

UCLA

UCLA Electronic Theses and Dissertations

Title

Biomimetic Electrostatics for Submerged Oceanic Sensing, Communication, and Coordination

Permalink

<https://escholarship.org/uc/item/4mn2g3m2>

Author

Friedman, Jonathan Katzenstein

Publication Date

2012

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

**Biomimetic Electrostatics for Submerged
Oceanic Sensing, Communication, and
Coordination**

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Electrical Engineering

by

Jonathan Katzenstein Friedman

2012

© Copyright by
Jonathan Katzenstein Friedman
2012

ABSTRACT OF THE DISSERTATION

**Biomimetic Electrostatics for Submerged
Oceanic Sensing, Communication, and
Coordination**

by

Jonathan Katzenstein Friedman

Doctor of Philosophy in Electrical Engineering

University of California, Los Angeles, 2012

Professor Mani B. Srivastava, Chair

Navigation, localization, and targeting while completely submerged in the ocean are all extremely difficult due to the lack of a proficient sensor. The highly conductive nature of salt-water results in severe radio-wave attenuation precluding the use of RADAR. Naturally- occurring noise sources, high energy costs, long-wavelengths, and surface turbidity restrict the use of SONAR imaging to low- resolutions, depths, and far fields. The oceans dark, turbulent, and silty disposition impedes optical imaging.

Nature knows another better way. *Apterionotus Albifrons* is a nocturnal oceanic fish that cannot rely upon optical notions of vision to navigate, hunt, or avoid predators. Instead, it relies upon an electroreceptive capability achieved through a dense grid of electric field (Voltage) sensors arrayed along both sides of the body and concentrated around the head. It emits an electric field into the water and senses the self-induced forces down its sides. Objects in the vicinity that differ in conductivity from the background ocean environment disturb the field, redistribute the current, and hence the spatial distribution of voltages measured by the fish.

This dissertation chronicles the effort to produce an engineered sensor which

mirrors the biological phenomenon of electroreception and demonstrate its ability to visualize targets with different conductivities from the background ocean environment at very high resolution by detecting perturbations in a quasi-static electric field (electrostatics). This culminates in the first Biomimetic Electrostatic Images (BEI) and demonstrates the potential of the technology to provide significant advances in underwater scientific enterprises, military applications, as well as in medicine.

The dissertation of Jonathan Katzenstein Friedman is approved.

Danijela Cabric

Young H. Cho

Mario Gerla

William J. Kaiser

Mani B. Srivastava, Committee Chair

University of California, Los Angeles

2012

May you always use a bigger hammer...

TABLE OF CONTENTS

| | | |
|----------|---|-----------|
| I | Introduction | 1 |
| 1 | Ocean Systems | 2 |
| 1.1 | Submerged Wireless Sensor Networks | 2 |
| 1.2 | Can We Learn In-Ocean Networking From Fish? | 2 |
| 1.3 | Electrosensory Organs in Fish | 5 |
| 1.4 | Contributions | 6 |
| 2 | Physical Channels | 9 |
| 2.1 | Overview | 9 |
| 2.2 | Concentration Gradient (Chemotaxis) | 10 |
| 2.3 | Acoustics (Baroreception) | 11 |
| 2.4 | Electromagnetics (including Photoreception) | 13 |
| 2.5 | Electrostatics (Electroreception) | 17 |
| 2.6 | Conclusion | 19 |
| 3 | Related Work | 21 |
| 3.1 | Overview | 21 |
| 3.2 | Commercial Capacitive Sensing Integrated Circuits | 22 |
| 3.3 | Electrical Impedance Tomography | 25 |
| 3.4 | Electrical Capacitance Tomography | 27 |
| 3.5 | Submerged Antennas | 29 |
| 3.6 | Body Area Networks | 29 |
| 3.7 | Conclusion | 31 |

| | | |
|-----------|--|-----------|
| II | Biomimetic Electrostatics | 33 |
| 4 | Early Experiments with Oceanic Electrostatics | 34 |
| 4.1 | Overview | 34 |
| 4.2 | Towards Artificial Ampullae of Lorenzini | 34 |
| 4.3 | Initial Test Configuration | 35 |
| 4.4 | Metal Plate Transconduction | 35 |
| 4.5 | Tank Attenuation Experiments | 38 |
| 4.5.1 | SNR and BER | 38 |
| 4.5.2 | Angular Dependency | 40 |
| 4.6 | Ocean Deployment | 42 |
| 4.7 | Conclusion | 45 |
| 5 | The Folded Unipolar Model | 48 |
| 5.1 | Overview | 48 |
| 5.2 | Objectives | 49 |
| 5.3 | The Folded Unipole Antenna | 50 |
| 5.3.1 | Numerical Analysis of FUA's | 52 |
| 5.3.2 | History of the Folded Unipole Antenna | 53 |
| 5.3.3 | The Raines Model of the FUA | 53 |
| 5.3.4 | The Electromagnetic Field | 56 |
| 5.3.5 | Power and Input Impedance | 56 |
| 5.3.6 | FUA Analysis | 57 |
| 5.4 | Applicability to our Ocean Studies | 58 |
| 5.4.1 | The Impedance of the Artificial Ampullae | 59 |

| | | |
|----------|--|-----------|
| 5.4.2 | The Effect of Submergence | 60 |
| 5.4.3 | An Explanation for Prior Discrepancy | 61 |
| 5.4.4 | A Note of Caution Regarding FUA Analysis | 63 |
| 5.5 | Mixed-Mode Propagation | 65 |
| 5.5.1 | Experimental Configuration | 66 |
| 5.5.2 | Experimental Findings | 68 |
| 5.6 | Conclusion | 69 |
| 6 | Electrostatic Transceiver Design | 71 |
| 6.1 | Overview | 71 |
| 6.2 | Design Considerations | 72 |
| 6.2.1 | Material Considerations | 72 |
| 6.2.2 | Salinity and the Electric Field | 74 |
| 6.2.3 | Conductivity | 76 |
| 6.2.4 | Power | 78 |
| 6.2.5 | Electrical Noise | 80 |
| 6.3 | The Single Differential Channel Analog Signal Path | 83 |
| 6.4 | Conclusion | 87 |
| 7 | Modeling Differential Electrostatics | 88 |
| 7.1 | Overview | 88 |
| 7.1.1 | The Electric Field in Theory | 88 |
| 7.1.2 | The Electric Field in the Ocean | 90 |
| 7.2 | Modeling the Differential Disturbance Field | 97 |
| 7.2.1 | Modeling Electrostatics | 97 |

| | | |
|------------|---|------------|
| 7.2.2 | Basic Assumptions of BEI | 99 |
| 7.2.3 | BEI Expectations | 100 |
| 7.3 | Conclusion | 102 |
| III | Waveform Design | 103 |
| 8 | Waveform Shape | 104 |
| 8.1 | Overview | 104 |
| 8.2 | Waveform Classification | 104 |
| 8.2.1 | Tank Setup | 105 |
| 8.2.2 | Pulse-type EOD | 107 |
| 8.2.3 | Wave-type EOD | 109 |
| 8.3 | Conclusion | 110 |
| 9 | Waveform Foveation of a Single Dipole | 112 |
| 9.1 | Overview | 112 |
| 9.2 | Directivity | 112 |
| 9.3 | The Differential Field Pattern | 115 |
| 9.4 | Foveation Through Selection | 117 |
| 9.5 | Conclusion | 121 |
| 10 | Waveform Foveation through Array Synthesis | 122 |
| 10.1 | Overview | 122 |
| 10.2 | Drawing Parallels with RADAR | 123 |
| 10.3 | Numerical Analysis of Electromagnetic Structures | 126 |
| 10.4 | Arbitrary Array Analysis | 128 |

| | | |
|-----------|---|------------|
| 10.4.1 | From an RPD Point Cloud to an RPD Model | 129 |
| 10.4.2 | Optimizing RPD Models to Speed Analysis | 133 |
| 10.4.3 | Determining the Array RPD Model by Compositing | 137 |
| 10.5 | Array Steering and Optimization | 141 |
| 10.5.1 | Near-field Electrostatic Beamforming | 141 |
| 10.5.2 | Beam Steering | 143 |
| 10.5.3 | Array Optimization | 145 |
| 10.6 | Conclusion | 147 |
| IV | Waveform Encoding | 148 |
| 11 | Bio-inspired Informatics | 149 |
| 11.1 | Overview | 149 |
| 11.2 | Bio-inspired Informatics | 150 |
| 11.3 | Angle-of-Arrival-Assisted Relative Interferometry | 153 |
| 11.4 | Conclusion | 154 |
| 12 | Continuous-Wave AoA-Assisted Relative Interferometry | 155 |
| 12.1 | Overview | 155 |
| 12.2 | Angle-of-Arrival Assisted Radio Interferometry | 156 |
| 12.2.1 | Nomenclature | 158 |
| 12.2.2 | Transmission | 158 |
| 12.2.3 | Reception | 159 |
| 12.2.4 | RADAR Classification | 160 |
| 12.2.5 | Prior Art | 161 |

| | | |
|-----------|--|------------|
| 12.3 | Design Considerations | 162 |
| 12.3.1 | Modelling Phase and Frequency Error | 163 |
| 12.3.2 | Frequency Interleaving | 166 |
| 12.3.3 | Deployed Platforms and Doppler | 168 |
| 12.4 | Noncooperative Targets | 168 |
| 12.5 | Cooperating Targets | 172 |
| 12.6 | Software Defined Radio Approach to ARI | 174 |
| 12.6.1 | Transmitter | 175 |
| 12.6.2 | Receiver | 175 |
| 12.6.3 | Phase Recovery | 178 |
| 12.6.4 | SDR Feasibility | 179 |
| 12.7 | Implementation | 182 |
| 12.8 | Related Work | 184 |
| 12.9 | Conclusion | 186 |
| 13 | Pulse-Wave AoA-Assisted Relative Interferometry | 187 |
| 13.1 | Overview | 187 |
| 13.2 | Pulse-Wave ARI | 191 |
| 13.2.1 | Transmission | 194 |
| 13.2.2 | Reception | 194 |
| 13.2.3 | EW Considerations | 196 |
| 13.3 | Performance Expectations | 198 |
| 13.3.1 | Timing Requirements | 198 |
| 13.3.2 | Common Design Parameters | 198 |
| 13.3.3 | The Resolution-Coverage Trade-off | 199 |

| | |
|--|------------|
| 13.4 Conclusion | 200 |
| V Biomimetic Electrostatic Imaging | 202 |
| 14 3-Electrode Biomimetic Electrostatic Imager | 203 |
| 14.1 Introduction | 203 |
| 14.2 Proximity Sensing | 204 |
| 14.2.1 Resolving Self-Interference | 206 |
| 14.2.2 Results | 207 |
| 14.3 From Proximity Sensing to Imaging | 209 |
| 14.4 The Testing Environment | 210 |
| 14.5 Waveforms and Scanning | 210 |
| 14.6 Early Biomimetic Electrostatic Imaging | 212 |
| 14.7 Results | 216 |
| 14.8 Conclusion | 220 |
| 15 16-Electrode Biomimetic Electrostatic Imager | 221 |
| 15.1 Introduction | 221 |
| 15.2 The 16 Electrode Configuration | 221 |
| 15.3 Experimental Setup | 222 |
| 15.3.1 The Connections | 223 |
| 15.3.2 The Analog Signal Path | 224 |
| 15.3.3 The Electrodes | 225 |
| 15.4 Results | 225 |
| 15.5 Conclusion | 231 |

| | |
|--|------------|
| 16 17-Electrode Biomimetic Electrostatic Imager | 234 |
| 16.1 Overview | 234 |
| 16.2 The AFE17 Single Card 16 Channel Receiver | 234 |
| 16.2.1 The Control Bus | 236 |
| 16.2.2 Sample and Hold | 242 |
| 16.2.3 Bias Control | 242 |
| 16.2.4 Gain Control | 243 |
| 16.2.5 Frequency Control | 244 |
| 16.2.6 Annunciation Unit | 247 |
| 16.3 Evaluation | 248 |
| 16.3.1 Experimental Configuration | 249 |
| 16.3.2 Model Validation | 251 |
| 16.3.3 The Effect of Output Power | 254 |
| 16.3.4 Sensing Efficacy | 256 |
| 16.3.5 Detection Range Composite Findings | 259 |
| 16.3.6 Channel Sensitivities | 261 |
| 16.3.7 Improved Electrodes | 262 |
| 16.4 Design Conclusions | 265 |
| 16.5 Conclusion | 269 |
| | |
| 17 68-Electrode Biomimetic Electrostatic Imager | 270 |
| 17.1 Overview | 270 |
| 17.2 The BEI68 Platform | 270 |
| 17.3 BEI Detection and Tracking Performance | 276 |
| 17.4 Evaluation | 282 |

