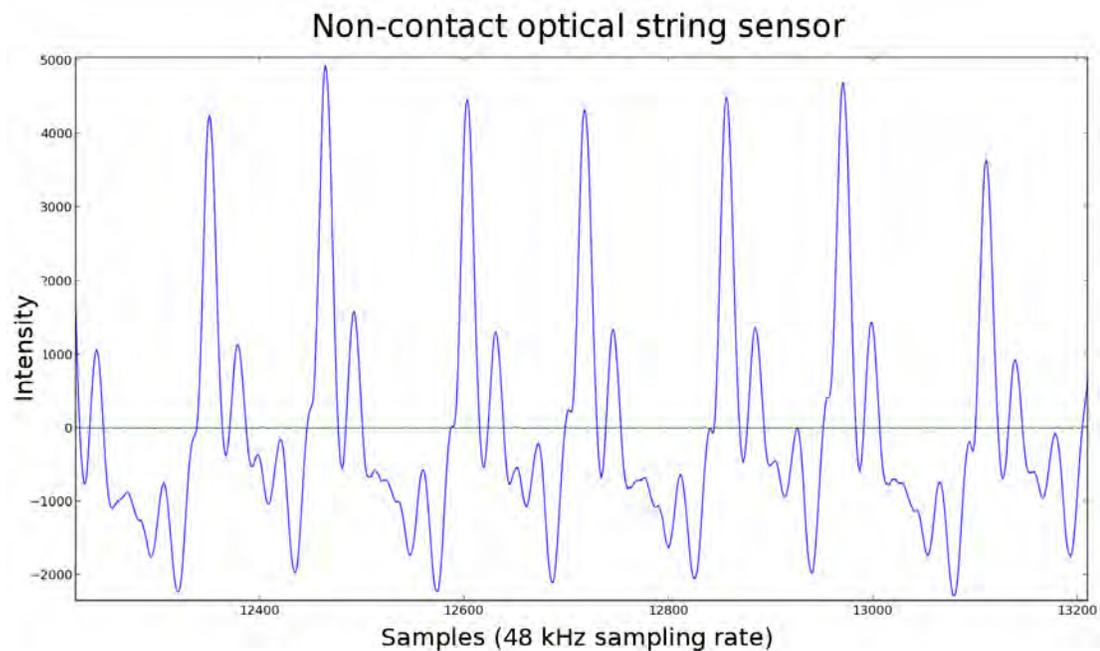


Figure 6.47: Non-contact optical bridge pickup, down-bow on open D-string (48 kHz sampling rate) [115].

tronic control and acquisition board inside and outside of the MRI scanner. The sensors were mounted between the bow hair and the stick, one at the frog and another at the tip. Pressure exerted on the bow against the string resulted in compression of the sensor, as the stick and hair were brought closer together, thus reducing the effective bend radius and causing bending losses in the fibre. The sensors were pre-loaded with a compression of approximately 3 mm when mounted in the bow, in order to make use of their more linear range.

Pressure at the frog is greatest at the start of down-bows, while pressure at the tip is greatest at the start of up-bows. This is elucidated in Fig. 6.51 showing a down-bow, the resulting string displacement, and the measured force at the frog and tip, as well as the calculated total bow force and estimated bow position between the frog and tip. It is interesting to note that the transfer of pressure between the frog and the tip can clearly be seen at around 65.5 seconds, as the bow-string contact point moves past the middle of the bow.

A Polhemus Liberty electromagnetic motion tracking system [205] was used to determine the accuracy and linearity of estimating the bow-string contact position from

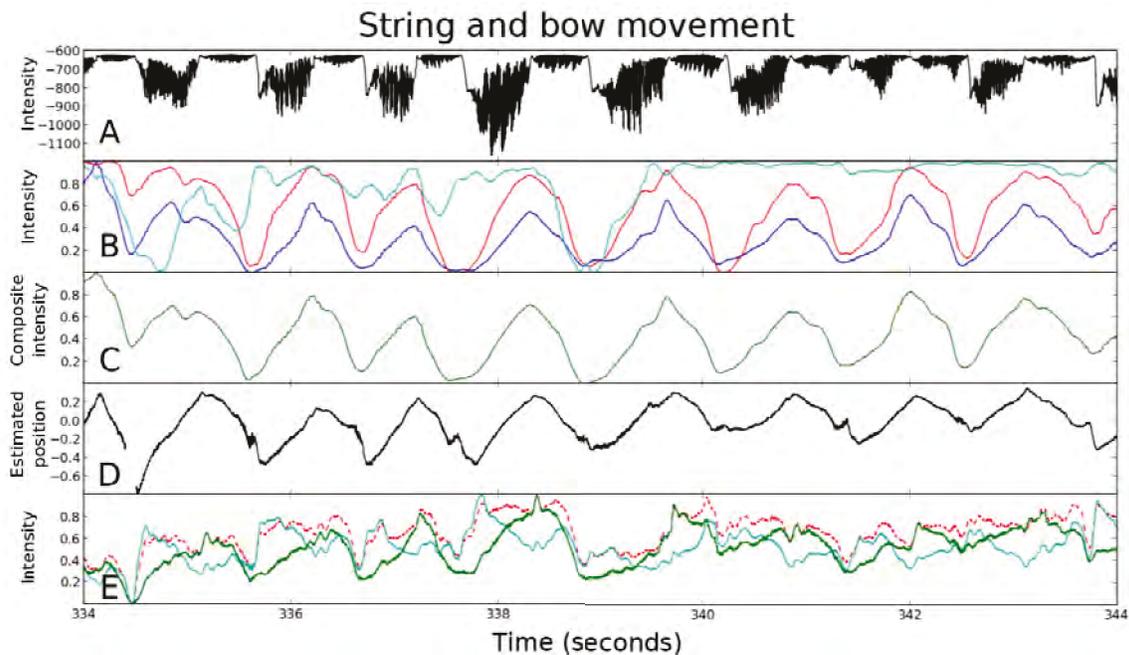


Figure 6.48: Synchronous trace of a cellist performing a repeated bowing gesture in MRI scanner: (A) string vibration using non-contact optical pickup, (B) three orthogonal magnetic field components measured using 3D Faraday rotation sensor, (C) principal component of magnetic field sensors, (D) estimated bow-string contact position calculated between frog and tip, (E) frog (cyan), tip (green), and total (dashed) bow pressure measured using two coils of plastic optical fibre embedded in silicone. Normalized data shown in plots (B) through (E) [115].

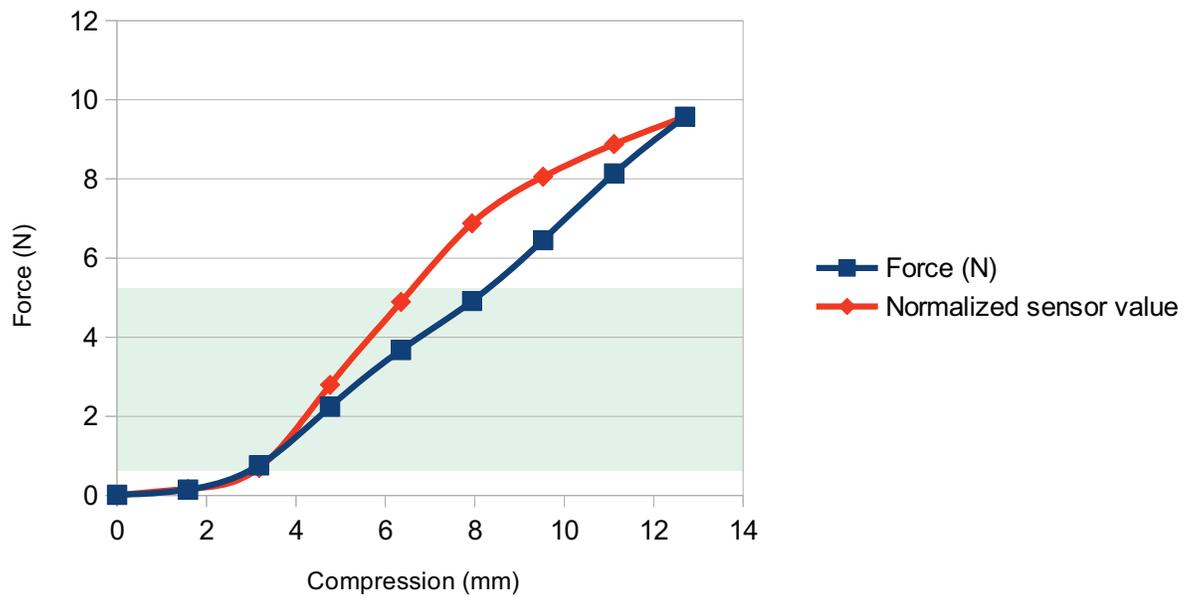


Figure 6.49: Characterization of Toroidal force sensor under compressive strain, with shaded area highlighting region of linear response.

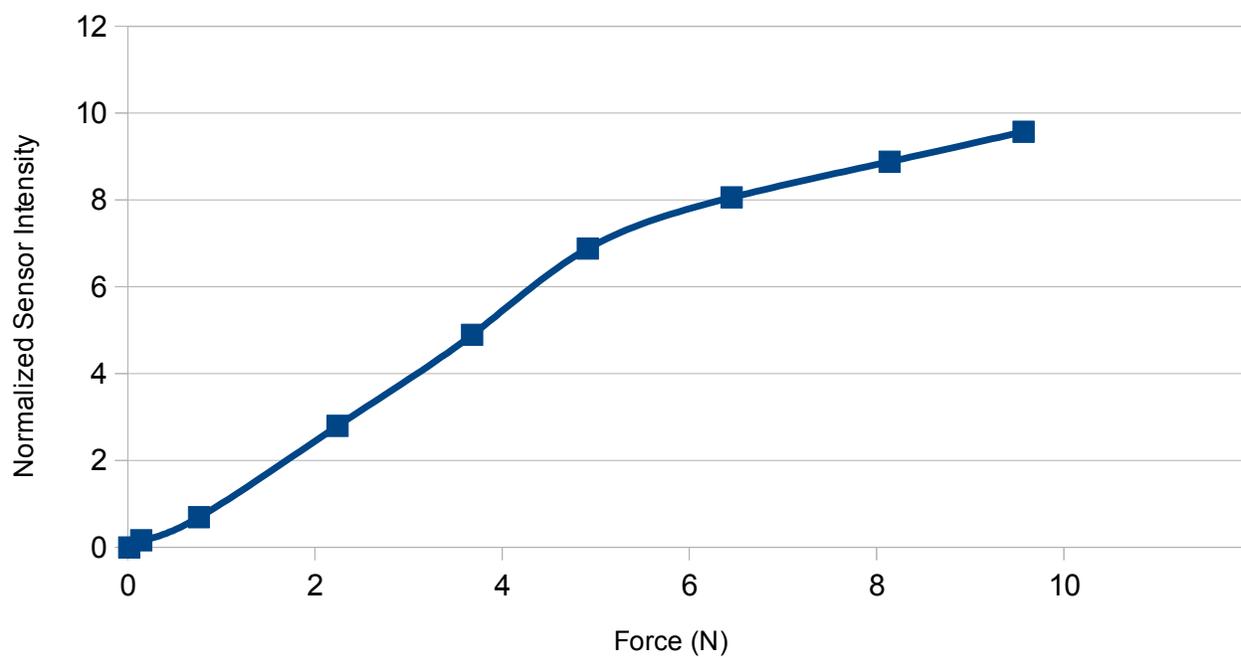


Figure 6.50: Characterization of toroidal force sensor nonlinearities under compressive strain.

the optical force sensors mounted in the frog and tip of the bow. In addition to the optical force sensors, two Polhemus tracking markers were mounted at the frog and tip of the bow. A third, static, tracking marker was mounted on a stationary surface in front of the Polhemus emitter. Polhemus Liberty motion capture and optical sensors signals were captured simultaneously. With a range of force similar to that used in cello performance, the bow was brought across the static marker, from frog to tip. The bow position was measured with the Polhemus Liberty alongside the bow-hair deflection force using the toroidal silicone optical sensors, acquired with a *fungible-2* optoelectronic acquisition board. The bow contact position of the bow and the static marker, as a function of the distance between the frog and tip sensors, was then calculated as a normalized difference between the optical sensor values. The relative bow displacement was calculated as an average of the displacement of the Polhemus markers mounted on the frog with respect to the static marker. Fig. 6.52 shows the estimated bow position along with the position measured using the Polhemus as a function of time. The root-mean-square error for this example was 2.3%. Fig. 6.53 shows the linearity of the estimate with respect to the measured bow position, with the greatest nonlinearity where the weight is transferred from frog to tip. The root-mean-square error could be improved by linearizing this transfer function using a lookup table determined through calibration.

Optical Flow Sensor

The optical flow sensor was connected to a *fungible-2* optoelectronic control and acquisition board and accessed over SPI. A 1 mm emitter fibre was mounted on the bow, connected at one end to a red laser diode ($\lambda = 650 \text{ nm}$, 1 mW) and directed at a lensed multi-core receiver fibre mounted on the bridge. The opposite end of the receiver fibre was focused on the optical flow sensor, conveying the spot of light hitting the receiver fibre onto the two-dimensional pixel array. The bow was drawn back and forth over the strings while maintaining alignment of the emitter and receiver fibres. A plot of the resulting measurement is shown in Fig. 6.54. In addition to the one-dimensional “horizontal” bow translation, the bow-bridge distance was also measured. The distance between the bow and bridge modulated the “vertical” position of the spot of light projected onto the receiver, as the angle between the emitter and receiver fibres

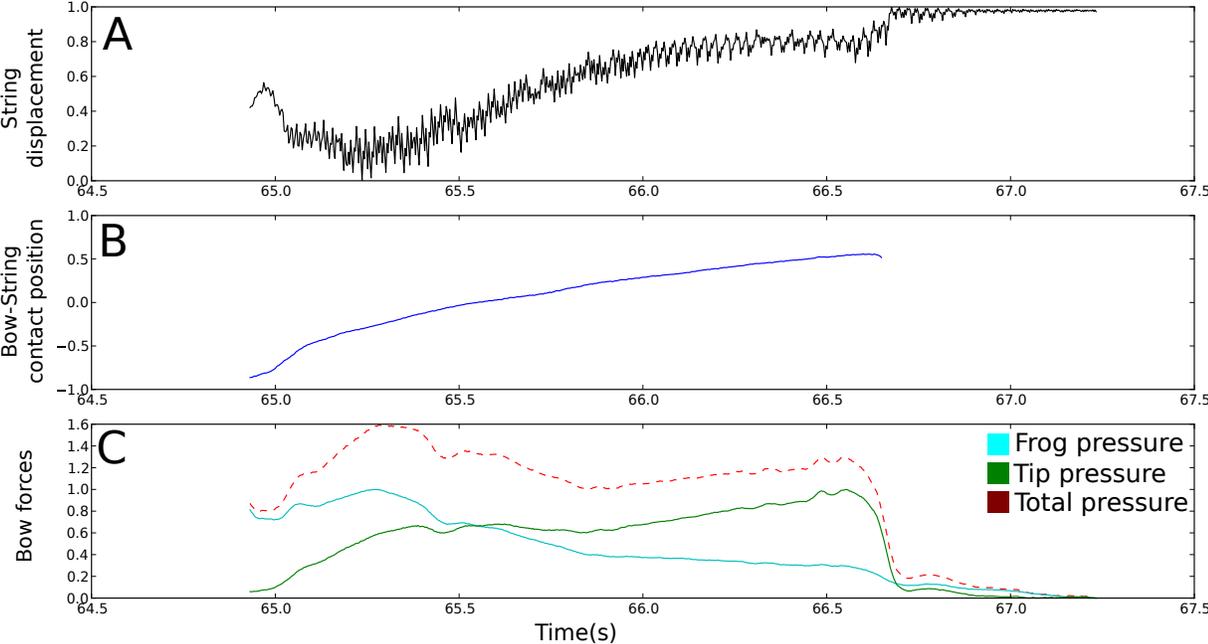


Figure 6.51: Bow-string interaction during down-bow: (A) string displacement using bridge pickup; (C) frog (cyan), tip (green), and total (dashed-red) bow force; and (B) estimated bow position, between frog and tip [239].