| | | |
|-----------|--|------------|
| 17.4.1 | Experimental Configuration | 282 |
| 17.4.2 | Tracking Moving Targets | 284 |
| 17.5 | Conclusion | 288 |
| VI | Conclusion | 289 |
| 18 | Future Applications | 290 |
| 18.1 | Overview | 290 |
| 18.2 | Real-Time In-Vivo Guidance of Surgical Tumor Resection | 291 |
| 18.3 | Submarine Collision Avoidance and Navigation | 293 |
| 19 | Synopsis | 295 |
| 19.1 | Overview | 295 |
| 19.2 | Contributions to Submerged Electromagnetics | 296 |
| 19.2.1 | The Folded Unipole Antenna | 297 |
| 19.2.2 | Ocean-Terminated Folded Unipoles | 298 |
| 19.2.3 | Summary | 301 |
| 19.3 | Contributions to Far-field Analysis | 301 |
| 19.3.1 | The Twisted Quadratic Distribution | 301 |
| 19.3.2 | Array Steering and Optimization | 302 |
| 19.4 | Contributions to RADAR | 304 |
| 19.5 | Contributions to Submerged Electrostatic Fundamentals | 305 |
| 19.6 | Contributions to Electrostatic Imaging | 308 |
| | References | 311 |

LIST OF FIGURES

| | | |
|-----|--|----|
| 1.1 | Electroreception is the ability of some fish species to sense the environment around them by either passively listening or actively detecting changes in a self-made electric field emitted around them. The presence of objects with different conductivities from the ocean redistribute the field lines casting a detectable electrical shadow on the body of the fish. | 4 |
| 1.2 | The Ampulla of Lorenzini is a fundamental component of the electric-field sensing organ in some Teleost and Chondrichthyes fish species. | 5 |
| 2.1 | The communication efficiency of a representative sample of commercial and research underwater pressure-wave communication systems. From top: [Int], [WYH06], [Lin], [Lin],[Ins06], [Lin], [Lin], [Lin], [Tri09], [Lin], [Lin], [Lin], [Aqu09]. (<i>Inset</i>) The communication system designed and implemented in this work (section 4.6) is at least two orders of magnitude more efficient. | 13 |
| 3.1 | Cypress Semiconductor’s TrueTouch touch-screen controller Application Specific Integrated Circuit (ASIC) shown here as reproduced from the company’s marketing materials. | 23 |
| 3.2 | A model of Stanford’s proposed EIT rig for earth exploration. Side view (<i>left</i>) and top view (<i>right</i>) of the electrode and environment arrangement. 48 electrodes were attached to a Berea sandstone core using conductive silver. The core is 5 cm in diameter and 10 cm tall. Reproduced from [Sta06]. | 26 |
| 4.1 | Our artificial Ampullae of Lorenzini electrode system inspired by the dimensions of an adult male Oman Shark. | 35 |

| | | |
|-----|---|----|
| 4.2 | Electrostatic transconduction demonstrated in a metallic plate. (<i>left</i>) Horizontal translation along the plate and (<i>right</i>) vertical translation above the plate. | 37 |
| 4.3 | Experiment configuration for signal loss over distance measurements. | 39 |
| 4.4 | Distance vs. signal to noise ratio of a BPSK signal. We can observe that the signal is about 0.5 dB stronger in salt water of 20ppt salinity, as compared with fresh water. There is no propagation if the water is removed from the tank, e.g., the SNR stays at the level of the receiver noise floor. The transmitted signal has a 20mV peak-to-peak voltage swing. | 40 |
| 4.5 | Variation in the SNR vs. distance is difficult to visualize, while the bit error rate (BER) clearly changes with increasing distance in fresh water and remains constant in salt water. | 41 |
| 4.6 | The radiation power density pattern of a generic ideal linear dipole antenna is isotropic about the axis of the dipole (in this case, Z), but experiences strong nulls along the collinear bearings. The plane of the antenna, as properly applied to our in-tank studies, is indicated by the shaded plane. | 42 |
| 4.7 | The angular dependence of the received signal to noise ratio given a fixed receiver and rotated transmitter. The transmitter and receiver probes are aligned at 0°. This test was performed in fresh water. . | 43 |
| 4.8 | Deployment test in Will Roger’s State Park, California. | 43 |
| 4.9 | Digital communication using GMSK at 8MHz and 250kbps data-rate at only 23mW of transmission power. Transmission with 100 bytes per packet (<i>left</i>). Transmission with 10 bytes per packet (<i>right</i>). A reduction in the packet length helps re-open the link. | 44 |

| | | |
|------|--|----|
| 4.10 | Our next generation underwater Software Defined Radio System (WURI). It is completely submersible and while powered by one multi-cell LiPoly battery pack. The computer is built around a dual core 1.6GHz Intel Atom 330 CPU using the CISC x86 architecture. | 46 |
| 4.11 | Quasi-static electric field spectrogram taken in the Pacific Ocean in Marina Del Rey, California. Only noise is visible. The injected 20 MHz wide sinusoidal sweep signal is not detectable. | 47 |
| 5.1 | <i>(left)</i> Basic schematic representation of the Folded Unipole Antenna (FUA) showing its four principal lumped elements. <i>(right)</i> The bottom of an FUA in New Jersey, USA operating at 1.4MHz and 1kW output power [Haw11]. | 50 |
| 5.2 | The Raines approach to folded unipole antenna analysis applied to the coaxial feedline connecting the Artificial Ampulla (AA) to the Universal Software Radio Peripheral (USR) when the AA are not submerged. Note the incredibly poor antenna behavior. This structure will not radiate efficiently. | 60 |
| 5.3 | The Raines approach to folded unipole antenna analysis applied to the coaxial feedline connecting the Artificial Ampulla (AA) to the Universal Software Radio Peripheral (USR) when the AA are submerged in the ocean. The ocean and AA act collectively as a capacitive top-termination. The retuned structure will radiate very efficiently at several points. | 62 |
| 5.4 | Sweep of Impedance and Input Current vs. height. This is my recreation of Fig. 2-8 from [Rai07], which shows some disagreement with the results presented here. | 64 |

| | | |
|-----|--|----|
| 5.5 | Sweep of Impedance and Input Current vs. height. This is my recreation of Fig. 2-10 from [Rai07], which shows some disagreement with the results presented here. | 65 |
| 5.6 | Experimental configurations for investigating mixed-mode propagation through water and air | 66 |
| 6.1 | The salt transfer function. Experimental (<i>blue, solid</i>) and ideal (<i>red, dashed</i>). Ideally, a straight line, the good linearity indicates good relative instrumentation accuracy. The more positive slope of the experimental trace indicates a deficit in absolute accuracy. | 75 |
| 6.2 | The electrical resistance of the water path between the transmitter electrodes as their separation distance is increased. Resistance varies almost linearly with separation and exponentially with decreasing salinity. | 76 |
| 6.3 | Effective transmit power into the water as both the distance between the transmit electrodes and salinity of the water is increased. The conductivity of the water plays the greatest role in determining output power. | 79 |
| 6.4 | An amplified electrostatic signal in the Pacific Ocean. Sampled after differential amplification across a pair of receive electrodes spaced 5cm apart. The actual signal is the 1kHz waveform riding on a 60Hz noise-form with high-frequency power-supply transients interspersed every 10ms. | 80 |
| 6.5 | The six stages of the Analog Front End (AFE) designed for our Biomimetic Imaging System: (a) platinum electrodes in the ocean, (b) bridge bias balance network, (c) instrumentation amplifier, (d) high-quality active BPF, (e) passive LPF, and (f) data acquisition – Agilent 33401A. | 82 |

| | | |
|-----|--|----|
| 6.6 | Expectation and achieved performance of the HQ-BPF. The dotted line represents the simulated high quality band-pass filter while the solid blue line shows the measured filter output. Despite lower gain, the filter achieves the desired $Q > 25$ criterion. | 85 |
| 6.7 | The receive electrodes and mounting fixture shown just prior to submergence. The electrodes are arranged into two differential channels on orthogonal axes. | 87 |
| 7.1 | Instrumentation and measurement equipment, shown on location, includes: an oscilloscope, data processing workstation, analog front ends, bipolar power supplies, and Agilent digital multimeters. . . . | 91 |
| 7.2 | A custom-built motion control platform and cantilever frame was constructed to gather positionally precise high density electric field measurements at depth in the Pacific Ocean. | 93 |
| 7.3 | The electrostatic field of a BEI transmitter electrode pair as calculated at depth in the Pacific Ocean. | 95 |
| 7.4 | The electrostatic field of a BEI transmitter electrode pair as observed at depth in the Pacific Ocean gathered during an expedition to Marina Del Rey, California. Note the strong agreement with the simulated data of figure 7.3. | 96 |
| 7.5 | A model of two sensor channels (cyan and blue) being disturbed by an induced dipole. The red X's are the the transmitting electrodes. | 97 |
| 7.6 | A detail view of the receiver channel located at the midpoint between the two transmitting electrodes. Only the center part of the response is shown. | 98 |

| | | |
|-----|---|-----|
| 7.7 | Simulated field analysis showing the transmit electrodes as black X's, the receiving electrode array as red squares, and the poles of the induced dipole as blue circles. The actual units and magnitudes are unimportant, the purpose is to illustrate the isovoltaic features to scale. | 100 |
| 8.1 | A diagrammatic view of the tank electrode configuration. | 106 |
| 8.2 | The experimental salt-water tank and supporting test, excitation, and measurement equipment. | 106 |
| 8.3 | Conductivity study of constant polarity excitation in our ocean-simulant tank. | 107 |
| 8.4 | The electrochemistry in our tank system disturbs the background measurements when excitation is unipolar. | 108 |
| 8.5 | Conductivity study of alternating polarity excitation in our ocean-simulant tank. | 110 |
| 9.1 | Electric fish have thousands of voltage sensors arrayed along their body and concentrated around their head. During prey capture, they curl their bodies to focus the electric field toward their prey. Figure derivative of [Nel11b]. | 113 |
| 9.2 | A Delrin (non-conductive) plate installed behind the plane of the transmission electrodes. The gantry, visible above the tank, is installed in its y-axis configuration (across the short axis of the tank) | 114 |
| 9.3 | Directivity along the long dimension of the tank (<i>top</i>) and across the short dimension of the tank (<i>bottom</i>) | 114 |
| 9.4 | Signal gain across the short-axis of the tank as a result of foveation. | 115 |
| 9.5 | The Pacific ocean recording | 116 |

| | | |
|------|--|-----|
| 9.6 | The expeditionary dataset with four thresholds overlaid. From top: (1) the limit above which the SINR is clearly sufficient for detection ($> 30\text{mV}$), (2) a reasonable guess as to the limits of detectability for our current receiver hardware ($> 20\text{mV}$), (3) the expected limit of detection as the receiver hardware is improved, and (4) the absolute zero (noise floor) of our measurement equipment. | 117 |
| 9.7 | The inherently directional nature of the differential electric field pattern. The central red region indicates the area of assured signal reach, while the multi-colored transitional regions indicate areas of intermediate detection probability. The white regions are areas where the electric field has virtually no influence. | 118 |
| 9.8 | The effect of transmitter location on the differential electric field pattern | 119 |
| 9.9 | A concept biomimetic jellyfish robotic vehicle with a circular array of electrodes around its bell. The pattern direction may be chosen simply by selecting the appropriate electrode pair. | 120 |
| 10.1 | Static electric fields generated from a dipole produce curved voltaic field lines as shown in the model. Unlike isotropic radiation, traveling along the horizontal dashed line results in non-linear changes in voltage. | 123 |
| 10.2 | Propagation mechanisms for electric charges in the ocean. | 125 |
| 10.3 | The MPED algorithm predicts radiation power density in an unknown direction by selecting the closest known value in a Mercator projection. Here an example dataset is shown as a flat Mercator projection (a), a Mercator projection with relief (b), and in the traditional spherical coordinates (c). Circle data markers indicate known points. | 130 |