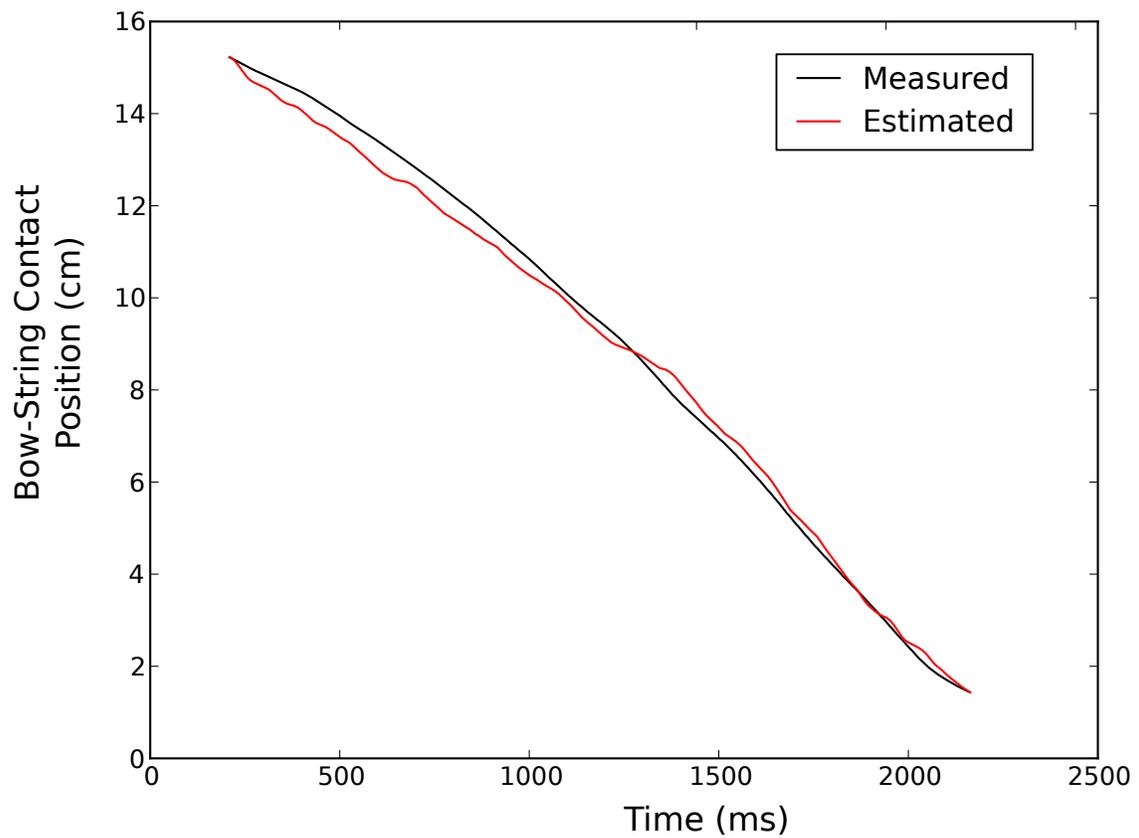


Figure 6.52: Measured (Polhemus Liberty) and estimated (optical force sensors) bow displacement during down-bow. Slope of curve indicates bowing velocity, and variance of actual human gesture.

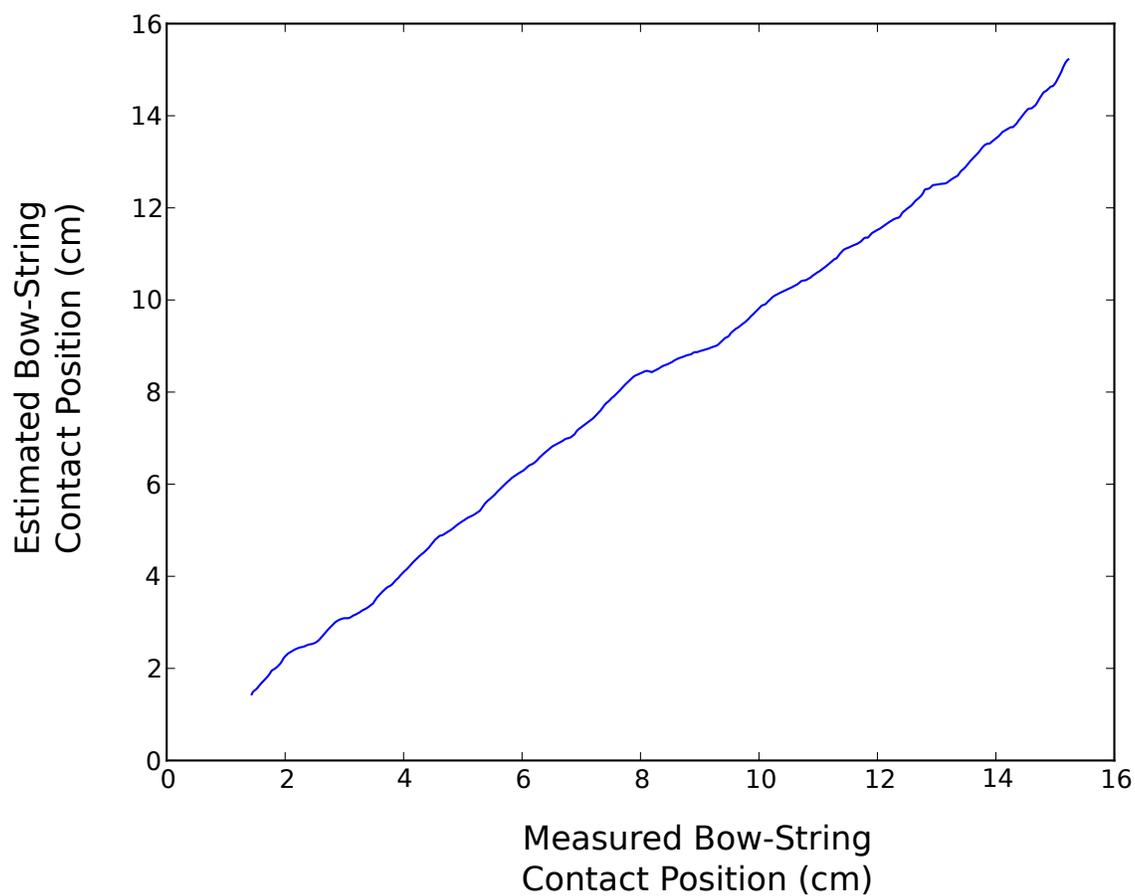


Figure 6.53: Linearity of measured (Polhemus Liberty) and estimated (optical force sensors) bow displacement.

was maintained.

As the optical flow sensor reported only the relative displacement since the previous report, the sensor tended to drift with large gestures and the accrued position calculated and stored on the microprocessor had to be reset quite often. It was also difficult to maintain alignment of both lensed sensor heads while performing natural cello bowing gestures. The sensor provided a measure of signal “quality” which was useful for judging the accuracy of the captured gesture as a whole, as a type of confidence measurement. A small graphical user interface was developed to visualize the sensed horizontal and vertical positions which were mapped to the position of a square on the screen using Cartesian co-ordinates, while the signal quality was represented by the size and colour of said square.

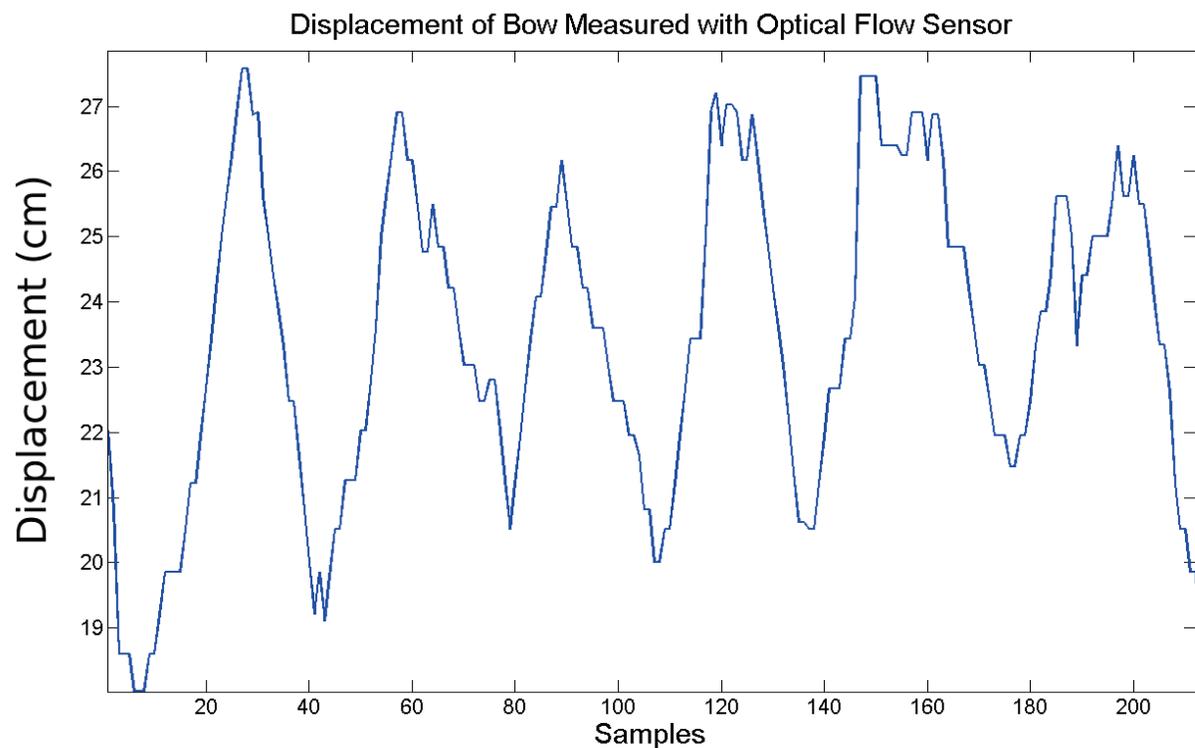


Figure 6.54: Bowing gesture captured with optical flow sensor (30 Hz sampling rate) [115].

Magnetic Field Sensor

The magnetic field sensor was tested inside the MRI scanner. Fig. 6.55 shows a cellist in the MRI scanner with the sensorized MRI-compatible optoacoustic cello.



Figure 6.55: MRI-compatible cello and cellist in MRI scanner [115].

The 3D magnetic field sensor mounted on the bow was brought into the MRI scanner and placed on the fully-retracted patient table, centred in the head coil. The patient table was then moved 1252 *mm* into the bore of the scanner, in about 6 seconds. The analog signal from the optical magnetic field sensor is shown in 6.56, as well as the rotation of polarization angle as the sensed magnetic field was increased.¹² The combined

¹²The light intensity through crossed polarizers is given by Malus' law: $I = I_0 \cos^2(\theta)$, where I is the output light intensity, I_0 is the input light intensity (in this case the intensity of the linearly-polarized beam), and θ is the angle of polarization of the outgoing wave with respect to the analyzer's primary axis. The phasor angle then has to be unwrapped to infer the rotation past $\pi/2$ radians. As the crossed polarizers of the sensor were offset by 45 degrees in order to provide a mid-scale output in the absence of

SNR of the sensor and acquisition electronics was measured at 62.2 dB.

Using the approximation given in [278], $V = 0.121 + 0.056 \cdot \lambda^{-2} + 0.0292 \cdot \lambda^{-4}$ at $\lambda = 660$ nm, the Verdet constant was estimated as $V = -117.403$ rad/Tm, and given a TGG crystal of length 11 mm and a 3 T magnetic field, the polarization angle ought to undergo a change of 222 degrees. In other words, the signal intensity would range from a full-scale signal to its minimum, or vice versa, equivalent to a 90 degree rotation, for every magnetic field change of 1.216 T. With the sensor placed on the fully-retracted patient table, centred in the head coil, the magnetic field was estimated to be 0.694 T through an interpolation of data points given in [243]. Thus, a change in 2.306 T would have been seen as the patient table was moved into the scanner bore, resulting in a change in polarization angle of 170.6 degrees. Results showed rotation of 172 degrees, as seen in Fig. 6.56, which translates into a 2.32 T change in magnetic field.

It can be seen that the sensor provided an ambiguous output for large changes in magnetic field and thus was too sensitive for measuring field strength changes in excess of 1.2 T (equivalent to a rotation of 90 degrees). However, the change in magnetic field within the scanner as seen by the sensor as it was moved radially (orthogonal to the static magnetic field oriented along the z-axis) and located at arm's length from the isocentre, was well within the sensor's range. An SNR of 62.2 dB and a 1.2 T dynamic range results in a magnetic field resolution in excess of 1 mT. Fig. 6.48 shows bowing gestures captured with the magnetic field sensor during actual cello performance in the MRI scanner.

6.4.4 MRI-Compatibility

In collaboration with Melanie Segado and Robert Zatorre, magnetic field maps of an MRI phantom, as described in Section 5.5, were acquired with and without the cello and bow (placed next to the imaging volume just outside the head coil) in order to quantify any changes in field homogeneity. A cylindrical phantom was imaged with a Siemens Magnetom Trio 3 T (0.1 ppm) scanner using a gradient field mapping sequence ($TR = 1000$ ms, $TE_1 = 2.99$ ms, $TE_2 = 5.45$ ms). The magnitude and phase images of

a magnetic field, the output signal was clipped as the intensity increased (which corresponded to more negative values) causing the photodiode-transimpedance amplifier circuit to saturate. Due to clipping of the magnetic field intensity signal between 4–5 s, the angle was computed by cubic interpolation, thus the slight angle artifact at about 4.6 s.

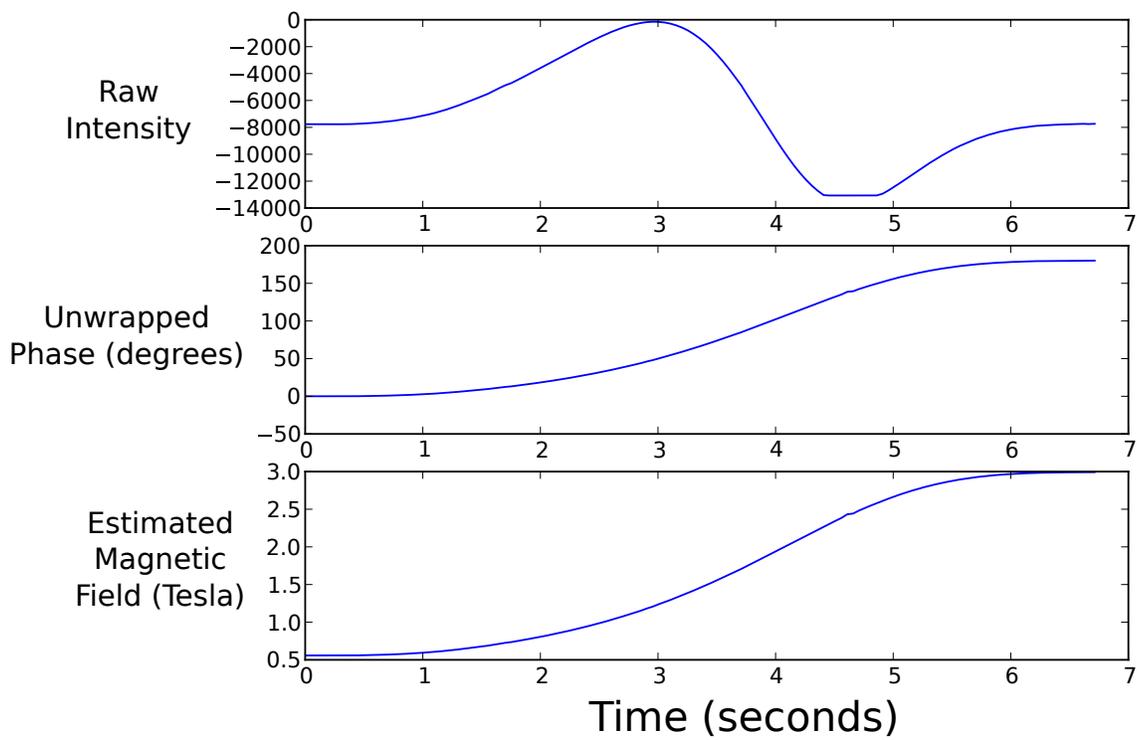


Figure 6.56: Measurement of the MRI scanner's static magnetic field (Z-axis) as a function of time (in seconds) as the sensor was moved into the scanner on the patient table: (*top*) raw light intensity from TGG sensor, (*middle*) rotation of polarization angle in degrees, and (*bottom*) magnetic field in Teslas.

the phantom, without and with the cello and bow, are shown in Figures 6.57 and 6.58; no differences are visible. The results are summarized in Table 6.1. The cello does not disrupt the field homogeneity and no significant changes in images were detected with and without the cello in the scanner. Thus, the cello is MRI compatible. A field mapping scan was also performed with a subject, as shown in Fig. 6.59, and no image artifacts were detected.

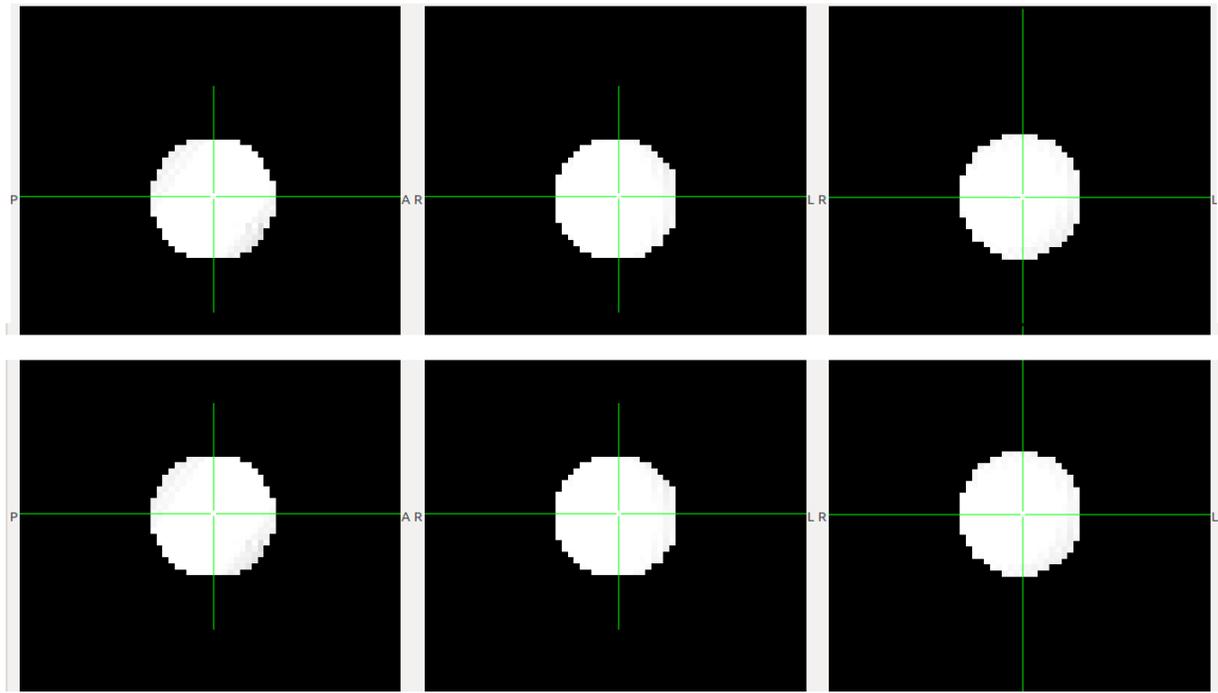


Figure 6.57: Gradient magnetic field maps of MRI phantom (masked magnitude images): (*top row*) without device, and (*bottom row*) with cello and bow in scanner bore. No differences are visible.