| | | |
|------|---|-----|
| 10.4 | Plot of the azimuthal distribution using the conventional double-integral distribution (<i>left</i>) and our twisted-quadratic distribution (<i>right</i>). The plot on the right contains 33 fewer points in total, yet looks denser. Its distribution is 320% more uniform by our D metric definition. | 134 |
| 10.5 | An arbitrary heterogenous array composed of four elements with different radiation patterns. From the left: a uniform distribution, a primary lobe radiator, a random distribution, and a fan pattern. The elements are shown using a hybrid coordinate system consisting of cartesian element origins and spherical power densities. | 139 |
| 10.6 | The FAAAST analysis procedure can consider arrays composed of arbitrary size, placement, heterogenous antenna types, and element weighting. Shown here: (a) the resulting RPD for the array taken in far-field and shown in a 3D isometric perspective, and (b) the resulting RPD in a polar azimuth perspective for clarity. | 140 |
| 10.7 | An example of the use of an array transmitter to increase effective range and minimize side lobes. The individual transmit electrodes have a relative strength and polarity of 1, 2, -1, 1, -2, -1, left-to-right, respectively. The result is an enlarged primary lobe and reduced (relative to the primary) secondary lobes. | 142 |
| 10.8 | Focusing of the field via complication. The six charges from the previous figure have been reweighted to a new relative polarity/amplitude (left-to-right) of: 0.1, 0.1, -1, 1, -0.1, -0.1. The side lobes have been reduced by approximately 50%, while the primary lobe has been reduced in amplitude by only 14.4%. | 144 |

| | | |
|------|---|-----|
| 10.9 | The difference between the reference and focused beam patterns (previous two figures). Note the strong attenuation in the side lobes while the close end of the hyperbolic null region has been filled in by additional field strength. | 145 |
| 11.1 | A spectrogram from a telling experiment conducted and reported by [TZ05] in which two Brown Ghost Knifefish (<i>Apteronotus Albifrons</i>) were placed at opposite ends of a shared tank. Reproduced from [TZ05]. | 150 |
| 11.2 | Spectrogram (<i>top</i>) and time-domain (<i>bottom</i>) of a Brown-Ghost Knifefish interacting with an artificial electrode pair. Reproduced from [TZ05]. | 152 |
| 11.3 | (<i>Top</i>) Two waves at the same frequency, no matter the phase, produce a wave at the same frequency. (<i>Bottom</i>) Two waves at different frequencies, no matter the phase, produce a wave with an envelope at the frequency difference between the two. | 153 |
| 12.1 | An Angle-of-arrival-assisted Radio Interferometry (ARI) based system consists of transmitting beacons and passive receive-only objects each of which is equipped with a steerable directional antenna. The beacons transmit unique pairs of frequencies creating a beat envelope which is not unique to any beacon. | 157 |
| 12.2 | The proposed superheterodyne down-converter and signal-processing blocks for the ARI receiver. | 159 |
| 12.3 | The principal ARI data signal is relatively unaffected by initial phase error in the carrier signal or in generating a quadrature variant of it. | 165 |

| | | |
|------|---|-----|
| 12.4 | Phase error among the generation of the signal frequencies themselves has a more direct manifestation on the final phase at f_{com} . The relative error introduced by changing the phase of f_{x_1} and f_{x_2} is shown. | 166 |
| 12.5 | Linear resolution as a function of $f_{com}/2\pi$, or 1 degree of phase. $f_{com} > 150MHz$ is required for most room-scale applications, while navigation might accept frequencies as low as a few MegaHertz. . | 167 |
| 12.6 | Stanford Research Institute's forward looking ground penetrating radar system prototype (Reproduced from [SL05]) for near real-time detection of surface emplaced or buried anti-tank mines. This prototype was an outgrowth of military interest which began with [Cor01]. | 169 |
| 12.7 | ARI makes an initial estimate using just the intersection of the Lines-Of-Bearing (LOB) and then revises it by searching in the beam intersection volume trying to fit the position estimate to the measured phase difference of arrival. | 170 |
| 12.8 | The phase difference of arrival estimates the target position to within a smaller volume and a substantially smaller angle faster than traditional AoA RADAR approaches. The squares are transmitters, the dark circle is the receiver, and the central circle is a non-cooperating target. | 173 |
| 12.9 | Signal reception from analog to digital. (a) The recieved signal at the antenna; (b) The signal after mixing; (c) The signal after low-pass filter; (d) The signal after sampling. | 176 |

| | | |
|-------|--|-----|
| 12.10 | <i>(top)</i> The proposed full-analog superheterodyne down-converter and signal-processing blocks for the ARI receiver [FCS08]. <i>(middle)</i> An SDR approach utilizing time-domain filters. <i>(bottom)</i> An improved SDR approach utilizing frequency-domain filters. | 177 |
| 12.11 | The Band-Pass Filter bank in the Time Domain. | 178 |
| 12.12 | The Band-Pass Filter bank in the frequency domain. | 179 |
| 12.13 | Linear resolution as a function of $f_{com}/2\pi$, or 1 degree of phase. Values below the horizontal limit line (<i>red</i>) are achievable with current low-cost SDR systems. | 180 |
| 12.14 | The end-to-end phase transfer function of the SDR-ARI model including up-conversion, down-conversion, and signal processing. A linear fit is plotted in the figure | 181 |
| 12.15 | The end-to-end phase transfer error with respect to a linear fit. The error comes predominantly from quantization noise during signal reconstruction. | 181 |
| 12.16 | The principal ARI data signal is relatively unaffected by initial phase error in the carrier signal or in generating a quadrature variant of it. In the time domain, phase shift is not apparent in the data signal despite significant shift in the baseband signal. | 182 |
| 12.17 | SDR-ARI is immune to carrier incoherency between transmitters and receivers and only minorly effected, $< \pm 1^\circ$, by incoherency among the beacon carriers. Consequently, time synchronization need only occur among beacons at their substantially lower data frequencies. | 183 |
| 12.18 | The Universal Software Radio Peripheral reconfigurable hardware running gnuRadio software during SDR-ARI transmit-recvie testing. | 183 |

| | | |
|-------|---|-----|
| 12.19 | The gnuRadio data flow diagram for the transmitter (left) and the receiver (right). | 184 |
| 12.20 | Operation of our gnuRadio data flow diagram at the ARI unique (left) and common offset (right) frequencies. Here, $f_{com} = 10MHz$ | 184 |
| 13.1 | PW-ARI phase contours (<i>top</i>) and time-of-flight contours (<i>bottom</i>). PW-ARI recovers the phase-difference of arrival (<i>top</i>) and the time-of-arrival (<i>bottom</i>) simultaneously making fusion possible. | 190 |
| 13.2 | The PW-ARI phase contours and time-of-flight contours overlaid as observed at the receiver. Note the grid-like coordinate system created from the fusion of these two contour sets. PW-ARI recovers the phase-difference of arrival and the time-of-arrival simultaneously making fusion possible. | 191 |
| 13.3 | A Pulse-Wave Angle-of-arrival-assisted Radio Interferometry (PW-ARI) based system consists of pairs of transmitting beacons, grouped into stations, and passive receive-only assets which make use of the transmitter timing and RADAR echo information to locate targets. The beacons transmit unique pairs of frequencies creating a beat envelope which is not unique to any beacon allowing phase comparison. | 193 |
| 13.4 | PW-ARI coverage area vs. inter-beacon distance (<i>top</i>). Average accuracy for a target 50km in range using 20MHz (<i>bottom</i>). | 197 |
| 13.5 | A PW-ARI RADAR system imaging a target at 20MHz, from 50km. | 200 |
| 14.1 | (<i>Left</i>) Output of our purpose-built analysis tool given a dipole transmitter and a linear array of receivers . The isovoltaic field lines are indicated. (<i>Right</i>) Voltage versus position along the transect indicated by the dashed line in the left panel. | 205 |

| | | |
|--------|--|-----|
| 14.2 | <i>(Left)</i> Output of our purpose-built analysis tool given a dipole transmitter and a linear array of receivers with an induced dipole (conductive object) present. The isovoltic field lines are indicated. <i>(Right)</i> Voltage versus position along the transect indicated by the dashed line in the left panel. Note the distortion. | 205 |
| 14.3 | A comparison of tank-acquired and simulated electric field measurements. The darker traces are the background scans. The lighter traces are with a dipole present. | 206 |
| 14.4 | The design of a proximity detector requires the isolation of environmental dielectric disturbances (targets) from those self-induced. Subtracting the reference electrode reveals the disturbance field resulting from the target's presence. | 207 |
| 14.5 | Preliminary findings from our object proximity/imaging sensor. | 208 |
| 14.6 | A diagrammatic, target-facing, view of the tank electrode configuration. All dimensions are in inches. | 209 |
| 14.7 | A model of the motion control platform with the relative positions of the tank, electrodes, and conductive pipe target indicated. | 211 |
| 14.8 | The experimental salt-water tank and supporting test, excitation, and measurement equipment. | 212 |
| 14.9 | Raw sensor recordings from P_L (<i>top, red</i>), P_C (<i>middle, blue</i>), P_R (<i>bottom, green</i>) with respect to the position of a conductive target. | 213 |
| 14.10A | A spatial view of the two differential channels P_{dL} (<i>upper</i>) and P_{dR} (<i>lower</i>). The use of differential measurement improves the sensitivity of measurements as it reduces the dynamic range requirements allowing higher gain in the receiver. | 214 |

| | | |
|--------|---|-----|
| 14.11A | <i>Biomimetic Electrostatic Image</i> expresses the likelihood that the observed disturbances to the voltages measured originated from an object located at the pixel affiliated spatial coordinate. The images shown here are two-dimensional expressing a top-down look at a plane parallel to the surface of the tank and submerged to a depth of about 9cm. | 215 |
| 14.12 | Estimated target location area (the image's white pixels) as a percentage of the entire scan area. The X-Y coordinates indicate the location of the pipe at the time the image was constructed and analyzed. | 216 |
| 14.13 | Various biomimetic electrostatic images from the 2-channel (3 electrode) imaging system. Poor noise isolation performance in the front-end results in a high degree of variability and limited position resolution. | 217 |
| 14.14 | Mean position estimation error of the white pixels with respect to the center of the true location of the disturbing pipe target. The X-Y coordinates indicate the location of the pipe at the time the image was constructed and analyzed. | 218 |
| 14.15 | The prior figure viewed end-on from the Y-axis such that it now expresses the relationship between position across the long-dimension of the tank (parallel to the sensor array) and the mean position error. The parabolic nature of this lateral error function is readily apparent. | 219 |
| 15.1 | The 16 electrode array submerged in the test tank. | 222 |

| | | |
|------|--|-----|
| 15.2 | The experimental configuration consisting of (from left) an instrument rack (sitting on the table), a custom 2-axis motion control platform with the water tank installed below, and a control station. (Inset) The stainless-steel pipe used as the target. | 223 |
| 15.3 | Two representative channels of the sixteen electrode (8 channel) imager recording a metal pipe target swept through a salt-water tank. | 226 |
| 15.4 | One of the differential channels from the sixteen electrode (8 channel) imager recording a metal pipe target swept through a salt-water tank. | 227 |
| 15.5 | The fit lines without their corresponding datasets. The general correlation between the left-to-right channel positions and their corresponding fit lines is apparent. | 228 |
| 15.6 | Reducing each channel to the value detected at its center results in the green (lighter) piece-wise linear curve. The actual channel center locations are plotted as the blue (darker) piece-wise linear curve. Smooth curves are then fitted to the piece-wise linear data. | 229 |
| 15.7 | Unlike the permanent dipoles of the mathematical model, the steel pipe is an induced dipole where its dipole moment is always aligned with the curved field lines. This discrepancy is the source of the channel position error, but, as a deterministic phenomenon, may be compensated for. | 230 |
| 15.8 | A top-down image of the tank in two dimensions using real-time sensor data, matched to the calibration data, and scaled by the fit ratio. The brighter areas represent areas of high probability in which we expect the pipe target to be found. The pipe's actual boundary is indicated by the yellow circle with red x center. | 231 |

| | | |
|-------|--|-----|
| 15.9 | Various top-down images of the tank in two dimensions using real-time sensor data, matched to the calibration data, and scaled by the fit ratio. The brighter areas represent areas of high probability in which we expect the pipe target to be found. The pipe’s actual boundary is indicated by the yellow circle with red x center. The imager tracks the target as it traverses the tank. | 232 |
| 15.10 | The ability of the imager to localize the target degrades as the target moves beyond the horizontal extent of the array and, therefore, exceeds the sensor’s high confidence area. | 233 |
| 16.1 | The AFE17 as built. Shown with its shielding covers removed. Each of the square areas contain a complete analog front end. In the distance, near the edge connector fingers, the large power filtering components are visible. | 235 |
| 16.2 | A simplified schematic of the AFE17 16 channel receiver card. The card’s inputs consist of power and digital control signals. The analog outputs from the 16 analog processing units are available on the card’s edge connector. | 236 |
| 16.3 | Key resistive elements in the single element receiver were replaced by digitally controllable resistors. This provides the controlling software real-time control over the bias, gain, and center frequency of the channel. | 237 |
| 16.4 | The Ragobot was designed by the author and David C. Lee to enable research in mobile, motile, wireless sensor networks. This image first appeared in [Lee06a]. | 238 |

| | | |
|-------|---|-----|
| 16.5 | The CB2 bus protocol is a simplified version of the more generic 595 series 8-bit shift register with 3-state parallel output registers protocol shown here. The CB2 requires only the SER, SRCLK, and RCLK signals, which are renamed DataIn, Clock, and Update, respectively. Excerpt from [Ins09]. | 240 |
| 16.6 | The gain of an instrumentation amplifier is inversely proportional to the value of its gain control resistor and highly non-linear. The gain control unit in the BEI platform is optimized for precise control around 10×. | 244 |
| 16.7 | The actual and entire gain control configuration space as measured on an AFE17 operating dry in the laboratory. | 245 |
| 16.8 | Fresh from the factory, an AFE17's 16 channels have a wide distribution of center-frequencies due to component tolerances. These center frequencies may be corrected to within 1 Hz using software control over the filtering hardware. | 246 |
| 16.9 | The AFE17 provides digital control over the frequency response of its contained filters. 11 possible frequency responses are shown here, distributed evenly over the 1024 possibilities. The AFE17's 10-bit control results in a frequency control resolution of just under 1 Hz. | 247 |
| 16.10 | The AFE17's frequency adjustment function validated in hardware. Note the fine frequency resolution, high quality filter function, and non-linear change in frequency with control input – all as expected. | 248 |

| | | |
|-------|--|-----|
| 16.11 | A single AFE17 installed in a test fixture with a rack of support equipment. (from top) Two power supplies provide the four voltages needed for operation, an Arbitrary Waveform Generator (AWG) provides the transmit signal and output amplification, a multi-core computer (equipped with sampling hardware) handles data acquisition, digital processing, and visualization functions. | 250 |
| 16.12 | The F68 electrode array. 68 electrodes are arranged in four rows. The middle row is for transmit (excitation) electrodes. | 251 |
| 16.13 | A reference metal pipe is mounted to a custom 2-axis gantry and raster-swept through the tank in front of an array of receive (sensing) electrodes with two interspersed transmit (excitation) electrodes. . | 252 |
| 16.14 | A reference metal pipe is raster-swept through the tank. The resulting differential sensor voltage is plotted. The X-Y plane corresponds to the position of the pipe (the sensor is stationary). A model demonstrating expectation (<i>top</i>) and data from an actual laboratory trial (#P20) (<i>bottom</i>). Note the strong agreement. . . | 253 |
| 16.15 | A range-axis (Y-Z) view of the electrostatic disturbance field caused by a metal pipe for various output power levels. The y-axis is the received RMS amplitude versus its baseline. | 255 |
| 16.16 | The dataset of figure 16.15 viewed end-on (X-Z) and sliced through at Y=-0.5 cm. The y-axis is the received RMS amplitude versus its baseline. | 256 |
| 16.17 | The dataset of figure 16.15 viewed end-on (X-Z) and sliced through at Y=-0.5 cm after compensation via a discrete-spatial derivative filter. The effect of output power on the differential disturbance field is apparent, proportional, and monotonic. The y-axis is the received RMS amplitude versus its baseline. | 257 |

| | | |
|-------|---|-----|
| 16.18 | A comparison between a scan with no target present (<i>top</i>) and the reference metal pipe target present (<i>bottom</i>). The y-axis is the detected V_{RMS} versus its baseline value, which explains the negative values. | 258 |
| 16.19 | The profile of the differential disturbance field as observed in the laboratory tank. The increase in transmit power results in an increase in output noise (<i>top</i>) as compared with a lower output power level (<i>bottom</i>). The x-axis is range in centimeters and the y-axis is V_{RMS} as in prior figures. | 259 |
| 16.20 | Output voltage (<i>black</i>) and output power (<i>blue</i>) versus the maximum detection distance observed by the F68. The point of optimality is indicated by the dashed red line. | 260 |
| 16.21 | The received voltage (<i>blue</i>) and output power (<i>green</i>) versus the output voltage into the water. The point of optimality is indicated by the dashed red line. | 261 |
| 16.22 | The channels closest to the excitation (transmit) electrodes register the strongest changes to the differential disturbance field. | 262 |
| 16.23 | A revised array was constructed making use of stainless-steel 1/4-inch round electrodes to replace the reacting steel of those used with the F68. | 263 |
| 16.24 | Output voltage (<i>black</i>) and output power (<i>blue</i>) versus the maximum detection distance observed as in figure 16.20, but with new electrodes and a newly calibrated AFE17. | 264 |
| 16.25 | The received voltage (<i>blue</i>) and output power (<i>green</i>) versus the output voltage into the water utilizing the revised electrodes. | 264 |

| | | |
|-------|--|-----|
| 16.26 | A comparison of the two designs through a measure of efficiency – how much sensor range was achieved for a given amount of output power. | 266 |
| 16.27 | The output voltage-current characteristic. The non-linearity is expected from the situational chemistry. This dataset is taken with the improved stainless-steel electrodes. | 266 |
| 16.28 | The difference in detection range when the transmit axis is just 7.1 cm (<i>top</i>) and across the full length of the fish body at 18.3 cm (<i>bottom</i>). | 268 |
| 17.1 | The BEI68 imaging platform – as built. | 271 |
| 17.2 | The BEI68 interfaces with this data acquisition sub-assembly. . . | 272 |
| 17.3 | The electrodes attach to the analog processing cards (AFE17's) via a commodity 36-pin mini-SAS connector. | 273 |
| 17.4 | The receive and control portions of the BEI68 architecture. The four AFE17's are grouped into pairs and routed to one of the two data acquisition cards in the host PC. The darkest gray elements are located on the F68 electrode array. The next lightest elements are located on the BEI68 motherboard. The lightest gray elements are located on the AFE17. The white elements are located in the host computer. | 274 |
| 17.5 | The F68 electrode array shown connected to a single AFE17. When fully wired, the F68 requires five separate connections: four (one each) to the individual AFE17 receivers and one to the TX2 transmitter card (not shown). | 275 |
| 17.6 | The time-space variability of the target's influence on a single channel of the BEI sensor array. | 276 |

| | | |
|-------|---|-----|
| 17.7 | The console data visualization algorithm is a representation of the difference between a short and long running low-pass filter. This pathway is per channel and duplicated 16 times to cover the 17 electrodes of the AFE17. | 277 |
| 17.8 | The Short-Long Sliding Window Filter evaluated over its design space in four dimensions. | 279 |
| 17.9 | Two isosurfaces from the design space of the Short-Long Sliding Window Filter. | 281 |
| 17.10 | Annotated experiment configuration. | 283 |
| 17.11 | The reference target is a metal sponge coerced into a spherical shape roughly 3.8 cm (1.5”) in diameter. Shown with a U.S. Quarter Dollar. | 284 |
| 17.12 | BEI68 console operator’s view without a target present (<i>left</i>) and with a target on direct approach (<i>right</i>). Color is auto-ranging. . . | 285 |
| 17.13 | BEI68 console operator’s waterfall view of a conductive target oscillating back and forth in front of the array. The slant of the tracks indicates the direction. | 286 |
| 17.14 | BEI68 console operator’s view of a target making a direct-on slow approach and then retreating. | 286 |
| 17.15 | BEI68 console operator’s view of a target making a direct-on approach, this time faster than in the previous figure, and then retreating. | 287 |
| 18.1 | Traditional neuroprobes require electrodes to be placed in direct contact with cells. BEI proffers an alternative less-invasive approach while also providing anatomical imaging. Obviously not to scale. . | 293 |
| 18.2 | Damage to the Los Angeles class US Navy SSN Greenville post collision with the Japanese civilian fishing boat Ehimemaru. <i>Source: US Navy</i> | 294 |

19.1 Electromagnetics is a dynamic phenomenon in which quanta of energy travel to the target. Electrostatics is a static phenomenon in which the quanta interact from their origins. 306

LIST OF TABLES

| | | |
|------|--|-----|
| 5.1 | Comparison of Raines and my models for the resonance point below 1/4 wavelength | 63 |
| 5.2 | Propagation through air and water | 67 |
| 6.1 | The active bandpass filter | 86 |
| 14.1 | Improvement in the performance of our biomimetic electrostatic imaging systems over time. | 203 |
| 16.1 | Trial Data Summary | 254 |

ACKNOWLEDGMENTS

This dissertation represents the culmination of almost eight years of my work. However, no one completes an advanced post-secondary degree alone. It is a collaborative process that takes the support of friends, family, colleagues, and mentors. I take great pride in being worthy of that support and it is, therefore, with a requisite amount of humility that I would like to acknowledge a few of the many people who have played a role in the realization of this work:

To my mother, who got her Ph.D. at 23, who has always been there for me – even as I finished mine seven years “late”.

To my brother Ari, thank you for not finishing your MD/PhD program at the age of 23. Had that happened, mom would never let me hear the end of it! Keep up the good work. The world needs you.

To my sister-in-law Tara, thank you for making my brother so very happy and for making me an uncle!

To my brother Michael who is both a Masters in Mechanical Engineering and a master of mechanical engineering. Thank you for all of the late-night, last-minute, Jonathan-stop-panicking consults. I think all those drafting lessons you provided me paid off!

To my sister Lisa, the biz wiz, may your amazing smile always shine as bright.

To all of my interns (listed below, multi-year in italics) and students (not listed) throughout the years, from whom I have learned so much about myself, engineering, and about being a teacher: (2004): Sophia Wong, David Levin, Dennis Chao (2005): Michael Petralia, Laura Corman, Eric Whittenmeir (2005): *Kris Lin, Jennifer Yu*, Jin Li, *Stanley Hsu, David Tung, Kevin Leung* (2006): *Anna Davitian, Jaclyn Saito*, Frank Kuo, Nikki Chukah-Obah, Keith Mayoral, Patrick

Daly (2007): *Newton Truong*, Jon Lau (2008): Chenni Qian (2009): Yan Wang (2010): *Tanner Miller*

To all of the involved and impassioned faculty at the Georgia Institute of Technology (Georgia Tech) who instilled in me an infatuation with mixed-signal circuit design, including: Dr. John Peatman, Dr. Ron Schaefer, and Dr. John Bordelon.

To Lynn McCroskey, Jim Cawthon, Minton O’neal, Doug Anglin, my friends and mentors from IMAX, Inc. and Sonics Associates, Inc., whose tutelage, opportunity, and guidance made me the engineer I am today. I am forever in their debt.