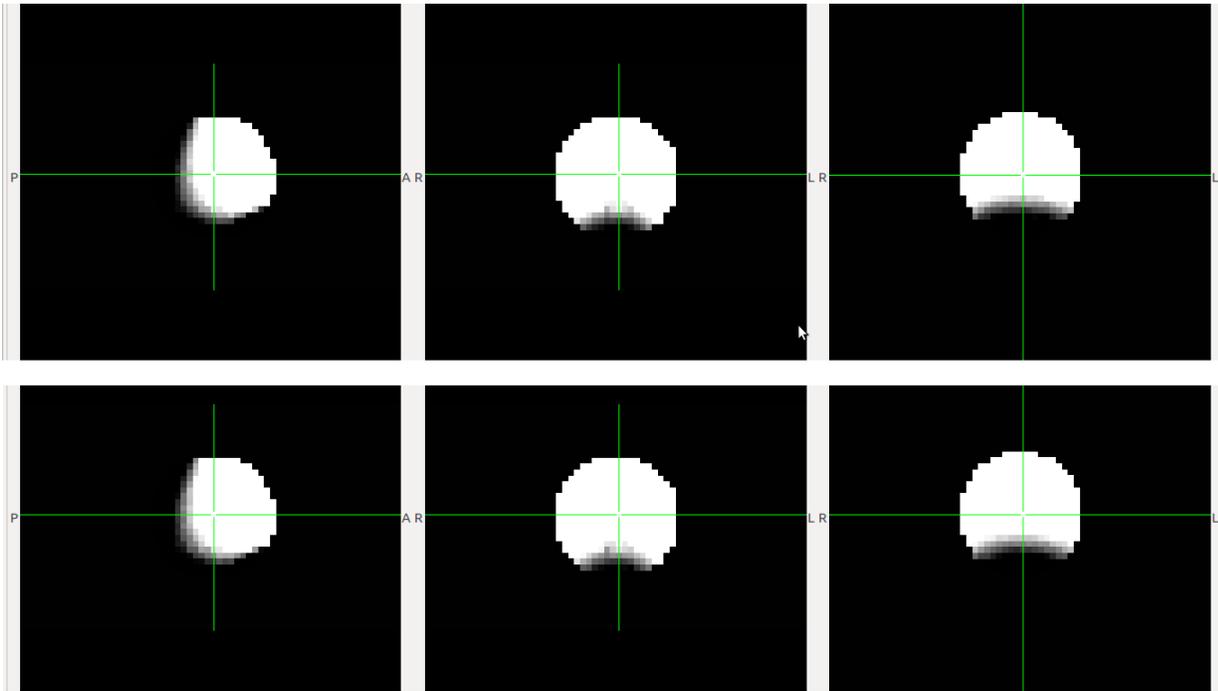


Figure 6.58: Gradient magnetic field maps of MRI phantom (masked phase images): (*top row*) without device, and (*bottom row*) with cello and bow in scanner bore. No differences are visible.

Table 6.1 MRI compatibility of optoacoustic cello.

| Description | Value |
|--|-------------|
| Maximum change in B_0 field homogeneity ($\frac{\delta B_0}{B_0}$) | 0.02 ppm |
| Larmor (resonant) frequency shift | 3.2 Hz |
| Number of voxels | 4819 voxels |
| Mean image intensity without device | 1608.9 |
| Change in mean intensity with device | 0.9 |
| Mean voxel-by-voxel percent change | 0.11% |

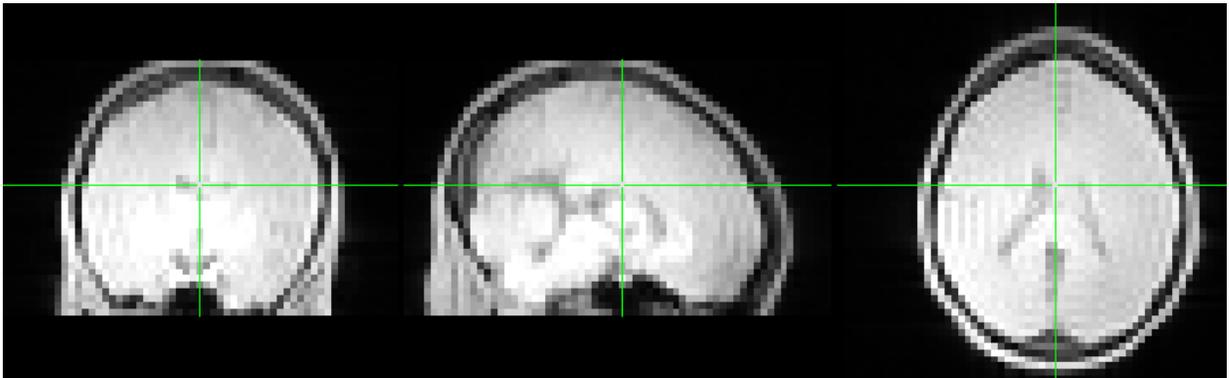


Figure 6.59: MRI field map of cellist with MRI-compatible cello (magnitude image). No image artifacts are visible.

6.5 Discussion

The first experimental setup for capturing cellist behaviour involved a pitch modification and motor compensation task, similar to that performed with vocalists by Zarate and Zatorre [293, 293], conducted in collaboration with Segado et al. [239]. The experiment involved a cellist listening to a target tone, then recreating that tone on the MRI-compatible cello with and without auditory feedback. Auditory feedback was delivered to the cellist using MRI-compatible headphones and was produced by amplifying the optoelectronic signal captured from the vibrating string using the non-contact sensorized bridge. During a portion of trials that included feedback, the auditory feedback was electronically shifted in pitch by up to a semitone in either direction, after the tone produced by the cellist had enough time to stabilize. The cellist, upon hearing the change in pitch through headphones, was then directed to compensate motorically, by shifting with the left hand so that the final perceived pitch matched the initial target tone. It was found that cellists were able to compensate for the perceived pitch shifts in a timely and accurate manner, as seen in Figures 6.45 and 6.60¹³.

This study used an array of equipment, including the MRI-compatible cello and Mr-COSBBE acquisition system (see Section 3.4.5 and Fig. 3.14). Commercial audio

¹³The disturbances seen in this signal are expected and similar to what is seen when this task is performed with an acoustic instrument. The sound produced by the cello is not a pure tone and contains transients that cause errors in the pitch-tracking algorithm. As well, some roughness in the shift is expected as this is an imperfect (i.e., human) performance.

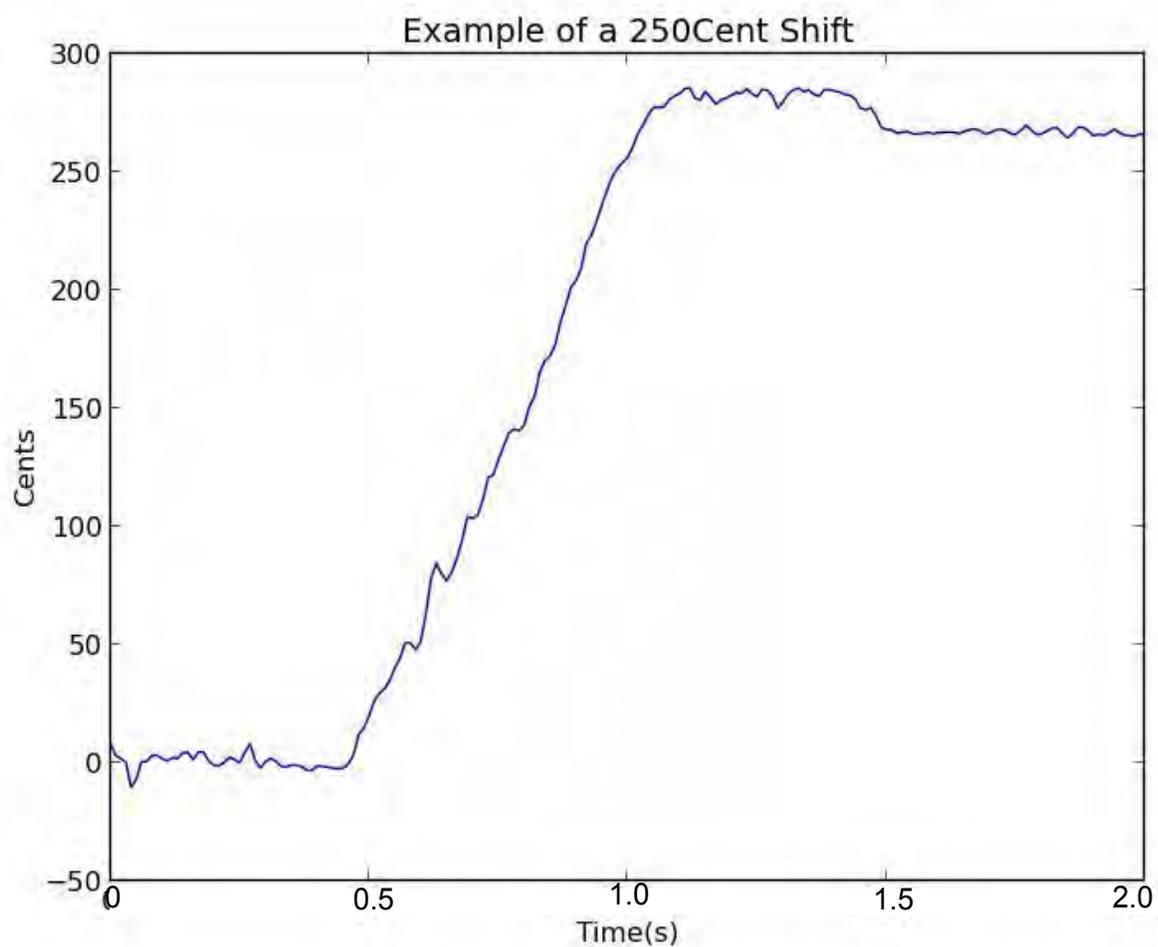


Figure 6.60: Fundamental frequency shift from D^\sharp to F^\sharp on the D-string, captured with the non-contact sensorized bridge [239].

systems, including a laptop computer, multi-channel recorder, audio mixer, and pitch shifter were integrated into the stimuli and feedback presentation system. The laptop computer managed the ordering and timing of trials, while logging the timing of MRI acquisition sequences. The multi-channel recorder was used to log the auditory stimuli, cello performance, and auditory feedback. The MR-COSBBE acquisition system was used to capture and log behavioural data, including: string displacement (non-contact sensorized bridge), finger-to-fingerboard contact position (3D-printed transparent fingerboard), bow pressure (toroidal silicone fibre optic sensors) and movement (TGG magnetic field sensors). Initial pilot testing was performed with an acquisition rate of 100–300 Hz per channel with data reported to a laptop computer over the USB-serial link, and internal *fungible-2* logging to a micro-SD card at 1 kHz per channel. Subsequent testing made use of the multi-channel USB audio mode, which allowed for a higher sampling rate and unified acquisition of all channels simultaneously.

This rather complete and flexible system, with its specialization for cello performance capture within the MRI scanner has never before been documented or demonstrated. There are, however, certain limitations, improvements to be made, as well as future testing and calibration that are discussed in this section. Future studies will continue to push the limits of what is possible with this unique instrument, and perhaps answer the question of MEG compatibility.

6.5.1 Cello Body

Improvements to the cello body were proposed after discussion with cellists and luthiers. The cello neck, having an underside cross-section composed of an angled V-shape and compound circular curve to allow for the passage of optical fibres exiting the fingerboard, should be reduced in thickness and be redesigned with a curvature to better match the characteristic shape of most acoustic cellos. For this, the aluminium mould would have to be re-designed and re-tooled, and a third composite prototype would need to be laid-up. The next prototype cello body could benefit, as well, from a top plate laid over the main body portion between the fingerboard and bridge to increase the stiffness and eliminate any decrease in the nut-bridge distance as the strings are tightened, or any out-of-plane torsional deformation. It is important to note that while the acoustic projection of the cello tone would be improved with increased body stiff-

ness, for neuroimaging studies this augmentation would be detrimental to manipulated auditory feedback tasks where the sound level of masking noise would have to be increased in kind.

6.5.2 Fingerboard

The 3D-printed fingerboard prototype was utilized for the initial neuroimaging study for a combination of reasons, including flexibility and ease of use given the timing of the experimental process, as well as limitations in manufacturing resources. Without clear, explicit, and unwavering experimental goals during the design and manufacturing process, the flexibility of having been able to manufacture a single fingerboard and later, when the experimental methodology was firmly cemented, secure the optical fibres where and only as numerous as necessary, was of great benefit. The initial material clarity and subsequent post-processing of the 3D-printed fingerboard was found to be inferior than initially anticipated, especially after quantifying the quite refined results obtained using the cast transparent fingerboard. The 3D-printed fingerboard suffered from non-uniform fibre-fingerboard interfaces due to air bubbles and fibres shifting during epoxy curing. As mentioned, the disadvantage of the cast fingerboard was the requirement that all embedded optical fibres had to be cast at once and together with the fingerboard itself. Additionally, a method was required to ensure the alignment and depth of the POF was maintained during casting and curing of the fingerboard. While the design of several templates and fixtures were discussed, an adequate solution had not been found given constraints of design time, manufacturing costs, and re-usability. Thus, a hybrid cast and 3D-printed fingerboard was envisioned. The fingerboard would incorporate a relatively low-cost 3D-printed fixture wherein the polished POF would be potted during the initial phase of assembly. The fixture, designed similarly to the 3D-printed fingerboard containing many evenly-spaced holes, would allow the fibres to be placed wherever they were needed and the depth controlled by the length of jacket stripped from the POF. With the POF located and secured properly within the fixture, the casting phase would commence by mixing and degassing the urethane resin. The urethane would then be poured into the aluminium fingerboard mould and the fixture with its fibres, having a complementary curvature, would be placed in the mould and submerged in the resin. After curing, all components will

have been potted together. This method has several advantages: (1) the casting process is relatively fast, simple, and inexpensive; (2) fibres can be placed as needed, where needed; (3) fibres stay aligned during curing; (4) several identical templates can be manufactured at once, independent of the number and location of potted POF, reducing time and costs; (5) high SNR is guaranteed due to proper fibre location and depth, obviated post-processing with a high-transparency finished fingerboard, and optimal fibre-fingerboard optical coupling.

6.5.3 Bridge

While the non-contact optical bridge pickup played an admirable role in the experimental setup, the asymmetric output was a shortcoming. A simple improvement using real-time wave-shaping—that is, using a calibrated look-up table to linearize the output—would produce a more symmetric output. The intrinsic non-monotonic behaviour, however, could be eliminated by using a photomask. A triangular photomask placed between the string and receiver lens causing a shadow, proportional in area to the displacement of the string, would linearize the output signal. Another improvement, through the use of a multi-core receiver fibre and four-quadrant photodiode was attempted, though additional testing and improved lensing of the fibre onto the photodiode is needed.

The 3D-printed bridge enabled rapid prototyping and controlled parametrization of mechanical parts. A limitation was discovered in how the multiple parts were attached, as dimensions of fitted parts needed to be adjusted to compensate for the tolerances of the actual printed part. Thus, using tapered joints (e.g., on the tubular pillars seen in Fig. 6.25) would allow for a more robust fit, where tolerances along the tapered axis are sufficiently relaxed.

6.5.4 Bow

Bow Force and Position Sensing

With two optical bow force sensors attached to the bow, one at the tip and one at the frog, as shown in 6.30, it was possible to estimate the bow-string contact position using a linear model. Further work is required to characterize and calibrate the bow force

and magnetic field sensors within the scanner.

A more comprehensive bow force calibration procedure similar to that presented [99, 164, 235] is suggested using a Polhemus electromagnetic field tracking system and a load cell. Synchronous acquisition of bow force (using a load cell), bridge string deflection sensor, the bow position (using the Polhemus spatial tracking system), and the device under test (the silicone bow force sensors) is essential. This calibration would allow the comparison and analysis of the potential mismatch between frog and tip sensors. Further analysis of the Polhemus measurements, specifically the deflection of the bow hair, would help to characterize the compliance of the bow along its length. Lastly, moving the bow against the load cell in a bowing motion, at various pressures, would improve the calibration of the bow-string position estimate, and further validate the linear model.

Further development and testing of the optical flow sensor is needed to ensure the sensor is sufficiently robust to be used within the MRI scanner. A calibration of bow displacement at differing bow-bridge distances is required, and could be performed synchronously with the above-mentioned calibration procedure.