To my friends Chris Collier, Jerod Stoehr, Rick McClain, Duane Benson, and Michelle McGillivray, of MEC Northwest Manufacturing and Nolan Johnson and Nancy Viter of Sunstone Circuits with whom I have worked for years and who have given me a broad and deep understanding of electronic manufacturing processes over the many scales of their business. They have always gone the extra mile to ensure that my PCB designs are built well and delivered on-time – even when it’s my fault for the delay – and I can not thank them and their staff enough.

To Jessica Horvath who, despite being from the UCLA english department, was quite adept as extra “lab hands” and for her editorial assistance that manifests in various parts of this work.

To Deena Columbia, Mandy Smith, Brenda Trujillo, and Michelle Welch, of the UCLA electrical engineering department’s graduate affairs office, who have routinely gone above and beyond the call of duty on my behalf and have worked tirelessly over many years to facilitate my matriculation. Although they remain unaware, I secretly harbour a supreme love of paperwork, deadlines, and institutional bureaucracy.

To my many friends, co-authors, and colleagues at the Networked Embedded Systems Laboratory, Thomas Schmid, Zainul Charbiwala, Rahul Balani, Henry

Herman, Newton Truong, Dustin Torres, David Lee, Chenni Qian, Younghun Kim, Heemin Park, Ram Kumar, Saurabh Ganeriwal, Simon Han, Anna Davitian, and several others whom I have surely failed to mention, who have enriched my graduate experience in ways I may never fully appreciate.

To Dr. Young H. Cho (of my committee) and Andrew Goodney for facilitating access to the Information Sciences Institute's Marina Del Rey facility where most of the in-ocean studies were conducted.

To Dr. Robert Brizolara of the United States Office of Naval Research, Dr. John Walrod of SAIC, and professor Dr. Shashank Priya of the Virginia Institute of Technology (Virginia Tech) whose personal and institutional support proved vital in the direction and completion of this dissertation. The MURI program in general, and their specific directives in particular – to simultaneously aim high and risk failure while seeking inspiration from biology, have had far-reaching and long-lasting consequences on my research interests. I am truly grateful.

And finally, to my advisor and mentor, Professor Dr. Mani Srivastava, for his guidance on the project, this dissertation, and my entire graduate school experience as a whole. His extensive knowledge on a wide range of topics is inspiring and has brought together all aspects of embedded computing and engineering for me. I would also like to thank Professors Dr. William Kaiser, Dr. Young Cho, and Dr. Danijella Cabric for their similar role and Professor Dr. Mario Gerla, from the computer science department, for serving on my thesis committee.

This material was supported in part by the U.S. Office of Naval Research under MURI Award CR-19097-430345 and the UCLA Center for Embedded Networked Sensing. I am deeply grateful for their support. Any opinions, findings, conclusions, or recommendations expressed in this material are mine and do not necessarily reflect the views of these agencies – even if they should!

VITA

- 2000 Director of Information Technologies. Oxmoor Corporation, LLC, Birmingham, Alabama.
- 2001 Founder. HalcyonIT, LLC, Atlanta, Georgia.
- 2002 Associate Professor (equivalent). Chernigov State Institute for Economics and Management, Chernigov, Ukraine.
- 2003 Bachelors of Science in Computer Engineering with High Honors. Georgia Institute of Technology (Georgia Tech).
- 2003 Graduate Student Researcher. Networked and Embedded Systems Laboratory, University of California, Los Angeles.
- 2005 Award Winner. ACM-IEEE International Symposium on Low Power Electronic Design Student Design Contest for “Helimote: Enabling Self-Sustaining Wireless Sensor Networks Through Solar Energy Harvesting.”
- 2005 Founder and Chief Executive Officer. Atla Labs, Inc., Los Angeles, California.
- 2006 Award Winner. International Solid State Circuits Conference and Design Automation Conference Student Design Contest, Operational System Design Category for “Illumimote: A High Performance Light Sensor Module for Wireless Sensor Networks.”
- 2007 Teaching Assistant. Thick Film and Hybrid Fabrication Laboratory, University of California, Los Angeles.
- 2007 Masters of Science in Electrical Engineering. University of California, Los Angeles.

- 2008 Award Winner. International Solid State Circuits and Design Automation Conference Student Design Contest, Operational System Design Category for “XCXO: Crystal Compensated Crystal Oscillator.”
- 2008 “Best Paper” Award. “NAWMS: Nonintrusive Autonomous Water Monitoring System,” The 6th ACM Conference on Embedded Networked Sensor Systems
- 2009 “Best Paper” Award. “Software-Defined Underwater Acoustic Networking Platform,” Proceedings of the ACM International Workshop on UnderWater Networks.
- 2011 Visiting Scholar. Golshani Laboratory for Neuroscience, University of California, Los Angeles.

Part I

Introduction

CHAPTER 1

Ocean Systems

1.1 Submerged Wireless Sensor Networks

The development and deployment of wireless underwater networks, unmanned marine robots, and sensors [Rho07], while receiving much attention [Int] [Aqu09] [WYH06] [Lin] [Ins06], has been limited due to the lack of an accommodating physical channel [Tri09]. Acoustic pressure-wave systems suffer from frequency-dependent bandwidth and attenuation, time-varying multi-path, and the low speed of sound in water. Optical systems do not work well in the dark turbid oceanic waters. Conventional radio systems depend on electromagnetic propagation that, experimentally, results in extremely high signal attenuation [SK73] given the frequency-dependent permittivity of seawater [ASS04].

1.2 Can We Learn In-Ocean Networking From Fish?

Nature, on the other hand, has had millions of years to develop and optimize life in the ocean. Many fish species are organized into social communities and society requires effective and reliable communication to flourish. Survival itself dictates that jamming-avoidance (multiple access), self-recognition (modulation), localization, and all-weather availability (active signaling) are requisite qualities in a physical channel – so evolution has naturally selected for them.

Nocturnal ocean animals and those that live at depth are unable to employ

light-dependent forms of perception to navigate their surroundings, hunt, or avoid predators [Eva97]. Rather, the primary means of perception for many Teleost and Chondrichthye fish species involves the passive detection of electrical fields necessarily created by muscle activity in their prey (a process similar to the human electrocardiogram) [Nel11a]. In some species, the process is an active one in which the generation of a Voltage gradient in the intermediate environment enables the discernment of non-emissive navigational hazards (such as rocks) by the manner in which they disturb the self-generated field [Kra94].

Electroreceptivity is achieved through a dense grid of electric field (Voltage) sensors which are concentrated heavily around the head of the fish and tightly arrayed along either side of the body [FB99]. Anatomically, each sensor is constructed of a glycoprotein gel-filled tube with nerve endings at either end [Lor05][Eva97]: altogether this system of sensors is labeled the Electro-Sensory Organ (ESO) and is used in the detection of active, e.g. emissive, targets. The ESO, by itself, is unable to perceive non-emissive objects [BHG69]. The inclusion of an Electric Organ (EO) rectifies this problem. The EO, through the direct conduction of current generated in the EO, establishes the aforementioned Voltage gradient in its surroundings, an event which is known as an Electric Organ Discharge [MSM01].

To comprehend the flow of current around the fish in the ocean picture two electrodes placed at either end of a wide flat conductive plate. Once the plate is energized, the majority of the current will flow from cathode to anode, with little or no deviation from the most direct path between the electrodes [PR96]. However, as current in any direction is comprised of like-polarity charges, a repulsive force exists between them causing some charges to take an indirect, more circuitous route. The current is following the isolines of the charge concentration gradient created by the electric field established by the electrodes [HRW97].

If an object less conductive than the ocean water (rocks, plastics, etc) enters the field, the current, following the path of least resistance will curve around it. The

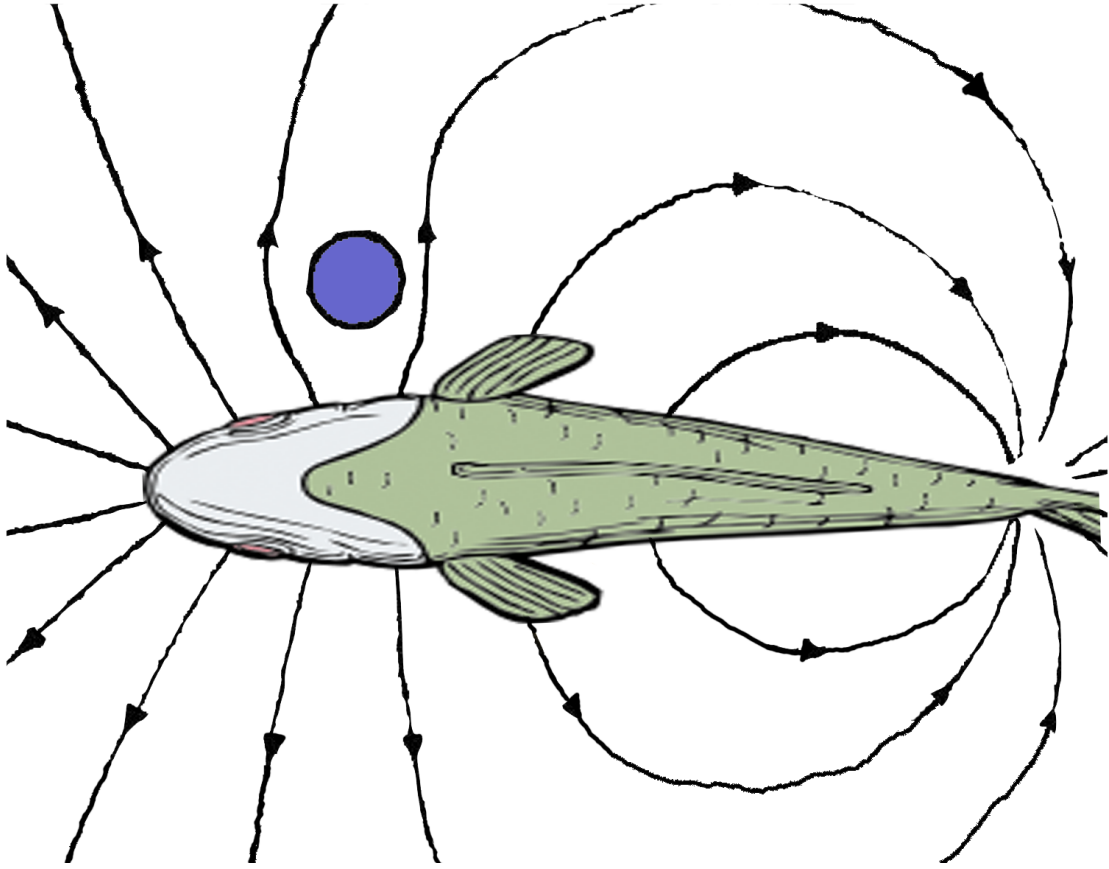


Figure 1.1 – *Electroreception is the ability of some fish species to sense the environment around them by either passively listening or actively detecting changes in a self-made electric field emitted around them. The presence of objects with different conductivities from the ocean redistribute the field lines casting a detectable electrical shadow on the body of the fish.*

redistribution of current necessarily spreads out the field lines which changes the location of the isovoltic lines' intersection with the fish's body: it effectively casts an electrical "shadow" by creating a region where the Voltage is more constant per unit distance along the body [Kra94] [Bul99] [MSM01]. Objects more conductive than the surrounding water have the opposite effect; they concentrate the field lines, establishing an electrical bright spot – a region of rapid Voltage change per unit distance along the body. As such, when electroreception is an active sensor, this mimetic system can not only detect objects and their location, but may classify them as well.

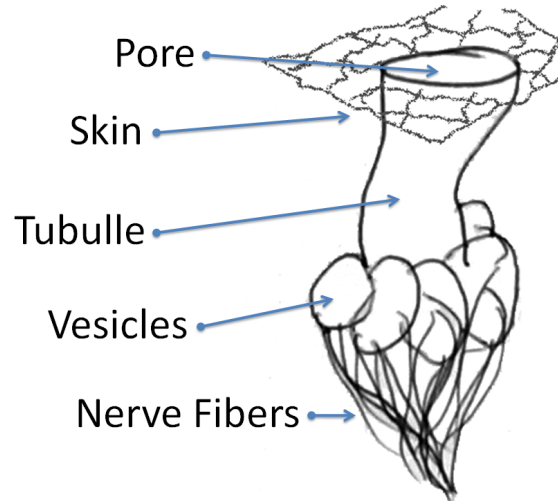


Figure 1.2 – *The Ampulla of Lorenzini is a fundamental component of the electric-field sensing organ in some Teleost and Chondrichthyes fish species.*

1.3 Electrosensory Organs in Fish

In the eighteenth century, Italian biologist Stephan Lorenzini observed a peculiar pore and organ system in crampfish [Lor05] and theorized that they might be used for navigation and hunting. These structures were later named the Ampullae of Lorenzini in his honor. However, the purpose of the ampullae was not clearly understood. Electrophysiological experiments suggested a sensibility to temperature, mechanical pressure and possibly salinity [FB99]. It was not until 1960 that the ampullae were clearly identified as specialized receptor organs for sensing electric fields [BHG69] [Kra94] [TZ05].