Calibration of Faraday Sensor

Capturing small gestures using the Faraday sensor within the MRI scanner has been accomplished, and enables the comparison of bowing speed and acceleration, though the characterization of the *position* measurement would validate a pseudo-linear model for small gestures. Retrieving unique positional information within the MRI scanner from the Faraday rotator requires perhaps the most involved technical setup. The sensor output itself provides insufficient information to infer the absolute position given free movement and orientation within the scanner bore. This is due to: (1) the sensor output is ambiguous for changes in the axis of polarization beyond 90 degrees; (2) the static magnetic field is rotation-symmetric to the Z-axis; (3) unconstrained orientation of the sensor along the Z-axis results in non-unique field intensities; and (4) the static magnetic field intensity decays as a nonlinear function of displacement from the isocentre, especially in the fringe field. Other researchers have shown, however, that a unique solution for the position and orientation is possible, given the MRI gradient field control signals in addition to the sensed orthogonal magnetic field components

[231, 232].

Several techniques are proposed to improve the spatial accuracy of the system through calibration and constraining the system parameters. The simplest augmentation is to add a second set of three orthogonal sensors, one at the tip and the other at the frog, in order to model the bow as a rigid body, adding much needed constraints to the system. For instance, the interpretation of an ambiguous signal maxima from a single sensor may be due to: (a) a reversal in the direction of movement, or (b) continued movement of the sensor through a monotonic change in magnetic field causing a rotation of polarization through 90 degrees. With two sets of spatially separated sensors, a change in the direction of movement would be seen as a symmetrical signal, whereas the continuation of a gesture through a monotonically-varying magnetic field would cause either a single sensor to reach a maxima at a time, or given the rare case where both sensors undergo a phase reversal at the same time, the nonlinear rate of change of magnetic field would present an identifiable differential signal. A third set of three orthogonal sensors would improve the rigid body model further, though space constraints would make that impracticable.

In order to linearize the mapping of magnetic field onto position, a characterization of the magnetic field is required. While the MRI scanner itself can measure the static magnetic field within the imaging volume, external sensors are required to perform a field mapping of the fringe field. By performing a synchronous characterization of the magnetic field using the Faraday rotator under test and an MRI-compatible Hall effect sensor as a known reference, the results could be compared and combined with the MRI scanner's intrinsic field map. From this procedure, an empirical validation of the Verdet constant of the TGG crystals is possible. Mapping the whole field, however, is a time-consuming process.

An alternative method to characterize the TGG crystal is to use the scanner's own gradient magnetic field. First, the bow with the Faraday sensors is placed at a known distance from the isocentre (within the imaging volume) of the MRI scanner, aligned along the Z-axis. The computer-controlled gradient coils can generate a maximum 40 mT/m spatially distributed magnetic field centred around the isocentre, superimposed on the static magnetic field. Synchronously capturing the time-varying gradient field control signals from the scanner itself, along with the output intensities from the Faraday rotators, would allow for the characterization of the Faraday sensors and thus the

calculation of the Verdet constant. During an fMRI task, the scanner is programmed with a specific sequence of time-varying spatial gradient magnetic fields. Without the need of a full field map, the bow position could be disambiguated with the synchronous capture of the Faraday sensor output and gradient field control signals. Within the sensor output signal, the bandwidth of smooth gestural signals resulting from bowing should be easily separable from the *kHz*-scale gradient field switching frequency. The total intensity of the sensed magnetic field will be the sum of the individual time-varying components, with the highest correlation between sensor output and gradient control signals arising from alignment of the sensors with the respective gradient fields. The intensity and polarity of the gradient field signals captured by the sensor would provide a better indication of the position of the sensor than relying on the static magnetic field alone. A rotation of approximately three degrees would be expected from an aligned sensor with a ± 40 *mT* gradient field, with an expected SNR of 38.6 dB. Fusing the data from six magnetic field sensor channels, the position data from the optical flow sensor, and the two bow force sensors would provide high-confidence measurements of bow force and position, as well as velocity and acceleration.

The Faraday effect is non-reciprocal. That is, if a beam of light were to pass from one end of a crystal to the other and then reflected back through the way it came, the angle of rotation would add for each transit, instead of cancelling itself out. Thus a crystal of half the length could have been employed with a mirror (or better yet, a prism) to produce a more compact and less expensive sensor (i.e., requiring only one lens and polarizer, and possibly one less fibre), although, co-locating the input and output fibres at one end would pose additional alignment issues.¹⁴ To further reduce the size of the sensor, a polarizer film could be employed instead of the glass-laminate polarizer that was used herein, however the optical quality and extinction ratio¹⁵ would suffer.

¹⁴Additionally, by using a semi-reflective mirror (or polarization-maintaining beam splitter), and two detectors—one to measure the once-rotated signal, the other to measure the doubly-rotated signal, and each with an appropriately-oriented polarization analyzer—we could cancel out the effects of the light source intensity fluctuations. The single-transit beam would be rotated by θ , and the second beam sent back through the crystal would be rotated by 2θ , to allow the calculation of a differential signal. Additionally, with a beam splitter and quarter-wave plate directly after the TGG crystal, temperature compensation can be accomplished as in [173].

¹⁵The extinction ratio is the proportion of light transmitted through crossed polarizers at zero and 90 degrees.

6.6 Conclusion

Adept cello performance is a skill acquired through countless hours of practice, shaping gestures and teasing out fine details of sound by fusing both motoric and auditory systems, each influencing the other, converging to move and excite. Our nascent understanding of the neural basis of cello performance can only progress by probing human performers as they play their instrument. Functional neuroimaging technologies, including fMRI and MEG, have encouraged the study of music perception and cognition, though the behavioural components of music performance have lagged behind music listening and imagery, as hitherto, the tools were non-existent.

The advent of an MRI-compatible cello, with which a cellist may perform within the MRI environment, stirs one to wonder what series of experiments might be accomplished with this new technology. Without a technological tool such as this, it is simple to resign oneself to diverge from a certain research path that is impossible to follow without said tool and say, “we won’t study that because we don’t have the equipment. At any rate, it’s too complicated.” While, yes, the complexity of the experimental setup is high, the rewards and depth of data are enormous.

With a sensorized cello able to capture both the acoustic and gestural components of cello performance, the tantalizing possibilities for generating a performance model from neural signals emerges—one that incorporates the all-important gestures, instead of attempting to proceed directly from the brain to sound. For scientists interested in the kinematics of musical performance, the sensors presented here offer a more direct approach to measuring performance, without the need for large motion-capture systems, with their own set of constraints. Music teachers, increasingly encouraging the video recording of practice sessions, might be interested in aiding their students by offering another avenue to visualize cello performance through the use of a sensorized instrument. And lastly, this type of instrument and fibre optic performance capture lends itself, quite rightly, to musicians interested in extended techniques, electronics, augmented instruments, and mixed media performance. As continuous gestural engagement with audio processing and synthesis is in the midst of overthrowing the event-based and pre-recorded composition style, it is hoped that composers will, once again, see the light that performers can shine on the interpretation of new music.

Chapter 7

Closing Remarks

The main contributions of this work are:

- Unique MRI-compatible musical interfaces
 - The Ballagumi, a novel flexible musical interface
 - Piano keyboard with position measurement throughout key's travel
 - Sensorized optoacoustic cello and bow
- Cost-effective fibre optic sensors for a variety of measurements
 - Position, displacement, vibration
 - Force, compression, flexion
 - Magnetic field intensity
- Optoelectronic hardware platform
 - Sensor interrogation and signal acquisition
 - Real-time processing and feedback
 - Communication and logging, with simple system integration

We have shown that it is not only feasible for musicians to perform on musical instruments inside the MRI scanner, but indeed necessary if we are to continue to use

fMRI for the study of the brain during musical performance. We have also demonstrated that it is possible to capture musicians' performance gestures and provide auditory feedback to subjects using specialized interfaces and electronics. As we have not had any MRI-compatibility issues throughout our entire design process, we are confident that our approach, using only fibre optic sensors, has been a good one. We look forward to pursuing new and exciting neuroimaging experiments that would not be possible without the sensors, instruments, and electronics presented here. While we are proud of what we have accomplished, there is still testing to be done and improvements to be made.

7.1 Limitations and Future Work

As described in more detail in each of the individual chapters, there are various limitations that need to be addressed as part of future work. Overwhelmingly, a better characterization of the optical sensors is required.

Piano Key Sensors

The continuous gradient photomasks, used as part of the piano key position sensors, have been validated with a subset of keys. All of the keys, however, need to be characterized independently in order for their derivative measures (e.g., instantaneous velocity and acceleration) to be meaningful, throughout each key's travel. The issue being that in order to calibrate every single key, a characteristic key-press must be applied to each one. This may be possible with a robotic device, though it is quite impossible for a human to apply such a repeatable input to the piano keys. One method to calibrate the keys in a simple and repeatable fashion, without the continued need for an external measurement system, is through the use of invariant data that we have already collected: the key free-release trajectory, measured with the Polhemus Liberty electromagnetic motion capture system. As each white (and similarly, each black) key should follow the same trajectory as it is released from a fully-depressed state, the keys can be calibrated using data that was collected in the lab for a single white and a single black key, knowing that each type of key is sufficiently similar and the same keybed angle is maintained. The key-release position curve may be stored within the acqui-

sition system, so that at each time the piano is in need of calibration, each key may be simply depressed and released, in order to measure each key's dynamic range (as is done already), and to fit the newly-acquired curve to the characteristic key-release curve, thereby calibrating and linearizing each key's response.

Cello and Bow Sensors

While the knowledge of finger-string-fingerboard contact position is useful, further analysis and characterization of the system would provide a more detailed picture of where and how many fingers are stopping the strings. Once the fingerboard is well characterized for single and multiple points of contact, a machine learning algorithm may provide a more accurate sense of where each finger is positioned.

The non-contact bridge sensor also underwent a limited characterization, and a more thorough analysis of the sensor's accuracy and linearity are well warranted. Nonlinearities of this string displacement sensor could impart harmonics or even inharmonic partials onto the acquired cello vibration signal that is used as auditory feedback. A linearized sensor signal would greatly reduce any potential sources of distortion. A bridge pickup using a multi-core optical fibre and a multiple element photodiode has been proposed to directly measure the absolute position of the string, as opposed to its relative displacement from rest. Additionally, an improved bridge pickup would also benefit from the characterization of the low frequency signal arising from the bow force applied to the string.

The bow force sensors and the resulting estimate of bow-string contact position have been validated. The characterization of the sensors could be further improved by measuring dynamic forces and position information of bowing gestures. Using well-calibrated data from the bow force measurements and the related estimates of bow position, the Faraday sensors could be characterized inside the MRI scanner. The magnetic field sensors could then be used to provide more robust measures of bow position and velocity, whether or not the bow and string were in contact.

7.2 Musical Applications

Musicians and composers are always searching for new sounds. The barrier for instrumental musicians to adopt DMIs will always be one of accessibility: the ability to pick up an instrument and immediately start playing, and the ability to spawn and shape sounds in a physically-meaningful and musically-applicable way. It would not be convenient, for instance, to wire up and wait for several minutes for a DMI to boot and then proceed through a calibration routine—nor is it reasonable to require a multi-purpose PC or extraneous graphical user interface. For these reasons, an embedded stand-alone system is required, one that provides the portability and physicality of the electric guitar and the flexibility and micro-timbral nuance of a Hammond organ.¹

Not the electronic piano keyboard, nor the general-purpose control surfaces and sensors found in portable devices like phones and tablets, nor the motion capture systems meant for gaming: none of these will provide the specialization, affordances, nor feedback necessary to create and control sounds in a way that approaches what is possible with acoustic musical instruments. I believe that physical modelling synthesis, with its inherent organic complexities, nonlinearities, and micro-timbral variations has the ability to ground the fundamental character of a DMI as a breathing, tangible device. The substantial computational cost of solving dynamical systems has limited the use of physical models to physically-inspired, modal, and digital waveguide synthesis. From a design standpoint, the increased accessibility of high-speed parallel processors, digital signal processors, and field-programmable gate arrays should pave the way for a community-based approach to embedded, real-time, solving of dynamical systems. The control and mapping aspects of physical modelling synthesis are themselves difficult problems, but through the application of system inversion, control theory, and machine learning in association with a rich control surface or interface, it is possible to create a musical instrument capable of evoking an emotional response.²

While the Ballagumi is not yet near that point, I have somehow hit a chord that resonates with those who have interacted with it—be it the feel, the look, or when it is

¹The original Hammond organ, with its mechanical “phonic wheels” [102] produced organic sounds with micro-timbral variations and reacted as a dynamical system under nonlinear control due to “defects” that were at the heart of its sonic character and feel [101, 222].

²And not just one of frustration.

actually played. A second iteration of the prototype is definitely worthwhile, and from there, perhaps a neuroimaging study may follow.

7.3 Applications in Neuroscience

The piano and cello, as I have mentioned, are presently being used in neuroimaging studies, and more are likely to follow. While fMRI delivers fine spatial acuity, MEG provides high temporal resolution. MEG, having similar material and electromagnetic sensitivities as fMRI, will require specially-designed interfaces. The use of MEG will continue to rise, along with hybrid MEG/fMRI studies, prompting the design of MEG/fMRI-compatible interfaces. As a next step, I would like to perform MEG-compatibility testing with the piano and cello interfaces. Notwithstanding motion artifacts during musical performance, the metal-free optoacoustic cello should have no issues, though the moving brass counterweights of the keyboard might need replacement.

Other musical interfaces conducive to the cylindrical bore of the MRI scanner are the recorder (*flute-à-bec*) and slide trombone. Both instruments are commercially-available in plastic and can be augmented with sensors to provide performance gestures of breath pressure and fingering, as well as embouchure and slide position, respectively. While the recorder has discrete tone holes, breath pressure can be used to shift the pitch and vary loudness. The slide of the trombone, on the other hand, enables continuously-variable control of pitch; overblowing allows jumping between registers; while breath pressure, embouchure, and the vocal tract play a role in the modulation of trombone tone [76]. These instruments highlight the motor control strategies that are the respective extensions of the piano and cello, in terms of progressing from bimanual coordination to the use of the hands, mouth, and vocal tract, possibly bridging the gap between instrumental and vocal music performance.

The low-cost hardware, free and open tools, and ample documentation of the Arduino and Raspberry Pi have eased the entry of neuroscientists into the design of their own electronic experimental equipment. These devices are often used in such a way as to mimic computer technology of decades gone by: emulating keyboard key presses and MIDI messages in order to conserve compatibility with outdated host computer hardware and legacy software running the experimental protocol. In these cases, sadly,

the source signals that were initially rich in data are truncated, leaving experimenters without the ability to look back upon the subjects' behaviours with the data mining and machine learning tools that are employed in brain-computer interfaces. The stand-alone embedded system as I've presented it, on the other hand, has been designed with the idea that one might as well collect all of the data that is acquired, whether or not it is immediately useful to the task or applicable to the current research question, as it might prove to be important in the future. Additionally, the use of an open platform to serve as the hub or host computer system supports the idea that neuroscientists, or their research assistants, might make use of the free and open community-based tools, and invest a modicum of time and effort at the outset, to establish a software environment that is interoperable between experiments and experimenters, rather than exerting an equal (or even greater) amount of energy creating a new stimuli-presenting and feedback-generating program for each new experiment. Having worked with over a dozen neuroscientists and research assistants, in the role of formal collaborator or as an *ad-hoc* consultant, I have seen far too often (and much to my chagrin, participated in) idiosyncratic implementations that are totally redesigned for the following, or another similar, experiment. With the advent of a unified hardware platform to serve the needs of an entire lab, and with a capable and flexible software environment that is community-supported, neuroscientists will be able to spend less time setting up an experiment and more time testing subjects and teasing apart datasets. Making the hardware and software available to the research community would also encourage the reproduction and incremental refinement of behavioural experiments.

7.4 Applications in Neuro-Rehabilitation

The sensors and electronics presented here are not only applicable to musical instruments, but could equally be adapted for use in other areas of behavioural neuroscience (e.g., motor learning), in clinical intervention (e.g., image-guided surgery), or in patient monitoring (e.g., respiration rate and tidal volume). As discussed in Section 2.1.4, however, there is a real need for musical interfaces that can track patients' fine and gross motor control, before, during, and at the end of a regimen of music-supported therapy. For instance, in the study performed by Rojo et al. [223], a battery of pre- and post-therapy motor control tests included the use of an ultrasonic movement analysis

system to measure gesture “smoothness,” however, there was no measure of gesture fluidity during the music-supported therapy sessions performed on a MIDI keyboard and electronic drum set, nor was there any gesture acquisition during the fMRI motor task. The MRI-compatible keyboard presented in Section 5.3, could be used to great effect to track the improvement in fine motor control throughout the course of therapy and as part of an fMRI music performance task within the scanner. The use of the magnetic field sensor (see Section 6.3.6) in combination with the compressible force sensor within the MRI scanner would enable the quantification of gesture smoothness in reaching tasks (i.e., gross motor control), periodicity and isometric force during rhythmic tapping tasks, as well as pinch and grip force.

Most importantly, augmenting motor therapy sessions in the clinic with monitored exercises practiced at home, through the use of electronic interfaces that track progress and provide real-time feedback to the user would be of great benefit to the clinician and the user. For in-home adoption, the interfaces would have to be: relatively inexpensive in comparison to medical-grade equipment, provide autonomous or caregiver-supported setup and operation, track clinically-actionable features, and provide adaptive and engaging feedback to the user. While a stand-alone rehabilitation device has some benefits over an interface which connects to a PC or a gaming system (i.e., namely configuration variability, and the lack of complete system control and adaptability, respectively), the device would require sensors, actuators and/or neuromuscular electrical stimulators, a control system, and likely include some form of audio-visual feedback; in essence, it would take the form of a human-centric robotic system. To engage users and ensure continued in-home exercise, especially in the case of children, a creative or game-based environment, where the focus is not directed at individual movements but higher-level goals would hopefully engender a sense of fun and perhaps distract from the very real difficulties they face. Adapted musical interfaces, embedded with the aforementioned technologies, would provide not only a sense of fun to the individual, but could be used in concert with others performing on musical instruments (be they adapted or conventional), integrating motor therapy sessions with the social engagement that might otherwise be lacking.

Appendices

Appendix A

Positioning of optical fibres within urethane cello fingerboard

Normal fret spacing where l is the distance (or span as luthiers call it) from nut to bridge, d is the distance from nut to fret, and n is the semi-tone number.

$$d = l - \frac{l}{2^{n/12}}$$

Now, let's say we put all the emitters where the frets should be and then we want to put the receivers mid-way (on a log scale) between the frets (i.e. essentially quarter tone markings) we divide the scale into 24 parts instead of the previous 12:

$$d = l - \frac{l}{2^{n/24}}$$

We'll likely need one more receiver than emitter, so in this case n goes from 1 to 25. Unfortunately for four strings that brings the total of fibres to 100, times the cross-sectional area of a 1 mm fibre (which is actually 2.2 mm if we count the jacket) giving 380 mm². In actuality we may need to use three fibres for each sensitive fret location, which takes the total to 144 fibres for 12 semitones on four strings.

If we take log-spaced fibres as above, we may want to vary the spot size as a function of position so that we can sense a wider area at frets closer to the nut and a narrower area closer to the bridge. Varying the spot size using the distance between the end of the fibre and the fingerboard-finger/string interface (or fingerboard-air interface when

there is no finger and/or string in the way) is a reasonable way to go, as long as for large fret spacing, the depth of the fingerboard is actually thick enough to support the fibre at that distance. If this is not the case, additional fibres will be needed to sub-divide the distance between actual finger positions for large fret spacing.

To calculate the spot size at the interface, the numerical aperture (NA) or acceptance angle is needed. This is a property of the fibre, such that

$$NA = \sqrt{n_1^2 - n_2^2} = n_1 \sin(\theta_{max})$$

where n_1 and n_2 are the core and cladding index of refraction, respectively, and θ_{max} refers to the acceptance half-angle. Given that the plastic fibre optic cables we are using have a core index of refraction of 1.49 and a numerical aperture of 0.5, and the Smooth-On Crystal Clear urethane plastic has an index of refraction of 1.491 (which we can approximate as an index of refraction of 1.49), then the fibre core and urethane indexes of refraction are ostensibly matched. Therefore we can effectively ignore refraction due to the optical fibre-urethane plastic interface, simplifying the calculation. Thus, $\theta_{max} = 19.6^\circ$ and the diameter of the spot,

$$D_s = 2(w * \tan(\theta_{max}))$$

where w is the distance between the end of the fibre and the urethane-air interface. There may be an additive term, approximately equal to the diameter of the fibre which should be added to D_s to correct for the fact that the fibre is not a point source, however the intensity of the spot at the fringes is decreased relative to the centre. A more accurate correction would be to add the beam waist to the distance term w .

Given that light needs to be conveyed from the emitter to the receiver through diffuse reflection, there needs to be some overlap between the spots of adjacent fibres. Ideally we'd like to have the spots overlapping such that the centre-to-centre distance between fibres is equal to the radius of the spot (i.e. $D_s/2 = d_{n+1} - d_n$). However, the fingerboard is likely not thick enough to support the fibre at this depth. Therefore a loss of signal by overlapping the fibres to a lesser degree could be feasible if the receiver is sufficiently sensitive. Otherwise some other configuration of emitter sharing (perhaps between/across strings) or using three fibres for each sensitive location would

be necessary. If we use three fibres per sensitive location, we can maximize w by taking the thickness of the fingerboard at the given location and subtract the minimum depth needed for mechanical/structural support. Then the centre-to-centre spacing of the three fibres would be made equal to $D_s/2$ to balance sensitivity and size of sensitive location; increasing the distance if the signal is strong enough. This method however would use the most number of fibres per sensitive location, especially if quarter-tone resolution is required for larger spacing (i.e. lower notes/positions).

Bibliography

- [1] D. C. Abeysinghe, S. Dasgupta, H. E. Jackson, and J. T. Boyd. Novel MEMS pressure and temperature sensors fabricated on optical fibers. *Journal of Micromechanics and Microengineering*, 12:229–235, May 2002.
- [2] T. Al Abdulmohsen. Aberration in hearing one’s own voice can cause not only stuttering but also depression. *Medical Hypotheses*, 74(5):784–788, May 2010.
- [3] E. Altenmüller, J. Marco-Pallares, T. F. Münte, and S. Schneider. Neural reorganization underlies improvement in stroke-induced motor dysfunction by music-supported therapy. *Annals of the New York Academy of Sciences*, 1169(1):395–405, July 2009.
- [4] Avago Technologies. AV02-1726EN: ADNS-9500 LaserStream Gaming Sensor. Technical report, Avago Technologies, 2011. Available: <http://www.avagotech.com/docs/AV02-1726EN>.
- [5] F. Avanzini. Efficient numerical modeling of vocal fold mechanics. In *Int. Conf. Forum Acusticum*, 2002.
- [6] A. P. Baader, O. Kazennikov, and M. Wiesendanger. Coordination of bowing and fingering in violin playing. *Cognitive Brain Research*, 23(2–3):436–443, May 2005.
- [7] D. Balslev, L. O. D. Christensen, J.-H. Lee, I. Law, O. B. Paulson, and R. C. Miall. Enhanced accuracy in novel mirror drawing after repetitive transcranial magnetic stimulation-induced proprioceptive deafferentation. *J. Neurosci.*, 24(43):9698–9702, Oct. 2004.
- [8] D. Balslev, J. Cole, and R. C. Miall. Proprioceptive deafferentation slows down the processing of visual hand feedback. *Journal of vision*, 7(5):12.1–12.7, 2007. PMID: 18217852 PMCID: 2570484.
- [9] M. Bangert and E. Altenmüller. Mapping perception to action in piano practice: a longitudinal DC-EEG study. *BMC Neuroscience*, 4(1):26, 2003.

- [10] M. Bangert, T. Peschel, G. Schlaug, M. Rotte, D. Drescher, H. Hinrichs, H.-J. Heinze, and E. Altenmüller. Shared networks for auditory and motor processing in professional pianists: Evidence from fMRI conjunction. *NeuroImage*, 30(3):917–926, Apr. 2006.
- [11] E. J. Barbero. *Introduction to Composite Materials Design*. CRC Press, 2011.
- [12] S. L. Bengtsson, H. H. Ehrsson, H. Forssberg, and F. Ullén. Effector-independent voluntary timing: behavioural and neuroimaging evidence. *The European Journal of Neuroscience*, 22(12):3255–3265, Dec. 2005. PMID: 16367791.
- [13] G. Bernatzky, P. Bernatzky, H. Hesse, W. Staffen, and G. Ladurner. Stimulating music increases motor coordination in patients afflicted with morbus parkinson. *Neuroscience Letters*, 361(1-3):4–8, May 2004.
- [14] M. Bock and F. K. Wacker. MR-guided intravascular interventions: Techniques and applications. *Journal of Magnetic Resonance Imaging*, 27(2):326–338, 2008.
- [15] M. Bock, J. Sikora, R. Umathum, and W. Semmler. A fiber-optical position sensor for interventional MRI. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 231, 2003.
- [16] M. Bock, J. Sikora, R. Umathum, and W. Semmler. A fiber-optical position sensor for interventional MRI. In *11th Scientific Meeting and Exhibition of the International Society for Magnetic Resonance in Medicine*, Toronto, Canada, 2003. ISMRM.
- [17] M. Bock, R. Umathum, J. Sikora, S. Brenner, E. N. Aguor, and W. Semmler. A faraday effect position sensor for interventional magnetic resonance imaging. *Physics in Medicine and Biology*, 51(4):999–1009, Feb. 2006.
- [18] G. Bolognini, M. A. Soto, and F. Di Pasquale. Fiber-optic distributed sensor based on hybrid Raman and Brillouin scattering employing multiwavelength Fabry-Pérot lasers. *Photonics Technology Letters, IEEE*, 21(20):1523–1525, Oct.15, 2009.
- [19] T. Boonstra, H. Clairbois, A. Daffertshofer, J. Verbunt, B. van Dijk, and P. Beek. MEG-compatible force sensor. *Journal of Neuroscience Methods*, 144(2):193–196, June 2005.
- [20] B. Borek, T. K. Shajahan, J. Gabriels, A. Hodge, L. Glass, and A. Shrier. Pacemaker interactions induce reentrant wave dynamics in engineered cardiac culture. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(3):033132, Aug. 2012.

- [21] N. Bottcher, S. Gelineck, and S. Serafin. PHYSMISM: a control interface for creative exploration of physical models. In *Proceedings of the 7th international conference on New interfaces for musical expression*, 2007.
- [22] D. D. Boyden. The violin and its technique in the 18th century. *The Musical Quarterly*, 36(1):9–38, Jan. 1950. ArticleType: research-article / Full publication date: Jan., 1950 / Copyright © 1950 Oxford University Press.
- [23] R. Brown, J. Chen, A. Hollinger, V. Penhune, C. Palmer, and R. Zatorre. Behavioral and neural basis for auditory-motor interactions in music performance. In *Annual Meeting of the Cognitive Neuroscience Society*, April 2010.
- [24] R. Brown, J. Chen, A. Hollinger, V. Penhune, C. Palmer, and R. Zatorre. Neurological and behavioral basis for auditory-motor transformations in music performance. In *11th International Conference on Music Perception and Cognition*, August 2010.
- [25] R. M. Brown, J. L. Chen, A. Hollinger, V. B. Penhune, C. Palmer, and R. J. Zatorre. Repetition suppression in Auditory–Motor regions to pitch and temporal structure in music. *Journal of Cognitive Neuroscience*, 25(2):313–328, Nov. 2012.
- [26] R. L. Buckner. Event-related fMRI and the hemodynamic response. *Human Brain Mapping*, 6(5-6):373–377, 1998.
- [27] Burdet, Gassert, Gowrishankar, Chapuis, and Bleuler. fMRI compatible haptic interfaces to investigate human motor control, 2006.
- [28] H. Burton, D. McLaren, and R. Sinclair. Reading embossed capital letters: An fMRI study in blind and sighted individuals. *Human Brain Mapping*, 27(4):325–339, 2006.
- [29] F. Buttkus, V. Baur, H.-C. Jabusch, M. de la Cruz Gomez-Pellin, W. Paulus, M. A. Nitsche, and E. Altenmüller. Single-session tDCS-supported retraining does not improve fine motor control in musician’s dystonia. *Restorative neurology and neuroscience*, 29(2):85–90, 2011. PMID: 21701060.
- [30] J. H. Callicott, V. S. Mattay, A. Bertolino, K. Finn, R. Coppola, J. A. Frank, T. E. Goldberg, and D. R. Weinberger. Physiological characteristics of capacity constraints in working memory as revealed by functional MRI. *Cereb. Cortex*, 9(1): 20–26, Jan. 1999.
- [31] J. H. Callicott, M. F. Egan, V. S. Mattay, A. Bertolino, A. D. Bone, B. Verchinski, and D. R. Weinberger. Abnormal fMRI response of the dorsolateral prefrontal cortex in cognitively intact siblings of patients with schizophrenia. *Am J Psychiatry*, 160 (4):709–719, Apr. 2003.

- [32] F. Carpi, A. Khanicheh, C. Mavroidis, and D. De Rossi. MRI compatibility of silicone-made contractile dielectric elastomer actuators. *Mechatronics, IEEE/ASME Transactions on*, 13(3):370–374, June 2008.
- [33] J. H. Carr and R. B. Shepherd. *Neurological rehabilitation: optimizing motor performance*. Elsevier Health Sciences, Aug. 1998.
- [34] D. Chapuis, R. Gassert, L. Sache, E. Burdet, and H. Bleuler. Design of a simple MRI/fMRI compatible force/torque sensor. In *International Conference on Intelligent Robots and Systems*, volume 3, pages 2593–2599, 2004.
- [35] D. Chapuis, R. Gassert, E. Burdet, and H. Bleuler. Hybrid ultrasonic motor and electrorheological clutch system for MR-Compatible haptic rendering. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pages 1553–1557, 2006.
- [36] J. Chen, M. Woollacott, and S. Pologe. Accuracy and underlying mechanisms of shifting movements in cellists. *Experimental Brain Research*, 174(3):467–476, Oct. 2006.
- [37] J. Chen, M. Woollacott, S. Pologe, and G. Moore. Pitch and space maps of skilled cellists: accuracy, variability, and error correction. *Experimental Brain Research*, 188(4):493–503, July 2008.
- [38] J. Chen, M. H. Woollacott, S. Pologe, and G. P. Moore. Pitch and space maps of skilled cellists: accuracy, variability, and error correction. *Experimental Brain Research*, 188(4):493–503, Apr. 2008.
- [39] J. L. Chen, R. J. Zatorre, and V. B. Penhune. Interactions between auditory and dorsal premotor cortex during synchronization to musical rhythms. *NeuroImage*, 32(4):1771–1781, Oct. 2006.
- [40] J. L. Chen, V. B. Penhune, and R. J. Zatorre. Listening to musical rhythms recruits motor regions of the brain. *Cerebral Cortex*, 18(12):2844–2854, Dec. 2008.
- [41] J. L. Chen, V. B. Penhune, and R. J. Zatorre. Moving on time: Brain network for Auditory-Motor synchronization is modulated by rhythm complexity and musical training. *Journal of Cognitive Neuroscience*, 20(2):226–239, Feb. 2008.
- [42] M. A. Choma, C. Yang, and J. A. Izatt. Instantaneous quadrature low-coherence interferometry with 3x3 fiber-optic couplers. *Optics Letters*, 28(22):2162, 2003.
- [43] M. S. Christensen, J. Lundbye-Jensen, N. Petersen, S. S. Geertsen, O. B. Paulson, and J. B. Nielsen. Watching your foot move—an fMRI study of visuomotor interactions during foot movement. *Cereb. Cortex*, 17(8):1906–1917, Aug. 2007.

- [44] P. Cook. Singing voice synthesis: History, current work, and future directions. *Computer Music Journal*, 20(4):38–46, 1996.
- [45] S. C. Cramer, G. Nelles, J. D. Schaechter, J. D. Kaplan, and S. P. Finklestein. Computerized measurement of motor performance after stroke. *Stroke*, 28(11):2162–2168, Nov. 1997.
- [46] S. C. Cramer, R. M. Weisskoff, J. D. Schaechter, G. Nelles, M. Foley, S. P. Finklestein, and B. R. Rosen. Motor cortex activation is related to force of squeezing. *Human Brain Mapping*, 16(4):197–205, Aug. 2002. PMID: 12112762.
- [47] Cypress Semiconductor. CY8C29466, CY8C29566, CY8C29666, and CY8C29866 PSoC programmable system-on-chip data sheet. Technical report, Cypress Semiconductor Corp., 2011.
- [48] Cypress Semiconductor. PSoC 5: CY8C55 family datasheet Programmable System-on-Chip (PSoC), 2011. Document Number: 001-66235.
- [49] T. Dai, J. Liu, V. Sahgal, R. Brown, and G. Yue. Relationship between muscle output and functional MRI-measured brain activation. *Experimental Brain Research*, 140(3):290–300, Oct. 2001.
- [50] J. Dakin and B. Culshaw, editors. *Optical Fiber Sensors Volume Four: Applications, Analysis, and Future Trends*. Artech, Boston, 1988.
- [51] C. Dalzell, T. Han, and I. Ruddock. Distributed optical fibre sensing of temperature using time-correlated two-photon excited fluorescence: theoretical overview. *Applied Physics B: Lasers and Optics*, 93(2):687–692, Nov. 2008.
- [52] B. L. Danielson and C. Y. Boisrobert. Absolute optical ranging using low coherence interferometry. *Applied Optics*, 30(21):2975–2979, July 1991.
- [53] L. Danisch. Fiber optic bending and positioning sensor with selected curved light emission surfaces. United States Patent, September 1995. US5633494.
- [54] C. De Luca, S. Comani, L. Di Donato, M. Caulo, M. Bertollo, and G. Romani. A-magnetic optic-mechanical device to quantify finger kinematics for fMRI studies of bimanual coordination. *Brain Topography*, 19(3):155–160, Mar. 2007.
- [55] F. Debaere, N. Wenderoth, S. Sunaert, P. Van Hecke, and S. P. Swinnen. Cerebellar and premotor function in bimanual coordination: parametric neural responses to spatiotemporal complexity and cycling frequency. *NeuroImage*, 21(4):1416–1427, Apr. 2004.