Each ampulla is a bundle of sensory cells containing multiple nerve fibers. These fibers are enclosed in a gel-filled tubule which has a direct opening to the surface through a pore. The gel is a glycoprotein based substance with the same resistivity as seawater [Eva97]. Consequently, the ampullae can detect electric fields in the water through the Voltage differential at the skin pore versus the base of the electroreceptor cells – that is, the difference in neurological activity between the terminal axons at the pore and those in the interior vesicles of the ampulla

[FB99] as shown in figure 1.2.

Two broad categories of electroreceptors have evolved. *Ampullary* electroreceptors are associated with passive detection of inadvertent signals emitted by biological organisms as they move in the ocean. Pore systems in the skin of fish and gaps in their scales permit muscle movement and other neurological activity to conduct out into the surrounding environment by allowing internal ion concentrations, established across cellular membranes, to create an electric field. *Tuberous* receptors are associated with active electroreception and are found in species with an Electric Organ (EO). The EO is responsible for actively emitting an electric field.

1.4 Contributions

In this work, I have endeavored to mimic the biological capacity for electroreception in an artificial engineered system and demonstrate its utility across a broad cross-section of sub-surface in-ocean application requirements. For generality, my work centers on the physical means of realizing electroreception – the submerged electrostatic physical information channel.

This thesis is broadly organized into three parts which follow a logical progression. The first part involves the physics, design, and optimizations necessary to access a submerged electrostatic channel. The second part treats waveform design – that is, encoding information in the channel. The third part covers my work in developing an active sensor. Necessarily, this entails the transmission of an encoded signal and the detection of environmental disturbances to that signal. Consequently, communication, coordination, localization, navigation, or even actuation through the channel is also possible.

The ability to locate and “visualize” targets with different conductivities from the background ocean environment proffers a great range of new possibilities for

underwater exploration. A working short-range proximity sensor could be used for underwater wireless network neighborhood discovery and station keeping. In the case of mobile network nodes, this sensor could assist in collision avoidance and formation management.

The need for such technology is clear. As of 2001, the United States Navy (USN) submarine surfacing procedure calls for its submarines to maintain position underwater for at least 180 seconds while actively pinging its sonar in order to establish the presence of surface vessels posing a collision hazard as the submarine comes to the surface [Nav01]. In at least one instance, procedure was ignored to disastrous effect when the USN *Greenville* struck a Japanese civilian fisherman training ship with high school students aboard while surfacing [NT01]. Had an electrostatic imaging system been developed and deployed aboard, the crew would have had access to instantaneous real-time video imagery of the electrical disturbances caused by the presence of the surface vessel and the collision might have been avoided.

The problem is pervasive. In the past ten years (2001-2011) there has been an average of one major collision event per year involving military submarines at an average repair cost of \$78 million USD per incident¹[Tin09]. Biomimetic electrostatic imaging shows tremendous promise in submerged salt-water navigation and collision avoidance since it is built on an effectively instantaneous phenomenon and does not require an active target.

The contributions of my work include (1) the application of theory to explain the observed performance of submerged oceanic electrostatics, (2) the development of a model to predict future design improvements, (3) experimental proof of the existence and utility of the phenomenon, (4) an engineering validation of the rationale for the naturally observed weak-electric fish waveforms, and (5) the design and implementation of the first *Biomimetic Electrostatic Imaging* (BEI) systems –

¹Repair estimates were not available for all incidents.

an engineered sensor which mirrors the biological phenomenon of electroreception. This dissertation includes validations of the 3-electrode [FTC10], 16-electrode [FHT11], and, ultimately, 68-electrode versions of the BEI hardware.

CHAPTER 2

Physical Channels

2.1 Overview

To confer maximum survival advantage, an ideal perceptive system should be fast (low-latency), available (all-weather), unnatural (low background noise), communal (widebandwidth, to support multiple organisms in close proximity), social (transmit and receive), and have minimal impact on the basal metabolic rate (low power).

As an aggregate whole, fish possess up to four primary senses employed in external environmental monitoring: (1) A chemotactic sense of smell/taste, (2) a baroreceptive sense of sound/touch, (3) a photoreceptive sense of sight, and, most interestingly to the purposes of this treatise, (4) an electroreceptive sense of static electric charges – e.g. *voltage measurement*. This is in addition to the host of proprioceptive senses, which aid survival through internal monitoring: equilibrioception (orientation), thermoception (temperature), nociception (pain), etc.

Not all fish possess all of these senses and the degree to which each has evolved, and its corresponding acuity, varies by the needs of the particular species. In general, fish that live in salty (oceanic), turbid conditions and/or exhibit nocturnal and benthic (bottom-dwelling) behavior cannot rely on optical notions of vision to navigate, forage, hunt, and evade capture. Similarly, the silty turbulent conditions of the near-bottom littoral region obscures most of the pressure cues and homogenizes the chemical ones. For this class of animal, the near-instant

phase velocity and self-determination of electroreception is extremely attractive. Throughout history, electroreception, as an organ system in fish, has evolved independently at least five times [Zel11]. Nature’s strong preference for electric field perception in this class of animal is more easily understood by considering its alternatives with respect to the aforementioned criteria.

2.2 Concentration Gradient (Chemotaxis)

Many animals use a sense of “smell” to locate food, mates, shelter locations, and avoid danger [Eva97]. Sensory notions of odor and smell refer to an organism’s ability to detect and identify chemical compounds. The intensity of the odor is related to the concentration of the underlying chemical [Kra94]. While odor plumes are complex, dynamic, three-dimensional structures, odor itself has no directional properties [GA07]. Animals must therefore use a variety of different senses to obtain directional information. Since most odor plumes are composed of dispersing odor patches and dissipating vorticity eddies, aquatic animals may localize odor sources by simultaneous analysis of chemical and hydrodynamic dispersal fields, a process referred to as *eddy chemotaxis* [GA07].

In order for chemotaxis to be possible, the odorous compounds must come into physical contact with chemoreceptors in highly specialized cells in the sensory membranes of the animal. Accordingly, chemoreception requires a physical transfer of material from the prey animal (or other target phenomenon) to the predator (the sensor system under design). In an experiment conducted by Jayne M. Gardiner and Jelle Atema of the University of South Florida, the scientists were, for the first time, able to demonstrate conclusively that sharks require both olfactory and lateral line input for efficient and precise tracking of odor-flavored wakes and that visual input can improve food-finding performance when lateral line information is not available [GA07].

Momentarily ignoring the finding that chemotaxis is impossible without joint reliance on baroreception or photoreception [GA07], a nuanced definition of sensor range may be developed. In the scenario of a stationary target in still water, the only means by which the chemical odors propagate is through diffusion. This is a very slow process with concentration obeying a square law – increasing in range, but decreasing in concentration at the square of range. On the other hand, with a fleeing target, the pursuing predator moves into (and tracks) the odor plume. Lastly, given the physical nature of the chemical (it has mass), odors remain in the area of their deposit for a substantial amount of time when compared to other energy-based sensory phenomenon. We may therefore classify chemotaxis as an extremely high latency, low bandwidth, passive, potentially extremely long-range communication/sensory channel. It is an inappropriate model for the class of problems I aim to redress in this dissertation.

2.3 Acoustics (Baroreception)

At a molecular-level, acoustic communication systems rely on displacing the mass of the molecules in their way. Accordingly, pressure waves can not exist without matter and propagate better (e.g. faster) as density increases, but at a higher energy per distance cost. Physical displacement requires mechanical action and is limited in frequency and amplitude by the capabilities of mechanical excitation. Long-range (large amplitude) vibration can not be achieved at frequencies in excess of the order of kHz and any vibrating sources in the environment will interfere (including reflections of the transmitted signal from the surface, bottom, and suspended objects).

Sound waves can travel great distance in the ocean and are unaffected by the high permittivity of ocean water, but lack the spatial resolution necessary to form the basis of vision [Pre07]. Thermal gradients, suspended particles, wave motion,

and general turbidity all severely hamper or distort sound wave propagation. Sound, by its pressure wave nature carries a great energy cost in generating signals as the organism must work against the mass of the ocean to initiate propagation and, due to the limited bandwidth, must outshout competing organisms for access to the channel [SSP00].

To better quantify the energy cost of mass-displacement (acoustic) systems with respect to energy-displacement (conductive or radiative) systems we define the term “Communication Efficiency”, η_{comm} .

$$\eta_{comm} = \frac{Power}{\frac{Distance}{DataRate}} = \frac{\mu Watts}{\frac{Meters}{\frac{bits}{second}}} = \frac{\mu Joules}{Meter Bit} \quad (2.1)$$

Communication efficiency is a measure of how much energy it takes to move a certain amount of data, in a certain amount of time, a particular distance. For perspective, consider figure 2.1, a representative survey of both commercially-available and research-prototype acoustic (pressure-wave) communication systems.

Figure 2.1 is a bit generous in that it was created by using the maximum reported range, communication rate, and power. However, acoustic systems often do not reach their maximum data rate and maximum range simultaneously. Data rate degrades to compensate for increased multi-path interference from the higher output powers required as range is increased.

The last two entries in figure 2.1, whose value is barely visible, are our UWET system¹ using biomimetic artificial ampullae at 8MHz and 1MHz carrier frequency respectively. An enlarged view of the last four entries of figure 2.1 appear in the inset section. *UWET is at least two orders of magnitude more efficient and can achieve communications bandwidth more than two orders of magnitude higher than the surveyed acoustic systems.*

¹The UWET system using Artificial Ampullae (AA) is covered in chapter 4 and evaluated in section 4.6.




Figure 2.1 – *The communication efficiency of a representative sample of commercial and research underwater pressure-wave communication systems. From top: [Int], [WYH06], [Lin], [Lin],[Ins06], [Lin], [Lin], [Lin], [Tri09], [Lin], [Lin], [Lin], [Aqu09]. (Inset) The communication system designed and implemented in this work (section 4.6) is at least two orders of magnitude more efficient.*

2.4 Electromagnetics (including Photoreception)

Electromagnetic (radio) systems do not really have any analogs in the biological realm, yet they form the overwhelming basis of human-engineered communication and sensory systems. As such, they deserve some consideration. Optics (light) is an extreme form of electromagnetics and suffers many of the same issues as its lower-frequency brethren. Electromagnetic systems rely on aligned charge movement releasing energy. As charges move, reverse direction, and move again the conservation of energy laws require that momentum dissipate prior to the reverse in direction. This extra energy is radiated outward in the form of electromagnetic (EM) waves. EM waves propagate best in a vacuum where there are no intermediate

particles to collide with, which would result in absorption and scattering.

Electromagnetic propagation through water is very different from propagation through air because of water's high permittivity and electrical conductivity. Plane wave attenuation is high compared to air and increases rapidly with frequency. The loss of energy occurs primarily through two complementary mechanisms. The mobile charges of a conductive environment respond to the incident EM wave, which by absorption increases their energy level. By acting to restore their former lower energy state, the charges radiate EM waves with opposite polarity (180deg out of phase). At some distance away (the far-field) the incident and re-radiated (scattered) waves appear near equal in amplitude and add destructively. At the same time, this very transfer of energy itself, from wavefront to mobile environmental charges, is subject to thermal losses.