- [56] M. Demoucron, A. Askenfelt, and R. Caussé. Measuring bow force in bowed string performance: Theory and implementation of a bow force sensor. *Acta Acustica united with Acustica*, 95(4):718–732, July 2009.
- [57] Deprag Inc. Press release: No metal: Amazing air motors made from ceramic. <http://www.deprag.de/english/aktuell/presse/2009-11a.htm>, Nov. 2009.
- [58] M. Dhamala, G. Pagnoni, K. Wiesenfeld, C. F. Zink, M. Martin, and G. S. Berns. Neural correlates of the complexity of rhythmic finger tapping. *NeuroImage*, 20(2):918–926, Oct. 2003.
- [59] E. O. F. Division. LH4001-N: Polyethylene jacketed multi-core plastic optical fiber cord. Specification Sheet SPC051208-01, Mitsubishi Rayon Co. Ltd., Tokyo, Japan, December 2005. Available: <http://i-fiberoptics.com/pdf/lh4001-n.pdf>.
- [60] J. P. Donoghue, A. Nurmikko, G. Friehs, and M. Black. Chapter 63 development of neuromotor prostheses for humans. In D. S. M. Hallett, L.H. Phillips and J. Massey, editors, *Supplements to Clinical Neurophysiology*, volume Volume 57 of *Advances in Clinical Neurophysiology Proceedings of the 27th International Congress of Clinical Neurophysiology, AAEM 50th Anniversary and 57th Annual Meeting of the ACNS Joint Meeting*, pages 592–606. Elsevier, 2004.
- [61] H. H. Ehrsson, A. Fagergren, T. Jonsson, G. Westling, R. S. Johansson, and H. Forssberg. Cortical activity in precision- versus power-grip tasks: An fMRI study. *J Neurophysiol*, 83(1):528–536, Jan. 2000.
- [62] H. H. Ehrsson, A. Fagergren, and H. Forssberg. Differential fronto-parietal activation depending on force used in a precision grip task: An fMRI study. *J Neurophysiol*, 85(6):2613–2623, June 2001.
- [63] A. Engel, M. Bangert, D. Horbank, B. S. Hijmans, K. Wilkens, P. E. Keller, and C. Keysers. Learning piano melodies in visuo-motor or audio-motor training conditions and the neural correlates of their cross-modal transfer. *NeuroImage*, 63(2):966–978, Nov. 2012.
- [64] Epoxy Technology Inc. Epo-Tek 301, 2012. Available: <http://access.epotek.com/ssdocs/datasheets/301.PDF>.
- [65] Eska Optical Fiber Division. SH 4001: Super Eska polyethylene jacketed optical fiber cord. Specification sheet DPF1122-17, Mitsubishi Rayon Co., Ltd., 2001. Available: <http://www.pofeska.com/pofeskae/download/pdf/c/SH4001.pdf>.

- [66] Eska Optical Fiber Division. Eska polymer optical fiber. Technical Datasheet 01, Mitsubishi Rayon Co., Ltd., 2010. Available: http://www.pofeska.com/pofeskae/download/pdf/technical%20data_01.pdf.
- [67] M. Faraday. XLIX. experimental researches in electricity. —Nineteenth series. *Philosophical Magazine Series 3*, 28(187):294–317, 1846.
- [68] W. Feindel. Herbert Henri Jasper (1906–1999): An appreciation. *The Canadian Journal of Neurological Sciences*, 26(3):224–229, Nov. 1999.
- [69] J. Felblinger, R. M. Müri, C. Ozdoba, G. Schroth, C. W. Hess, and C. Boesch. Recordings of eye movements for stimulus control during fMRI by means of electro-oculographic methods. *Magnetic Resonance in Medicine*, 36(3):410–414, 1996.
- [70] G. R. Fink, J. C. Marshall, P. W. Halligan, C. D. Frith, J. Driver, R. S. J. Frackowiak, and R. J. Dolan. The neural consequences of conflict between intention and the senses. *Brain*, 122(3):497–512, Mar. 1999.
- [71] R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, July 1961.
- [72] M. Flueckiger, M. Bullo, D. Chapuis, R. Gassert, and Y. Perriard. fMRI compatible haptic interface actuated with traveling wave ultrasonic motor. In *Industry Applications Conference, 2005. Fourtieth IAS Annual Meeting. Conference Record of the 2005*, volume 3, pages 2075–2082 Vol. 3, 2005.
- [73] J. R. Foster, D. A. Hall, A. Q. Summerfield, A. R. Palmer, and R. W. Bowtell. Sound-level measurements and calculations of safe noise dosage during EPI at 3 T. *Journal of Magnetic Resonance Imaging*, 12(1):157–163, 2000.
- [74] I. Fotsing-Djouwe, M. Gagné, J. Laurin, and R. Kashyap. Optical fibre musical instruments: making sense of the senseless. *Journal of Materials Science: Materials in Electronics*, 20(0):170–174, 2009.
- [75] M. D. Fox, A. Z. Snyder, J. L. Vincent, and M. E. Raichle. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, 56(1):171–184, Oct. 2007.
- [76] V. Fréour and G. Scavone. Investigation of the effect of upstream airways impedance on regeneration of lip oscillations in trombone performance. In S. F. d’Acoustique, editor, *Acoustics 2012 Nantes*, Nantes, France, Apr. 2012.

- [77] N. Friedman, D. Reinkensmeyer, and M. Bachman. A Real-Time interactive MIDI glove for domicile stroke rehabilitation. In J. A. Jacko, editor, *Human-Computer Interaction. Users and Applications*, volume 6764, pages 151–158. Springer, Berlin, Heidelberg, 2011.
- [78] C. H. Fu, G. N. Vythelingum, M. J. Brammer, S. C. Williams, E. Amaro, C. M. Andrew, L. Yaguez, N. E. van Haren, K. Matsumoto, and P. K. McGuire. An fMRI study of verbal self-monitoring: Neural correlates of auditory verbal feedback. *Cereb. Cortex*, 16(7):969–977, July 2006.
- [79] N. Gaab, C. Gaser, T. Zaehle, L. Jancke, and G. Schlaug. Functional anatomy of pitch memory—an fMRI study with sparse temporal sampling. *NeuroImage*, 19(4):1417–1426, Aug. 2003.
- [80] R. Gafsi and M. A. El-Sherif. Analysis of induced-birefringence effects on fiber bragg gratings. *Optical Fiber Technology*, 6(3):299–323, July 2000.
- [81] J. R. Gaier, Y. YoderVandenberg, S. Berkebile, H. Stueben, and F. Balagadde. The electrical and thermal conductivity of woven pristine and intercalated graphite fiber–polymer composites. *Carbon*, 41(12):2187–2193, 2003.
- [82] A. Galehdar, K. J. Nicholson, P. J. Callus, W. S. T. Rowe, S. John, C. H. Wang, and K. Ghorbani. The strong diamagnetic behaviour of unidirectional carbon fiber reinforced polymer laminates. *Journal of Applied Physics*, 112(11):113921–113921–6, Dec. 2012.
- [83] R. Gassert, R. Moser, E. Burdet, and H. Bleuler. MRI/fMRI-compatible robotic system with force feedback for interaction with human motion, 2006.
- [84] R. Gassert, A. Yamamoto, D. Chapuis, L. Dovat, H. Bleuler, and E. Burdet. Actuation methods for applications in MR environments. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 29B(4):191–209, 2006.
- [85] C. Gerloff, W. Grodd, E. Altenmüller, R. Kolb, T. Naegele, U. Klose, K. Voigt, and J. Dichgans. Coregistration of EEG and fMRI in a simple motor task. *Human Brain Mapping*, 4(3):199–209, Jan. 1996.
- [86] C. Germain. Diagnosing and making tonal improvements to stringed instruments. *VSA Papers*, 1(1):124, 2005.
- [87] N. Gershenfeld and J. Paradiso. Musical Applications of Electric Field Sensing. *Computer Music J.*, 21(2):69–89, 1997.
- [88] M. Ghamsari-Esfahani. Evaluating the role of mappings in skill acquisition with digital musical instruments. Master’s thesis, McGill University, 2013.

- [89] B. Giannesini, P. J. Cozzone, and D. Bendahan. In vivo MR investigation of skeletal muscle function in small animals. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 17(3):210–218, Dec. 2004.
- [90] D. R. Gitelman, T. B. Parrish, K. S. LaBar, and M. M. Mesulam. Real-time monitoring of eye movements using infrared video-oculography during functional magnetic resonance imaging of the frontal eye fields. *NeuroImage*, 11(1):58–65, Jan. 2000.
- [91] J. M. González de Suso, G. Bernús, J. Alonso, A. Alay, A. Capdevila, J. Gili, J. A. Prat, and C. Arús. Development and characterization of an ergometer to study the bioenergetics of the human quadriceps muscle by NMR spectroscopy inside a standard MR scanner. *Magnetic Resonance in Medicine*, 29(4):575–581, 1993.
- [92] S. Graham, W. Staines, A. Nelson, D. Plewes, and W. McIlroy. New devices to deliver somatosensory stimuli during functional MRI. *Magnetic Resonance in Medicine*, 46(3):436–442, 2001.
- [93] K. Grattan and Z. Zhang. Fiber optic fluorescence thermometry. In *Topics in Fluorescence Spectroscopy*, pages 335–376. 2002.
- [94] K. T. V. Grattan and B. T. Meggitt. *Optical fiber sensor technology*. Springer, July 1998.
- [95] K. T. V. Grattan and T. Sun. Fiber optic sensor technology: an overview. *Sensors and Actuators A: Physical*, 82(1-3):40–61, May 2000.
- [96] T. Großhauser and R. Trautner. Everyday sensor system for music instruments: possibilities and usage in daily musical instrument playing, exercising and teaching. In *Proceedings of the 5th conference on Applied electromagnetics, wireless and optical communications*, pages 123–126. World Scientific and Engineering Academy and Society (WSEAS), 2007.
- [97] C. Guastavino and B. F. G. Katz. Perceptual evaluation of multi-dimensional spatial audio reproduction. *The Journal of the Acoustical Society of America*, 116(2): 1105–1115, Aug. 2004.
- [98] C. Guastavino, B. F. Katz, J.-D. Polack, D. J. Levitin, and D. Dubois. Ecological validity of soundscape reproduction. *Acta Acustica united with Acustica*, 91(2): 333–341, Mar. 2005.
- [99] E. Gaus, J. Bonada, E. Maestre, A. Pérez, and M. Blaauw. Calibration method to measure accurate bow force for real violin performances, 2009.

- [100] F. H. Guenther. Cortical interactions underlying the production of speech sounds. *Journal of Communication Disorders*, 39(5):350–365, Sept. 2006.
- [101] L. Hammond. Electrical musical instrument. United States Patent, 1934. US1956350.
- [102] L. Hammond. Electric organ. United States Patent, May 1939. US2159505.
- [103] M. Hara, G. Matthey, A. Yamamoto, D. Chapuis, R. Gassert, H. Bleuler, and T. Higuchi. Development of a 2-DOF electrostatic haptic joystick for MRI/fMRI applications. In *Proceedings of the 2009 IEEE international conference on Robotics and Automation*, pages 4238–4243, Kobe, Japan, 2009. IEEE Press.
- [104] N. G. Harbach. *Fiber Bragg gratings in polymer optical fibers*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2008.
- [105] R. Hari and R. Salmelin. Human cortical oscillations: a neuromagnetic view through the skull. *Trends in Neurosciences*, 20(1):44–49, Jan. 1997.
- [106] G. S. Harrington, C. T. Wright, and J. H. D. III. A new vibrotactile stimulator for functional MRI. *Human Brain Mapping*, 10(3):140–145, 2000.
- [107] N. G. Hatsopoulos and J. P. Donoghue. The science of neural interface systems. *Annual review of neuroscience*, 32:249–266, 2009. PMID: 19400719 PMID: PMC2921719.
- [108] S. Haykin. *Neural Networks*. IEEE, Mar. 1999.
- [109] S. S. Haykin. *Kalman filtering and neural networks*. John Wiley and Sons, 2001.
- [110] J. Hidler, T. Hodics, B. Xu, B. Dobkin, and L. G. Cohen. MR compatible force sensing system for real-time monitoring of wrist moments during fMRI testing. *Journal of Neuroscience Methods*, 155(2):300–307, Sept. 2006.
- [111] K. O. Hill and G. Meltz. Fiber bragg grating technology fundamentals and overview. *Journal of Lightwave Technology*, 15:1263–1276, Aug. 1997.
- [112] P. C. D. Hobbs. *Building Electro-Optical Systems: Making It all Work*. John Wiley & Sons, Sept. 2011.
- [113] A. Hollinger. *Design of fMRI-compatible electronic musical interfaces*. Master's, McGill University, 2008.
- [114] A. Hollinger and M. M. Wanderley. Optoelectronic acquisition and control board for musical applications. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, May 2012.

- [115] A. Hollinger and M. M. Wanderley. MRI-compatible optically-sensed cello. In *IEEE Sensors 2013*, 2013.
- [116] A. Hollinger, C. Steele, V. Penhune, R. Zatorre, and M. Wanderley. fMRI-compatible electronic controllers. In *Proceedings of the 7th international conference on New interfaces for musical expression*, 2007.
- [117] A. Hollinger, J. Thibodeau, and M. M. Wanderley. An embedded hardware platform for fungible interfaces. In *Proceedings of the International Computer Music Conference*, pages 26–29. ICMA, May 2010.
- [118] A. J. Hornof, T. Rogers, and T. Halverson. EyeMusic: performing live music and multimedia compositions with eye movements. In *Proceedings of the 7th international conference on New interfaces for musical expression*, pages 299–300, New York, New York, 2007. ACM.
- [119] K. L. Hyde, J. Lerch, A. Norton, M. Forgeard, E. Winner, A. C. Evans, and G. Schlaug. The effects of musical training on structural brain development: A longitudinal study. *Annals of the New York Academy of Sciences*, 1169:182–186, July 2009.
- [120] K. L. Hyde, J. Lerch, A. Norton, M. Forgeard, E. Winner, A. C. Evans, and G. Schlaug. Musical training shapes structural brain development. *The Journal of Neuroscience*, 29(10):3019–3025, Mar. 2009.
- [121] H. Imamizu, T. Kuroda, S. Miyauchi, T. Yoshioka, and M. Kawato. Modular organization of internal models of tools in the human cerebellum. *Proceedings of the National Academy of Sciences of the United States of America*, 100(9):5461–5466, Apr. 2003.
- [122] D. Inaudi, A. Elamari, L. Pflug, N. Gisin, J. Breguet, and S. Vurpillot. Low-coherence deformation sensors for the monitoring of civil-engineering structures. *Sensors and Actuators A: Physical*, 44(2):125–130, Aug. 1994.
- [123] D. Jacob, M. Vallet, F. Bretenaker, A. Le Floch, and R. Le Naour. Small faraday rotation measurement with a Fabry–Pérot cavity. *Applied Physics Letters*, 66(26):3546, June 1995.
- [124] G. A. James, G. He, and Y. Liu. A full-size MRI-compatible keyboard response system. *NeuroImage*, 25(1):328–331, Mar. 2005.
- [125] S. W. James and R. P. Tatam. Optical fibre long-period grating sensors: characteristics and application. *Measurement Science and Technology*, 14(5):R49–R61, 2003.

- [126] E. Jansson. Violin frequency response – bridge mobility and bridge feet distance. *Applied Acoustics*, 65(12):1197–1205, Dec. 2004.
- [127] P. Jezard and R. S. Balaban. Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine*, 34(1):65–73, 1995.
- [128] P. Jezard and S. Clare. Sources of distortion in functional MRI data. *Human Brain Mapping*, 8(2-3):80–85, 1999.
- [129] I. Jolliffe. Principal component analysis. In *Encyclopedia of Statistics in Behavioral Science*. John Wiley & Sons, Ltd, 2005.
- [130] J. Jones and D. Keough. Auditory-motor mapping for pitch control in singers and nonsingers. *Experimental Brain Research*, 190(3):279–287, 2008.
- [131] C. Jutten and J. Karhunen. Advances in blind source separation (BSS) and independent component analysis (ICA) for nonlinear mixtures. *International Journal of Neural Systems*, 14(5):267–292, Oct. 2004. PMID: 15593377.
- [132] M. Karjalainen, V. Välimäki, and Z. Jánosy. Towards high-quality sound synthesis of the guitar and string instruments. pages 56–63, 1993.
- [133] O. Kazennikov and M. Wiesendanger. Bimanual coordination of bowing and fingering in violinists—effects of position changes and string changes. *Motor control*, 13(3):297–309, July 2009. PMID: 19799167.
- [134] A. D. Kersey, M. A. Davis, H. J. Patrick, M. Leblanc, K. P. Koo, C. G. Askins, M. A. Putnam, and E. J. Friebele. Fiber grating sensors. *Journal of Lightwave Technology*, 15:1442–1463, Aug. 1997.
- [135] A. Khanicheh, A. Muto, C. Triantafyllou, B. Weinberg, L. Astrakas, A. Tzika, and C. Mavroidis. MR compatible ERF driven hand rehabilitation device. In *Rehabilitation Robotics, 2005. ICORR 2005. 9th International Conference on*, pages 7–12, 2005.
- [136] J. Kim, G. Schiemer, and T. Narushima. Oculog: playing with eye movements. In *Proceedings of the 7th international conference on New interfaces for musical expression*, pages 50–55, New York, New York, 2007. ACM.
- [137] H. Kimmig, M. W. Greenlee, F. Huethe, and T. Mergner. MR-Eyetracker: a new method for eye movement recording in functional magnetic resonance imaging. *Experimental Brain Research*, 126(3):443–449, May 1999.
- [138] P. K. Y. Ko, M. S. Demokan, and H. Tam. Distributed temperature sensing with Erbium-doped fiber amplifiers. *Journal of Lightwave Technology*, 14:2236, Oct. 1996.

- [139] K. M. Koch, D. L. Rothman, and R. A. de Graaf. Optimization of static magnetic field homogeneity in the human and animal brain in vivo. *Progress in nuclear magnetic resonance spectroscopy*, 54(2):69–96, Feb. 2009. PMID: 20126515 PMCID: PMC2802018.
- [140] K. Kotilahti, I. Nissilä, T. Näsi, L. Lipiäinen, T. Nojonen, P. Meriläinen, M. Huotilainen, and V. Fellman. Hemodynamic responses to speech and music in newborn infants. *Human Brain Mapping*, 31(4):595–603, 2010.
- [141] D. M. Krainak, T. B. Parrish, and J. P. Dewald. A method to capture six-degrees-of-freedom mechanical measurements of isometric shoulder and elbow torques during event-related fMRI. *Journal of neuroscience methods*, 161(2):314–322, Apr. 2007. PMID: 17196255 PMCID: 1892847.
- [142] M. A. Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37(2):233–243, 1991.
- [143] J. Ku, R. Mraz, N. Baker, K. K. Zakzanis, J. H. Lee, I. Y. Kim, S. I. Kim, and S. J. Graham. A data glove with tactile feedback for fMRI of virtual reality experiments. *CyberPsychology & Behavior : the impact of the Internet, multimedia and virtual reality on behavior and society*, 6(5):497–508, Oct. 2003. PMID: 14583125.
- [144] P. R. Laird. *The baroque cello revival: an oral history*. Scarecrow Press, 2004.
- [145] F. J. Langheim, J. H. Callicott, V. S. Mattay, J. H. Duyn, and D. R. Weinberger. Cortical systems associated with covert music rehearsal. *NeuroImage*, 16(4):901–908, Aug. 2002.
- [146] LAP GmbH Laser Applikationen. LAP laser template projectors for composite production, prepregs, carbon fibre fiber layup, laser projection of templates, outlines, lines, points and crosshairs, templating. Online, 2013. Available: <http://www.lap-laser.com/general-industries/your-industry/composites/>.
- [147] B. T. Larson, A. G. Erdman, N. V. Tsekos, E. Yacoub, P. V. Tsekos, and I. G. Koutlas. Design of an MRI-compatible robotic stereotactic device for minimally invasive interventions in the breast. *Journal of Biomechanical Engineering*, 126(4):458–465, 2004.
- [148] A. Larue, C. Jutten, and S. Hosseini. Markovian source separation in post-nonlinear mixtures. In *Independent Component Analysis and Blind Signal Separation*, pages 702–709. 2004.
- [149] B. Lee, C. Y. Liu, and M. L. Apuzzo. A primer on Brain-Machine interfaces, concepts, and technology: A key element in the future of functional neurorestoration. *World Neurosurgery*, 79(3–4):457–471, Mar. 2013.

- [150] D. Lesnik and D. Donlagic. In-line, fiber-optic polarimetric twist/torsion sensor. *Optics Letters*, 38(9):1494–1496, May 2013.
- [151] T. J. Levy and J. D. Beaty. Revolutionizing prosthetics: Neuroscience framework. *Johns Hopkins APL Technical Digest*, 30(3):223–229, 2011.
- [152] D. D. Lewis. Naive (bayes) at forty: The independence assumption in information retrieval. In C. Nédellec and C. Rouveirol, editors, *Machine Learning: ECML-98*, number 1398 in Lecture Notes in Computer Science, pages 4–15. Springer Berlin Heidelberg, Jan. 1998.
- [153] P. A. Lewis, A. M. Wing, P. A. Pope, P. Praamstra, and R. C. Miall. Brain activity correlates differentially with increasing temporal complexity of rhythms during initialisation, synchronisation, and continuation phases of paced finger tapping. *Neuropsychologia*, 42(10):1301–1312, 2004.
- [154] Q. W. Li, Y. Li, X. F. Zhang, S. B. Chikkannanavar, Y. H. Zhao, A. M. Dangelewicz, L. X. Zheng, S. K. Doorn, Q. X. Jia, D. E. Peterson, P. N. Arendt, and Y. T. Zhu. Structure-dependent electrical properties of carbon nanotube fibers. *Advanced Materials*, 19(20):3358–3363, 2007.
- [155] S. Li, A. Frisoli, L. Borelli, M. Bergamasco, M. Raabe, and M. W. Greenlee. Design of a new fMRI compatible haptic interface. In *Proceedings of the World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 535–540. IEEE Computer Society, 2009.
- [156] L. C. J. Lim, B. H. Kok, and M. Y. Mau. Optical navigation utilizing speed based algorithm selection. United States Patent, Jan. 2013. US20130002609.
- [157] C. J. Limb and A. R. Braun. Neural substrates of spontaneous musical performance: An fMRI study of jazz improvisation. *PLoS ONE*, 3(2):1–9, Feb 2008.
- [158] J. Liu, L. Zhang, B. Yao, and G. Yue. Accessory hardware for neuromuscular measurements during functional MRI experiments. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 13(3):164–171, Oct. 2001.
- [159] J. Z. Liu, L. Zhang, R. W. Brown, and G. H. Yue. Reproducibility of fMRI at 1.5 T in a strictly controlled motor task. *Magnetic Resonance in Medicine*, 52(4):751–760, 2004.
- [160] Y. Liu, J. Shen, W. Kainz, S. Qian, W. Wu, and J. Chen. Computational study of external fixation devices surface heating in MRI RF environment. In *2012 IEEE International Symposium on Electromagnetic Compatibility (EMC)*, pages 612–616, 2012.

- [161] M. Lotze, G. Scheler, H. R. Tan, C. Braun, and N. Birbaumer. The musician's brain: functional imaging of amateurs and professionals during performance and imagery. *NeuroImage*, 20(3):1817–1829, November 2003.
- [162] C. W. Lu, P. G. Patil, and C. A. Chestek. Current challenges to the clinical translation of brain machine interface technology. *International review of neurobiology*, 107:137–160, 2012. PMID: 23206681.
- [163] X. Luo and D. Chung. Electromagnetic interference shielding using continuous carbon-fiber carbon-matrix and polymer-matrix composites. *Composites Part B: Engineering*, 30(3):227–231, Apr. 1999.
- [164] E. Maestre, J. Bonada, M. Blaauw, A. Pérez, and E. Guaus. Acquisition of violin instrumental gestures using a commercial emf device. In *In Proceedings of International Computer Music Conference, 2007*.
- [165] Mag Design and Engineering. fMRI-compatible products. Online, 2008. Available: <http://www.magconcept.com/>.
- [166] A. P. Mahtab Ghamsari and M. Wanderley. Combining musical tasks and improvisation in evaluating novel digital musical instruments. In *Proc. of the 10th International Symposium on Computer Music Multidisciplinary Research*, volume CMMR2013, pages 506–515, Marseille, France, October 2013.
- [167] C. Maidhof, M. Rieger, W. Prinz, and S. Koelsch. Nobody is perfect: ERP effects prior to performance errors in musicians indicate fast monitoring processes. *PLoS ONE*, 4(4):e5032, 04 2009.
- [168] C. Maidhof, N. Vavatzanidis, W. Prinz, M. Rieger, and S. Koelsch. Processing expectancy violations during music performance and perception: An ERP study. *Journal of Cognitive Neuroscience*, Aug. 2009. PMID: 19702473.
- [169] J. Malloch and M. M. Wanderley. The T-Stick: From musical interface to musical instrument. In *To appear in the Proceedings of the 2007 International Conference on New Interfaces for Musical Expression (NIME07)*, New York City, USA, 2007.
- [170] J. Malloch, S. Sinclair, and M. M. Wanderley. A network-based framework for collaborative development and performance of digital musical instruments. In R. Kronland-Martinet, S. Ystad, and K. Jensen, editors, *Computer Music Modeling and Retrieval - Sense of Sounds*, volume 4969 of LNCS, pages 401–425, Berlin / Heidelberg, 2008. Springer-Verlag.
- [171] J. Malloch, S. Sinclair, A. Hollinger, and M. Wanderley. Input devices and music interaction. In J. Solis and K. Ng, editors, *Musical Robots and Interactive Multimodal*

- Systems*, volume 74 of *Springer Tracts in Advanced Robotics*, pages 67–83. Springer Berlin / Heidelberg, 2011.
- [172] J. Malloch, S. Sinclair, and M. M. Wanderley. Libmapper: (a library for connecting things). In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '13, pages 3087–3090, New York, NY, USA, 2013. ACM.
- [173] S. V. Marchese, S. Wildermuth, O. Steiger, J. Pascal, K. Bohnert, G. Eriksson, and J. Czyzewski. Optical high voltage sensor with oil- and gas-free insulation. In *Imaging and Applied Optics Technical Papers*, OSA Technical Digest (online), page STu2F.5. Optical Society of America, June 2012.
- [174] V. S. Mattay, T. E. Goldberg, F. Fera, A. R. Hariri, A. Tessitore, M. F. Egan, B. Kolachana, J. H. Callicott, and D. R. Weinberger. Catechol o-methyltransferase val158-met genotype and individual variation in the brain response to amphetamine. *Proceedings of the National Academy of Sciences of the United States of America*, 100(10):6186–6191, May 2003.
- [175] P. K. McGuire, D. A. Silbersweig, and C. D. Frith. Functional neuroanatomy of verbal self-monitoring. *Brain*, 119(3):907–917, June 1996.
- [176] McLennan, John Ewan. *The violin music acoustics from Baroque to Romantic*. PhD, The University of New South Wales, Sydney, Australia, Aug. 2008.
- [177] I. Meister, T. Krings, H. Foltys, B. Boroojerdi, M. Müller, R. Töpfer, and A. Thron. Playing piano in the mind—an fMRI study on music imagery and performance in pianists. *Cognitive Brain Research*, 19(3):219–228, May 2004.
- [178] I. Meister, T. Krings, H. Foltys, B. Boroojerdi, M. Müller, R. Töpfer, and A. Thron. Effects of long-term practice and task complexity in musicians and nonmusicians performing simple and complex motor tasks: Implications for cortical motor organization. *Human Brain Mapping*, 25(3):345–352, Apr. 2005.
- [179] Micronor Inc. MR318 MRI-compatible non-metallic fiber optic encoder. Online, 2006. Available: http://www.micronor.com/products/files/MR318/MDS_MR318.pdf.
- [180] E. R. Miranda and M. M. Wanderley. *New Digital Instruments: Control and Interaction Beyond the Keyboard*. A-R Publications, Middleton, Wisconsin, 2006.
- [181] Mitsubishi Rayon Co., Ltd. Specificatino sheet: SH 2001 super eska polyethylene jacketed optical fiber cord. Technical report, Industrial Fiber Optics Inc., 2001. Available: <http://i-fiberoptics.com/pdf/sh2001.pdf>.

- [182] Mitsubishi Rayon Co., Ltd. Specificatino sheet: SH 4001 super eska polyethylene jacketed optical fiber cord. Technical report, Industrial Fiber Optics Inc., 2001. Available: <http://www.i-fiberoptics.com/pdf/SH4001.pdf>.
- [183] S. Monosoff and P. Walls. Violin fingering. *Early Music*, 13(1):76–79, Feb. 1985.
- [184] G. C. Morris, S. Leitman, and M. Kassianidou. SillyTone squish factory. In *Proceedings of the 2004 conference on New interfaces for musical expression*, pages 201–202, Hamamatsu, Shizuoka, Japan, 2004. National University of Singapore.
- [185] R. Mraz, J. Hong, G. Quintin, W. R. Staines, W. E. McIlroy, K. K. Zakzanis, and S. J. Graham. A platform for combining virtual reality experiments with functional magnetic resonance imaging. *Cyberpsychology & Behavior: The Impact of the Internet, Multimedia and Virtual Reality on Behavior and Society*, 6(4):359–368, Aug. 2003. PMID: 14511447.
- [186] G. Nelles, S. C. Cramer, J. D. Schaechter, J. D. Kaplan, and S. P. Finklestein. Quantitative assessment of mirror movements after stroke. *Stroke*, 29(6):1182–1187, June 1998.
- [187] J. Newmarch. OpenMAX/OpenSL. Online, 2013. Available: <http://jan.newmarch.name/LinuxSound/Sampled/OpenMAX/>.
- [188] J. M. Newton, Y. Dong, J. Hidler, P. Plummer-D’Amato, J. Marehbian, R. M. Albistegui-DuBois, R. P. Woods, and B. H. Dobkin. Reliable assessment of lower limb motor representations with fMRI: use of a novel MR compatible device for real-time monitoring of ankle, knee and hip torques. *NeuroImage*, 43(1):136–146, Oct. 2008.
- [189] P. Nordbeck, W. R. Bauer, F. Fidler, M. Warmuth, K.-H. Hiller, M. Nahrendorf, M. Maxfield, S. Wurtz, W. Geistert, J. Broscheit, P. M. Jakob, and O. Ritter. Feasibility of real-time MRI with a novel carbon catheter for interventional electrophysiology. *Circulation: Arrhythmia and Electrophysiology*, 2(3):258–267, June 2009. PMID: 19808476.
- [190] Novatech Measurements Ltd. Pinch and pull force measurement. Online, 2007. Available: <http://www.novatechloadcells.co.uk/ap2/z3031.htm>.
- [191] NVIDIA Corporation. Tegra 4 processors. Online, 2013. Available: <http://www.nvidia.com/object/tegra-4-processor.html>.
- [192] J. G. Och, G. D. Clarke, W. T. Sobol, C. W. Rosen, and S. K. Mun. Acceptance testing of magnetic resonance imaging systems : report of AAPM nuclear magnetic resonance task group no.6. *Medical physics*, 19(1):217–229, 1992.

- [193] T. Ono, S. Miyakoshi, and U. Watanabe. Acoustic characteristics of unidirectionally fiber-reinforced polyurethane foam composites for musical instrument soundboards. *Acoustical Science and Technology*, 23(3):135–142, 2002.
- [194] Y. Orlarey, D. Fober, and S. Letz. Faust: an efficient functional approach to dsp programming. *New Computational Paradigms for Computer Music*, 2009.
- [195] J. Pa and G. Hickok. A parietal-temporal sensory-motor integration area for the human vocal tract: Evidence from an fMRI study of skilled musicians. *Neuropsychologia*, 46(1):362–368, 2008.
- [196] Y.-L. Park, S. Elayaperumal, B. Daniel, E. Kaye, K. Pauly, R. Black, and M. Cutkosky. MRI-compatible haptics: Feasibility of using optical fiber Bragg grating strain-sensors to detect deflection of needles in an MRI environment. In *16th Scientific Meeting and Exhibition of the International Society for Magnetic Resonance in Medicine*, Toronto, Canada, 2008. ISMRM.
- [197] Y.-L. Park, S. Elayaperumal, S. Ryu, B. Daniel, R. J. Black, B. Moslehi, and M. R. Cutkosky. MRI-compatible haptics: Strain sensing for real-time estimation of three dimensional needle deflection in MRI environments. In *17th Scientific Meeting and Exhibition of the International Society for Magnetic Resonance in Medicine*, Honolulu, Hawaii, 2009. ISMRM.
- [198] T. R. Parker, M. Farhadiroushan, V. A. Handerek, and A. J. Rogers. Temperature and strain dependence of the power level and frequency of spontaneous Brillouin scattering in optical fibers. *Optics Letters*, 22(11):787–789, June 1997.
- [199] R. Parncutt and G. McPherson. *The Science and Psychology of Music Performance: Creative Strategies for Teaching and Learning*. Oxford University Press, Mar. 2002.
- [200] C. Parrish. *Treasury of Early Music: Masterworks of the Middle Ages, the Renaissance and the Baroque Era*. Courier Dover Publications, 1958.
- [201] V. B. Penhune, R. J. Zatorre, and A. C. Evans. Cerebellar contributions to motor timing: A PET study of auditory and visual rhythm reproduction. *Journal of Cognitive Neuroscience*, 10(6):752–765, Nov. 1998.
- [202] S. Phillips and L. Lessard. Application of natural fiber composites to musical instrument top plates. *Journal of Composite Materials*, 46(2):145–154, Jan. 2012.
- [203] E. Pinet, C. Hamel, B. Glisic, D. Inaudi, and N. Miron. Health monitoring with optical fiber sensors: from human body to civil structures. In T. Kundu, editor, *Health Monitoring of Structural and Biological Systems 2007*, volume 6532, pages 1–12, San Diego, California, USA, Apr. 2007. SPIE.

- [204] R. R. Plant and G. Turner. Millisecond precision psychological research in a world of commodity computers: New hardware, new problems? *Behavior Research Methods*, 41(3):598–614, Aug. 2009.
- [205] Polhemus. Polhemus Liberty. Online, 2014. Available: <http://polhemus.com/motion-tracking/all-trackers/liberty/>.
- [206] S. Pollens. Some misconceptions about the Baroque violin. *Performance Practice Review*, 14(1):1–13, Jan. 2012. Article 6.
- [207] P. Pope, A. M. Wing, P. Praamstra, and R. C. Miall. Force related activations in rhythmic sequence production. *NeuroImage*, 27(4):909–918, Oct. 2005.
- [208] S. D. Power, T. H. Falk, and T. Chau. Classification of prefrontal activity due to mental arithmetic and music imagery using hidden markov models and frequency domain near-infrared spectroscopy. *Journal of Neural Engineering*, 7(2):026002, Apr. 2010.
- [209] Py4A. Python-for-Android. Online, 2013. Available: <http://code.google.com/p/python-for-android/>.
- [210] J. Qiu and K. Hirao. The faraday effect in diamagnetic glasses. *Journal of Materials Research*, 13(05):1358–1362, 1998.
- [211] Qualcomm Inc. Snapdragon 800 processors. Online, 2013. Available: <http://www.qualcomm.com/snapdragon/processors/800>.
- [212] Y.-J. Rao and D. A. Jackson. Recent progress in fibre optic low-coherence interferometry. *Measurement Science and Technology*, 7:981–999, July 1996.
- [213] N. Rasamimanana and F. Bevilacqua. Effort-based analysis of bowing movements: Evidence of anticipation effects. *Journal of New Music Research*, 37(4):339–351, 2008.
- [214] Raspberry Pi Foundation. FAQs: Raspberry Pi. Online, 2013. Available: <http://www.raspberrypi.org/faqs>.
- [215] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, June 1960.
- [216] J. R. Reichenbach, S. Wurdinger, S. O. Pfliederer, and W. A. Kaiser. Comparison of artifacts produced from carbon fiber and titanium alloy needles at 1.5 T MR imaging. *Journal of Magnetic Resonance Imaging*, 11(1):69–74, 2000.