We can derive the power losses expected for an electromagnetic wave propagating in the ocean by following the approach of Andy Wolski of the University of Liverpool [Wol11]. We begin with Maxwell's equations:

$$\nabla \cdot \vec{E} = 0 \tag{2.2}$$

$$\nabla \cdot \vec{B} = 0 \tag{2.3}$$

$$\nabla \times \vec{E} = -\dot{\vec{B}} \tag{2.4}$$

$$\nabla \times \vec{B} = \mu\epsilon\dot{\vec{E}} + \mu\vec{J} \tag{2.5}$$

Given that the ocean environment is well modeled as an ohmic conductor, the current density, \vec{J} , becomes:

$$\vec{J} = \sigma\vec{E} \tag{2.6}$$

with σ representing the conductivity.

Substituting equation (2.6) into equation (2.5), the Maxwell-Faraday² and

²Equation (2.4)

Maxwell-Ampere³ equations may be combined to yield the wave equation for the electric field in a conducting material:

$$\nabla^2 \vec{E} - \mu\sigma \dot{\vec{E}} - \mu\epsilon \ddot{\vec{E}} = 0 \quad (2.7)$$

While equation (2.7) only describes the electric component of the electromagnetic wave, it can be shown that explicit consideration of the magnetic field is unnecessary. The only solution for Maxwell's equations at all possible positions and times occurs when the magnetic and electric fields have the same wave vector and frequency. The varying electric field is, therefore, invariably coupled to the magnetic field associated with it. As the electric field amplitude changes, so does the magnetic field's – solving for either is sufficiently illustrative.

To find this solution Wolski proposes a wave equation of the form:

$$\vec{E}(\vec{r}, t) = \vec{E}_0 e^{j(\omega t - \vec{k} \cdot \vec{r})} \quad (2.8)$$

which introduces two new terms: one to express the distance and direction into the conductive medium (\vec{r}) and a second term (\vec{k}) to express the wave vector. Combining equations (2.7) and (2.8) gives the dispersion relation:

$$-\vec{k}^2 - j\omega\mu\sigma + \omega^2\mu\epsilon = 0 \quad (2.9)$$

For simplicity, we separate the dispersion relation into its real ($\vec{\alpha}$) and imaginary ($\vec{\beta}$) parts:

$$\vec{k} = \vec{\alpha} - j\vec{\beta} \quad (2.10)$$

where

$$\alpha = \omega\sqrt{\mu\epsilon} \sqrt{\frac{1}{2} + \frac{1}{2}\sqrt{1 + \frac{\sigma^2}{\omega^2\epsilon^2}}} \quad (2.11)$$

³Equation (2.5)

and

$$\beta = \frac{\omega\mu\sigma}{2\alpha} \quad (2.12)$$

These equations form an expression for the wave vector which is solely a function of the wave's frequency and the properties of the material in which it propagates.

In its final form, and using phasor notation, we see a product of two phasors:

$$\vec{E}(\vec{r}, t) = \vec{E}_0 e^{j(\omega t - \vec{\alpha} \cdot \vec{r})} e^{-\vec{B} \cdot \vec{r}} \quad (2.13)$$

It is the second of these, $e^{-\vec{B} \cdot \vec{r}}$ that is most interesting as it provides the decay behavior. In an ideal non-conducting material there is no exponential decay of the amplitude. The electromagnetic wave will travel indefinitely due to the lack of any loss of energy. If the wave enters an electrical conductor, however, very different behavior results.

The electric field in the wave will cause currents to flow in the conductor. When a current flows in a conductor⁴ there will be some energy converted to heat. This energy must come from the wave. Therefore, we expect the wave to decay with distance into the conducting media.

The attenuation of EM waves in any conducting medium increases both with an increase in conductivity and an increase in frequency [SK73]. Within water, increasing the salt content (salinity) increases conductivity and its attendant EM attenuation. A further simplification of equation (2.13), in the form of an approximation, is equation (2.14):

$$\alpha = 0.0173\sqrt{f\sigma} \quad (2.14)$$

where α is attenuation in dB per meter, f is frequency in Hertz, and σ is conduc-

⁴The only exception is a super-conductor – a conductor without resistance.

tivity in Siemens per meter [Moo67].

As is evident from equation (2.14) (given that $\sigma \approx 4.8$ S/m at 20°C [Tur68]), attenuation in sea water is very high. To detect or communicate via EM waves at any depth at all, it is necessary to use very low frequencies (on the order of 10kHz) where attenuation is less than 5 dB per meter. Even at this level, the energy cost of detection/communication over distance will rapidly outpace that of acoustics despite having a short-range advantage [FTS10]. For this reason, human under-sea efforts, as well as biology, have largely abandoned EM-based sensing and communication.

2.5 Electrostatics (Electroreception)

The high conductivity, turbidity, absorption (of light and other EM waves), and mass of the ocean make reliance on light, radio, and sound waves for communication and perception impractical for many species. In response, many of these species have evolved an electroreceptive capability, wherein they are able to detect (and in some cases manipulate) an electric field. To distinguish this behavior from the human-engineered radio (EM detector), the area of electric field study is known as electrostatics.

Electrostatics, as the name implies, involves the use of electric (E) fields which are, traditionally, invariant with time. For our purposes it is illustrative to consider a sequence of time-invariant fields with each successive field having more (then less, then more, etc) strength than its predecessor. When this approximation is valid, the field is said to be *quasi-static* [MT76].

In order to prove the validity of this assumption, consider that when the electric field, and hence the current flowing in the field, changes with time two currents must be considered – the *conduction* current and the *displacement* current.

The conduction current is Ohmic [HRW97] resulting from the movement of

charges between atoms⁵, while the displacement current results from the movement of bound charges within an atom (typically caused by the application of some external field) [PR96].

The most obvious example of displacement current is the vacuum-gap capacitor (two plates with nothing between them) in which the same amount of conduction current enters one plate as leaves the other. Between the plates, in the vacuum, there can be no conduction, but a magnetic field exists there as if a conduction current were flowing [Max65]. This phantom current is the displacement [Max55].

$$j_{conduction} = \sigma E \quad (2.15)$$

The conduction current is defined by equation (2.15) [MT76], where $j_{conduction}$ is the current density, σ is the unit length conductivity, and E is the applied field strength. By inspection the similarity to Ohm's law ($I = V/R$), where $\sigma = 1/R$, is obvious. As conduction current is Ohmic, it is stateless and time-invariant. However, the displacement current, defined by equation (2.16)⁶ [MT76], is based on the first time derivative of the electric field intensity and, accordingly, is not time-invariant.

$$j_{displacement} = \epsilon \frac{\delta E}{\delta t} \quad (2.16)$$

In order to qualify as quasi-static, the field must be mostly (quasi) time-invariant (static). That is, for a sinusoidal time-varying field $E = E_0 e^{j\omega t}$:

$$\frac{j_{displacement}}{j_{conduction}} = \frac{\epsilon}{\sigma} \omega \ll 1 \quad (2.17)$$

Ocean water, with $\epsilon/\sigma \approx 10^{-9}$ [Rho07], is quasi-static to almost 100 MHz. This concurs with the observation that EM radiation is transmitted poorly in water,

⁵...or the movement of charged atoms (ions) in space, as is common in biological systems

⁶ $j_{displacement}$ is the current density, ϵ is the permittivity of the medium, and E is the field intensity

while direct conduction, as will be shown, works extremely well.

Finally, it is important to note that despite the applicability of static analysis, quasi-static fields are time variant. The phase velocity of electromagnetic phenomenon decreases as a function of the permittivity. For seawater, with a permittivity near 80 [Rho07], the propagation velocity is two orders of magnitude slower than in air (equation 2.18).

$$v_{\phi} = \frac{c}{n_r} \tag{2.18}$$

This lower velocity results in an equivalent reduction in the near-field radius and a corresponding improvement in resolution for localization algorithms applied over this channel – a fact we will exploit in later chapters.

2.6 Conclusion

From the biology emerges a clearer statement of the class of problems I intend to target. Sensing and communication at long ranges is best handled through acoustics and chemotaxis. At shorter ranges, in fresh and brackish waters, needs are best dealt with via optical means. In the ocean, at the close ranges where an electric field may be extended, electroreception is the clear choice. It is the work of this dissertation to define the optimal design and operating conditions for such an engineered biomimetic system and therefore demonstrate when and how it should be used to greatest effect.

A recurring point of confusion expressed by peer-reviewers and colleagues attendant to presentations on this work is its differentiation from radio, the propagation of electromagnetic waves. For intuition building purposes, one might consider electrostatic coupling as the low frequency end of radio communication since the majority of energy coupling to any receiver within a wavelength of a

radio transmitter is overwhelmingly electrostatic or magnetostatic (the near-field or inductive field region) [Wes91] [Mil93] [RWD94]. In air, where this discussion is usually applied, the main advantages of radio wave energy propagation, as compared to electrostatic coupling, are higher bandwidth and longer range.

Another facet of this longer range is energy efficiency. Given that range can typically be improved with increased output power, of prime interest is the ratio between the two rather than absolute numbers. In air and at range, EM transmission is a clear winner. At shorter ranges (within a wavelength or two), the energy advantage dissipates given the high carrier frequencies necessary to keep short range receivers more than several wavelengths away.

At 2 meters, the operating distance of the carrier frequency of a radio transmitter would need to be in excess of 150 MHz to operate at far-field in air – *regardless of the amount of data carried across the channel*. This comes with a large energy penalty as energy and frequency are linearly correlated to the first order. Compared with a base-band electrostatic transmission, which can operate at the frequency of the data payload, a huge inefficiency occurs. For a 10 kbit/second data stream, electrostatics achieves an advantage of more than 15,000 \times assuming the receiver is within range.

To capitalize on this potential, the next chapter begins part 2, which looks both theoretically and experimentally at electrostatic channels and how best to use them.

CHAPTER 3

Related Work

3.1 Overview

It is both typical and premature to encounter the related work at this juncture. It is premature in that the reader might not yet have a full appreciation of the technique and approach of this dissertation. If desired the reader might skip past this chapter and return with a deeper understanding of the my approach. However, if the reader desires an immediate understanding as to why this work exists, it is best to soldier forward with the comparative analysis of this chapter.

To assist in this aim, keep the following in mind. In short, this work may be broadly classified as Biomimetic Electrostatics (BE) to consist of a design with the following characteristics:

1. Submerged in-ocean operation – the goal of the BE system is in-ocean and submerged operation, mimetic of the lifeforms that inspired it
2. A surface array of electrodes – the targets are forward of, and not surrounded by, the sensor as is necessary for tomographic applications of electrostatics.
3. Conjoined transmit and receive (in the case of imaging) – transmitter and receiver may share a ground and timing reference of extreme quality.
4. High-impedance receive electrodes – high sensitivity measurements from electrodes in direct contact with the seawater. Input and output current (leakage) is minimized.

5. Low-impedance transmit electrodes – direct coupling of an alternating voltage source via the electrodes to the seawater
6. Determination and detection of the *differential disturbance* field (in the case of imaging) – the phenomenon of interest is the change in a self-generated calibrated background field as observed across two receive electrodes within the field.

The necessity and biomimicry of each of these characteristics will be justified and evaluated throughout the remaining chapters. At least a course familiarity with this list is necessary to properly comprehend placement and relation to each of the prior art categories that now follow.

3.2 Commercial Capacitive Sensing Integrated Circuits

From an engineering perspective, electric fields, in the form of capacitive sensing, have oft been explored [Bax96], [Lee06b]. The core concepts are similar in that an electric field is created between two electrodes and disturbed by environmental factors of interest.

Indeed, there are now Application-Specific Integrated Circuits (ASIC) microchips wholly dedicated to the task [Sem] (see figure 3.1). They find use as touch-screen controllers, volume sliders, proximity sensors, and as any number of other human machine interface devices. These ASIC’s are designed for in-air operation and make certain design choices and assumptions attendant their operative intention.

Typical capacitive sensor designs require a relatively short (thin) dielectric separation between electrodes and environmental disturbance (e.g. finger in touch sensing applications). According to noted application engineer Mark Lee of Cypress Semiconductor, “most specify an overlay of 3 mm or less.” Disturbance sensing



Figure 3.1 – *Cypress Semiconductor’s TrueTouch touch-screen controller Application Specific Integrated Circuit (ASIC) shown here as reproduced from the company’s marketing materials.*

through a dielectric becomes increasingly difficult as the thickness of that dielectric increases. That is why the current upper bound is widely accepted to sit at 10 mm of overlay thickness [Lee06b].