- [217] J. Reithler, H. Reithler, E. van den Boogert, R. Goebel, and H. van Mier. Resistance-based high resolution recording of predefined 2-dimensional pen trajectories in an fMRI setting. *Journal of Neuroscience Methods*, 152(1-2):10–17, Apr. 2006.
- [218] P. Renaud and M. de Mathelin. Kinematic analysis for a novel design of MRI-compatible torque sensor. In *IROS*, pages 2640–2646. IEEE, 2009.
- [219] R. Riener, T. Villgrattner, R. Kleiser, T. Nef, and S. Kollias. fMRI-compatible electromagnetic haptic interface. In *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, pages 7024–7027, 2005.
- [220] R. Riso. Strategies for providing upper extremity amputees with tactile and hand position feedback – moving closer to the bionic arm. *Technology and Health Care*, 7(6):401–409, Jan. 1999.
- [221] Robin Gareus. POSIX-C library for handling Linear/Logitudinal time code (LTC). Online, 2013. Available: <https://github.com/x42/libltc/>.
- [222] H. Robjohns. Hammond B3: Modelled electromechanical tonewheel organ. Online, July 2003. Available: <http://www.soundonsound.com/sos/jul03/articles/hammondb3.asp>.
- [223] N. Rojo, J. Amengual, M. Juncadella, F. Rubio, E. Camara, J. Marco-Pallares, S. Schneider, M. Veciana, J. Montero, B. Mohammadi, E. Altenmüller, C. Grau, T. F. Münte, and A. Rodriguez-Fornells. Music-Supported therapy induces plasticity in the sensorimotor cortex in chronic stroke: A single-case study using multimodal imaging (fMRI-TMS). *Brain Injury*, 25(7-8):787–793, July 2011.
- [224] T. D. Rossing. *The Science of String Instruments*. Springer, Dec. 2010.
- [225] I. S. Ruddock and T. P. J. Han. Continuously distributed sensing via two photon excited fluorescence in doped optical fibre. *Journal of Physics: Conference Series*, 15:83–87, Jan. 2005.
- [226] M. H. Ruiz, F. Strübing, H.-C. Jabusch, and E. Altenmüller. EEG oscillatory patterns are associated with error prediction during music performance and are altered in musician’s dystonia. *NeuroImage*, 55(4):1791–1803, Apr. 2011.
- [227] E. Samset, T. Mala, R. Ellingsen, I. Gladhaug, O. Søreide, and E. Fosse. Temperature measurement in soft tissue using a distributed fibre bragg-grating sensor system. *Minimally Invasive Therapy and Allied Technologies*, 10(2):89, 2001.

- [228] T. Savion-Lemieux and V. B. Penhune. The effect of practice pattern on the acquisition, consolidation, and transfer of visual-motor sequences. *Experimental Brain Research*, 204(2):271–281, July 2010. PMID: 20526710.
- [229] T. Savion-Lemieux, J. Bailey, and V. B. Penhune. Developmental contributions to motor sequence learning. *Experimental Brain Research*, 195(2):293–306, May 2009.
- [230] J. D. Schaechter, C. Stokes, B. D. Connell, K. Perdue, and G. Bonmassar. Finger motion sensors for fMRI motor studies. *NeuroImage*, 31(4):1549–1559, July 2006.
- [231] K. Scheffler and J. Korvink. Navigation with hall sensor device for interventional MRI. In *Proc. 11th ISMRM Meeting and Exhibition*, page 950, 2004.
- [232] J.-B. Schell, J.-B. Kammerer, L. Hebrard, E. Breton, D. Gounot, L. Cuvillon, and M. de Mathelin. Towards a hall effect magnetic tracking device for MRI. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2964–2967, 2013.
- [233] J. C. Schelleng. The bowed string and the player. *The Journal of the Acoustical Society of America*, 53(1):26–41, 1973.
- [234] S. Schneider, T. Münte, A. Rodriguez-Fornells, M. Sailer, and E. Altenmüller. Music-Supported training is more efficient than functional motor training for recovery of fine motor skills in stroke patients. *Music Perception*, 27(4):271–280, April 2010.
- [235] E. Schoonderwaldt. *Mechanics and acoustics of violin bowing: Freedom, constraints and control in performance*. PhD thesis, KTH, 2009.
- [236] E. Schoonderwaldt and M. Demoucron. Extraction of bowing parameters from violin performance combining motion capture and sensors. *The Journal of the Acoustical Society of America*, 126:2695, 2009.
- [237] G. Schröpfer, W. Elflein, M. de Labachellerie, H. Porte, and S. Ballandras. Lateral optical accelerometer micromachined in (100) silicon with remote readout based on coherence modulation. *Sensors and Actuators A: Physical*, 68(1-3):344–349, June 1998.
- [238] T. W. Schubert, A. D’Ausilio, and R. Canto. Using arduino microcontroller boards to measure response latencies. *Behavior Research Methods*, April:1–15, 2013.
- [239] M. Segado, A. Hollinger, R. J. Zatorre, and V. B. Penhune. Auditory-motor integration in cello players. In *IXth Progress in Motor Control Conference*, 2013.

- [240] R. D. Seidler, D. C. Noll, and G. Thiers. Feedforward and feedback processes in motor control. *NeuroImage*, 22(4):1775–1783, Aug. 2004.
- [241] J. Shao, G. Rong, and J. M. Lee. Learning a data-dependent kernel function for KPCA-based nonlinear process monitoring. *Chemical Engineering Research and Design*, 87(11):1471–1480, Nov. 2009.
- [242] C. Shi, J. Chen, G. Wu, X. Li, J. Zhou, and F. Ou. Stable dynamic detection scheme for magnetostrictive fiber-optic interferometric sensors. *Optics Express*, 14(12):5098–5102, June 2006.
- [243] Siemens AG Healthcare Sector. System owner manual: MR compatibility data sheet. Technical report, Siemens Corporation, 2009.
- [244] Siemens Healthcare USA. MagnetomTrio. Technical details, Siemens Corporation, 2010. Available: <http://usa.healthcare.siemens.com/magnetic-resonance-imaging/3t-mri-scanner/magnetom-trio-tim/technical-details>.
- [245] SL Laser. Composites. Online, 2013. Available: <http://sl-laser.com/index.php/en/applications/9-nicht-kategorisiert/127-composites-en>.
- [246] J. O. Smith. Physical modeling using digital waveguides. *Computer Music Journal*, 16(4):74–91, Dec. 1992.
- [247] J. O. Smith. Physical modeling synthesis update. *Computer Music Journal*, 20(2):44–56, July 1996.
- [248] J. O. Smith. *Viewpoints on the History of Digital Synthesis*. 1999.
- [249] Smooth-On Inc. Crystal Clear Series: Clear urethane casting resins. Online, 2011. Available: http://www.smooth-on.com/tb/files/CRYSTAL_CLEAR_200_TB.pdf.
- [250] W. Song, Z. H. Tan, and S. C. Tan. Optical navigation system with object detection. United States Patent, Nov. 2012. US20120274606.
- [251] H. G. Stassen, G. Johannsen, and N. Moray. Internal representation, internal model, human performance model and mental workload. *Automatica*, 26(4):811–820, July 1990.
- [252] C. J. Steele. Tides of change: Identifying the neural correlates of motor sequence learning. Master’s thesis, Concordia University, Montreal, QC, CA, November 2007.

- [253] C. J. Steele and V. B. Penhune. Specific increases within global decreases: A functional magnetic resonance imaging investigation of five days of motor sequence learning. *The Journal of Neuroscience*, 30(24):8332–8341, June 2010.
- [254] R. Stowell. *The Cambridge Companion to the Cello*. Cambridge University Press, June 1999.
- [255] H. A. Strobel. *Cello Making: Step by Step*. Henry Strobel Publisher, July 1995.
- [256] A. Stuart, J. Kalinowski, M. P. Rastatter, and K. Lynch. Effect of delayed auditory feedback on normal speakers at two speech rates. *The Journal of the Acoustical Society of America*, 111(5):2237–2241, May 2002.
- [257] C. Styliadis, C. Papadelis, E. Konstantinidis, A. A. Ioannides, and P. Bamidis. An MEG compatible system for measuring skin conductance responses. *Journal of Neuroscience Methods*, 212(1):114–123, Jan. 2013.
- [258] A. J. Suminski, J. L. Zimelman, and R. A. Scheidt. Design and validation of a MR-compatible pneumatic manipulandum. *Journal of Neuroscience Methods*, 163(2):255–266, July 2007.
- [259] M. Tada and T. Kanade. An MR-compatible optical force sensor for human function modeling. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004*, pages 129–136. 2004.
- [260] M. Tada, S. Sasaki, and T. Ogasawara. Development of an optical 2-axis force sensor usable in MRI environments. In *Sensors, 2002. Proceedings of IEEE*, volume 2, pages 984–989 vol.2, 2002.
- [261] M. Tada, K. Umeda, and T. Tokuno. High-precision MRI-compatible force sensor with parallel plate structure. *2008 2nd IEEE RAS EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 33–38, 2008.
- [262] N. Takahashi, M. Tada, J. Ueda, Y. Matsumoto, and T. Ogasawara. An optical 6-axis force sensor for brain function analysis using fMRI. In *Sensors, 2003. Proceedings of IEEE*, volume 1, pages 253–258 Vol.1, 2003.
- [263] F. Tam, N. W. Churchill, S. C. Strother, and S. J. Graham. A new tablet for writing and drawing during functional MRI. *Human Brain Mapping*, 32(9999):240–248, 2010.
- [264] F. Tam, N. W. Churchill, S. C. Strother, and S. J. Graham. A new tablet for writing and drawing during functional MRI. *Human Brain Mapping*, 32(2):240–248, 2011.

- [265] K. G. Tee, H. K. Lee, and C. F. Lum. Optical finger navigation device. United States Patent, Nov. 2012. US20120274607.
- [266] S. Thaerig, N. Behne, J. Schadow, D. Lenz, H. Scheich, A. Brechmann, and C. S. Herrmann. Sound level dependence of auditory evoked potentials: Simultaneous EEG recording and low-noise fMRI. *International Journal of Psychophysiology*, 67(3):235–241, Mar. 2008.
- [267] M. H. Thaut, G. C. McIntosh, and R. R. Rice. Rhythmic facilitation of gait training in hemiparetic stroke rehabilitation. *Journal of the Neurological Sciences*, 151(2):207–212, Oct. 1997.
- [268] J. Thibodeau and M. M. Wanderley. Trumpet augmentation and technological symbiosis. *Computer Music Journal*, 37(3):12–25, Sept. 2013.
- [269] J. L. N. Thibodeau. Trumpet augmentation: Rebirth and symbiosis of an acoustic instrument. Master's thesis, Music Technology Area, Schulich School of Music, McGill University, 2011.
- [270] G. W. Thickbroom, B. A. Phillips, I. Morris, M. L. Byrnes, and F. L. Mastaglia. Isometric force-related activity in sensorimotor cortex measured with functional MRI. *Experimental Brain Research*, 121(1):59–64, June 1998.
- [271] W. J. Tippet, J.-H. Lee, K. K. Zakzanis, S. E. Black, R. Mraz, and S. J. Graham. Visually navigating a virtual world with real-world impairments: A study of visually and spatially guided performance in individuals with mild cognitive impairments. *Journal of Clinical and Experimental Neuropsychology*, 31(4):447, 2009.
- [272] D. Tosi, M. Olivero, and G. Perrone. Low-cost fiber bragg grating vibroacoustic sensor for voice and heartbeat detection. *Applied Optics*, 47(28):5123–5129, Oct. 2008.
- [273] J. A. Tourville, K. J. Reilly, and F. H. Guenther. Neural mechanisms underlying auditory feedback control of speech. *NeuroImage*, 39(3):1429–1443, Feb. 2008.
- [274] A. Toyomura, S. Koyama, T. Miyamaoto, A. Terao, T. Omori, H. Murohashi, and S. Kuriki. Neural correlates of auditory feedback control in human. *Neuroscience*, 146(2):499–503, May 2007.
- [275] E. Udd. *Fiber optic sensors: an introduction for engineers and scientists*. Wiley, New York, 1991.
- [276] E. Upton and G. Halfacree. *Meet the Raspberry Pi*. John Wiley & Sons, July 2012.

- [277] H. van Duinen, M. Post, K. Vaartjes, H. Hoogduin, and I. Zijdwind. MR compatible strain gauge based force transducer. *Journal of Neuroscience Methods*, 164(2):247–254, Aug. 2007.
- [278] A. B. Villaverde and D. A. Donatti. Verdet constant of liquids; measurements with a pulsed magnetic field. *The Journal of Chemical Physics*, 71(10):4021, Nov. 1979.
- [279] A. B. Villaverde, D. A. Donatti, and D. G. Bozinis. Terbium gallium garnet verdet constant measurements with pulsed magnetic field. *Journal of Physics C: Solid State Physics*, 11(12):L495, June 1978.
- [280] J. Vogan, A. Wingert, J.-S. Plante, S. Dubowsky, M. Hafez, D. Kacher, and F. Jolesz. Manipulation in MRI devices using electrostrictive polymer actuators: with an application to reconfigurable imaging coils. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, volume 3, pages 2498–2504 Vol.3, 2004.
- [281] J. Wakatsuki, K. Akashi, and T. Uozumi. Changes in concentration of oxygenated hemoglobin in the prefrontal cortex while absorbed in listening to music. In *International Conference on Biometrics and Kansei Engineering, 2009. ICBAKE 2009*, pages 195–200, 2009.
- [282] J. E. Warren, R. J. Wise, and J. D. Warren. Sounds do-able: auditory-motor transformations and the posterior temporal plane. *Trends in Neurosciences*, 28(12):636–643, Dec. 2005.
- [283] N. Weiskopf, K. Mathiak, S. Bock, F. Scharnowski, R. Veit, W. Grodd, R. Goebel, and N. Birbaumer. Principles of a brain-computer interface (BCI) based on real-time functional magnetic resonance imaging (fMRI). *IEEE Transactions on Biomedical Engineering*, 51(6):966–970, June 2004.
- [284] T. Westin, K. Bierre, D. Gramenos, and M. Hinn. Advances in game accessibility from 2005 to 2010. In C. Stephanidis, editor, *Universal Access in Human-Computer Interaction. Users Diversity*, number 6766 in Lecture Notes in Computer Science, pages 400–409. Springer Berlin Heidelberg, Jan. 2011.
- [285] J. Woodhouse and P. M. Galluzzo. The bowed string as we know it today. *Acta Acustica united with Acustica*, 90(4):579–589, 2004.
- [286] L. Xu. Temporal BYY learning for state space approach, hidden markov model, and blind source separation. *IEEE Transactions on Signal Processing*, 48:2132–2144, July 2000.

- [287] K. Yamane and F. Mitsumori. Optical isometric force transducer for measurement of rat skeletal muscle contraction in the NMR spectrometer. *NMR in Biomedicine*, 10(6):292–295, 1997.
- [288] S.-S. Yoo, T. Fairney, N.-K. Chen, S.-E. Choo, L. P. Panych, H. Park, S.-Y. Lee, and F. A. Jolesz. Brain-computer interface using fMRI: spatial navigation by thoughts. *NeuroReport*, 15(10):1591–1595, July 2004.
- [289] D. Young. The hyperbow controller: Real-time dynamics measurement of violin performance. In *Proceedings of the 2002 conference on New interfaces for musical expression*, pages 1–6, 2002.
- [290] D. Young. *A methodology for investigation of bowed string performance through measurement of violin bowing technique*. PhD thesis, Citeseer, 2007.
- [291] M. Zadel and G. Scavone. Different strokes: a prototype software system for laptop performance and improvisation. In *Proceedings of the 2006 conference on New interfaces for musical expression*, pages 168–171, Paris, France, 2006. IRCAM — Centre Pompidou.
- [292] J. M. Zarate and R. J. Zatorre. Experience-dependent neural substrates involved in vocal pitch regulation during singing. *NeuroImage*, 40(4):1871–1887, May 2008.
- [293] J. M. Zarate and R. J. Zatorre. Neural substrates governing audiovocal integration for vocal pitch regulation in singing. *Annals of the New York Academy of Sciences*, 1060:404–408, 2005. The Neurosciences and Music II: From Perception to Performance.
- [294] R. J. Zatorre, A. R. Halpern, D. W. Perry, E. Meyer, and A. C. Evans. Hearing in the mind’s ear: A PET investigation of musical imagery and perception. *Journal of Cognitive Neuroscience*, 8(1):29–46, Jan. 1996.
- [295] R. J. Zatorre, J. L. Chen, and V. B. Penhune. When the brain plays music: auditory-motor interactions in music perception and production. *Nat Rev Neurosci*, 8(7):547–558, July 2007.