Capacitive-sensing systems come in many configurations. At their core, there is always a set of conductors which generate and detect electric fields. The human body is filled with water, which is contained either inside cells or outside of them [Zel11]. Cellular membranes are phospholipid bilayers which are non-conductive. Therefore the water inside cells, despite its high ion content (and, therefore, conductivity) is inaccessible to conduction currents in the extracellular fluid. Therefore the relatively low internal resistance of the human body is an almost exclusive result of the extracellular fluid that both circulates in the bloodstream and permeates the organs and intra-organ cavities [SPL00]. These conductive fluid-filled cavities are covered by a layer of skin, a non-conducting lossy dielectric.

In a terminal digit of a finger, the structure you most commonly touch with, the effect is most pronounced as there is little else nearby (major muscles, heart, brain, etc) [HRW97], [Lee06b].

It is the conductive property of fingers that makes capacitive touch sensing possible. Mobile charges in the form of ions in the extracellular fluid translate and rotate to align with the external field permitted through the skin. To better understand the interaction, consider the parallel plate capacitor. A simple parallel plate capacitor has two conductors separated by a dielectric. Most of the energy is concentrated in the space intermediate the plates, however some of the energy is captured in the area outside the plates [HRW97]. This external energy forms the “fringe” field because the electric field lines associated with this energy are highly curved between edges of the plates, resembling fringing hair follicles or other extraneously curly filaments [PR96].

In BE, as in capacitive sensing, the energy of maximum utility is that associated with these “fringe” fields. The design challenge is one of maximizing the area/volume coverage of these fields and generating them with reasonable efficiency.

Part of the challenge of making a practical capacitive sensor is to design a set of printed circuit traces which direct fringing fields into an active sensing area accessible to a user [Lee06b]. A parallel plate capacitor, as we investigate in subsequent chapters, is not a good choice for such a sensor pattern.

As just described, the finger is itself a capacitive structure. Placing a finger into the fringing electric fields adds capacitance because it adds additional charge storing surface area to the existing capacitive system. A common misconception about capacitive sensors is that the finger needs to be grounded for the system to work.

A finger can be sensed because it can hold a charge not because it provides a

return path to ground. It's also why dead flesh is less effective (less water content to suspend ions). For this reason, biometric fingerprint scanners are introducing capacitive sensing to combat the gruesome attack scenario often portrayed in the movies – using a severed finger to defeat a fingerprint scanner. With the fusion of capacitive sensing the dead flesh would be discernible.

That's where the similarities end. Capacitive sensing makes use of combined excitation and sensing electrodes, since the phenomenon of interest is the capacitance between those electrodes. Measuring the capacitance is really just a measure of the additional charge storage capability imposed on the reference capacitor (electrodes) by the target. When submerged, the ocean itself provides so much additional charge storage at such immediate proximity that the effect of the target's presence is barely observed.

For this reason BE makes use of separate transmit and receive electrode pairs, uses direct conduction into the environment for excitation (rather than coupling through a dielectric), and does not attempt to characterise capacitance, electing instead to focus on force.

3.3 Electrical Impedance Tomography

Electrical Impedance Tomography (EIT) is an imaging method, in which the distribution (image) of conductivity of a volume enclosed by a ring (or several rings) of electrodes is desired to be known.

Electrical currents are injected through electrodes into the encircled conducting volume and the resulting potential distribution is measured at the remaining electrodes [DR00]. From these potentials, an image of the electrical conductivity can be obtained using numerical reconstruction techniques [Mol03]. This non-linear reconstruction is mathematically difficult and computationally intensive [Pol02].

In practice, EIT is a measurement technique for obtaining information about

the contents of process vessels and pipelines [Sys11], [LW93], [MMD98], [Sta06]. In commercial application EIT can be used for rapid multi-phase systems such as the measurement of flowing oil water systems and hydraulic conveying of slurries. An example of EIT designed for earth exploration is shown in figure 3.2. It is also used in a limited number of medical applications [DR00], [Mol03], [Pol02].




Figure 3.2 – *A model of Stanford’s proposed EIT rig for earth exploration. Side view (left) and top view (right) of the electrode and environment arrangement. 48 electrodes were attached to a Berea sandstone core using conductive silver. The core is 5 cm in diameter and 10 cm tall. Reproduced from [Sta06].*

Most applications in medicine and industry rely upon a fast and accurate image acquisition [Pol02]. Consequently, most EIT research aims to find methods which improve the speed and accuracy of EIT by a range of improvements to the numerical methods used in the forward solution and inverse reconstruction.

This is not unlike the BE imaging (BEI) proposed in this dissertation. The key differences lie in the operational constraint that EIT requires multiple electrodes arranged around the boundary of the vessel at fixed locations in such a way that they make electrical contact with the fluid inside the vessel but do not affect the flow or movement of materials, whereas BEI uses electrodes exclusively at the surface (non-encircling).

In EIT, current between a pair of electrodes is introduced to the intermediate

material and voltage measurements are taken between remaining electrode pairs. The next pair of electrodes is then selected for excitation and the process repeated until all combinations have been exhausted. This interrogates an entire “slice” through the measurement zone - analogous to a “body-scan” in medical imaging.

In BEI, current is transmitted between a fixed reference pair of electrodes optimized for the task. The current however, is largely irrelevant. It is a consequence of establishing a charge gradient (force field) in the intermediate environment, which in turns effects the flow of current near the detection electrodes. The field is detected by a non-encircling (typically planar) array of electrodes. The background is subtracted out as the reference condition and only changes introduced to the field by the presence of an object differing in conductivity is recorded.

3.4 Electrical Capacitance Tomography

Although like EIT in every other respect, ECT operates with high input impedance (e.g. through a dielectric) and looks for changes in capacitance caused by the intermediate environment. Because it must sense subtle variations in capacitance the electrodes must be large, limiting the number deployable per unit volume and, therefore, the resolution of the images produced. Capacitance measurements require direct pairing of electrodes which limits the total number of independent measurements to:

$$n_m \frac{n_e(n_e - 1)}{2} \tag{3.1}$$

where n_m is the number of independent measurements and n_e is the number of electrodes.

In [HPX88], an instrumentation system for obtaining cross-sectional images of two-component flows via this method was described. Their system used measurements made with capacitance electrodes mounted on the outer surface of the fluid-conveying pipe, and an image reconstruction algorithm based on backprojec-

tion. Images of static physical models simulating a stratified solids/air flow were presented.

As in [HPX88], ECT images are normally of relatively low resolution (a 32 x 32 pixel grid is typical) but they can be captured at high speed. With existing technology, image data can be captured at 100 frames per second for a 12 electrode sensor and displayed in real time [Pro10]. Reducing the number of electrodes, as in BE imaging, reduces the computational burdern. Higher image capture rates (up to 300fps) can be achieved for 6 or 8 electrode sensors [Pro10], [YP03].

Since it is a form of tomography, an ECT image shows the cross-sectional distribution of the vessel contents. This distribution is averaged over the length of the sensor electrodes. As in BE imaging, the spatial resolution achievable depends on the size and radial position of the target object, together with the target contrast – its permittivity difference relative to that of the other material in the pipe.

Typically, target objects (or local changes in permittivity) with a diameter 5% of that of the pipe or vessel, around which the ECT system is deployed, can be detected provided there is sufficient contrast between the permittivity of the target and the surrounding media. The accuracy of the image depends on the method used to construct the image from the inter-electrode capacitance measurements [YP03].

At present, the only image construction algorithm which is fast enough to be used for on-line image display is the linear backprojection method [Pro10]. This produces approximate images which are of acceptable quality for many applications. Other methods can be used to produce improved images off-line from captured capacitance measurements [HPX88]. These methods involve the use of iterative computational methods, or alternatively, the application of neural network techniques [YP03].

BE imaging as described first in this chapter’s overview and again in subsequent chapters, does not rely on an encircling array of electrodes and uses low, rather than high, impedance excitation electrodes. It is capable of substantially higher resolution images at an equivalent speed to ECT.

3.5 Submerged Antennas

The use of electrostatics in the ocean is more muddled. Many authors have explored the direct submergence of aerial antennas ([SK73], [Moo67], [Wol11], [Rho07], [Wes91], [Mil93], [RWD94], et. al.). Ultimately these works conclude that the severe attenuation of electromagnetic radiation in sea water is too great for practical employment.

A few authors have reached outlying conclusions. Recent experimental work by [ASS04] has reinvigorated the debate when they reported very low losses intermediate to submerged and airborne propagation when emitting megahertz order frequencies in salt-water. Their explanatory figure of the propagation model bears striking resemblance to that of electrostatics. Given that we have observed similar intermediate losses during early expeditions (reported in the next chapter), it is my belief that they experienced mixed mode propagation (see section 5.5) rather than propagation via a direct underwater pathway.

3.6 Body Area Networks

In the mid 1990’s Thomas Zimmerman did extensive work with applied electrostatics as part of his concept of an on-body Personal Area Network (PAN) [Zim95] [Zim96]. A PAN can reduce redundancy, as Zimmerman notes, given that a person who carries a watch, pager, cellular phone, music player/recorder, PDA, and notebook computer is carrying five displays, three keyboards, two speakers,

two microphones, and as many as eight radios (multi-band cellular, GPS, Wi-Fi, Bluetooth, NFC, et. al.) [Shi93]. Sharing resources among more limited devices makes sense. To this end, he proposed a PAN transmitter and PAN receiver utilizing capacitive coupling via body-worn electrodes. The transmitter (oscillator) and receiver (differential amplifier) were two terminal battery powered devices. The devices each had a local electrical ground, but were electrically isolated from each other, and therefore did not have a non-galvanic reference between them¹.

In Zimmerman's view the result was a model consisting of four impedances, in his nomenclature, A – D. Impedances A and C represent the coupling capacitance from two different differential electrodes through the body to the respective inputs of a differential amplifier. Impedances B and D represent the capacitive coupling through the air to the earth (ground). There are a few other impedances of note. The intra-electrode impedance between the transmitter electrodes, which presents a load to the transmitter, is ignored as the transmitter oscillator is treated as an ideal voltage source. The intra-electrode impedance between the receiver electrodes is also ignored as it is dominated by the current amplifier which tries to maintain the receiver electrodes at an isopotential².

The circuit is a Wheatstone bridge where any imbalance of the relationship $A/B = C/D$ will cause a potential, and therefore a current, across the receiver. Since the ratios must be exactly equal to null the circuit, and body-based PAN devices are constantly in motion, there will always be an electrical communication path, as long as receiver is sensitive enough to detect the imbalance.

¹No shared ground plane or other low-impedance ground reference

²This results in a low impedance between the receiver plates.

3.7 Conclusion

In follow-on work circa 1995, and culminating in a 1999 Ph.D. dissertation, Joshua Smith proposed biomimetic electroreception as a model for in-air imaging of capacitive changes to an quasi-static electric field. The foundational contributions of Zimmerman and Smith's work are not to be underestimated and have, indeed, lent support and inspiration to this thesis. However, the target applications for both of these engineers lay in air-based schemes. This dissertation is a substantial departure. Our goal is operation in the highly conductive, permeable, and permissible world of the sub-surface ocean.

Zimmerman actually explored this concern as he wrote [Zim96]:

The earth ground provides the return signal. The circuit reveals that body capacitance to the environment E degrades PAN communication by grounding the potential that the transmitter T is trying to impose on the body. For example, in one experiment, standing barefoot reduced communication between wrist-mounted devices by 12 dB.

In the ocean, those losses would be many dB higher.

In air, due to its great impedance, Smith (and Zimmerman) can work with large voltages (30Vpp [Zim95]) and expect only small voltages at their receive electrodes. Accordingly, Smith relied upon a collection of single-ended measurements to describe the diagonal values of the capacitance matrix. In the ocean, this approach is highly impractical as the low impedance of the marine environment results in large received voltages, from which you are interested only in tiny disturbances. The dynamic range problem imposed requires a differential input architecture, in turn predicating an alternate analysis. Our electrodes record a voltage that corresponds to the capacitance difference between them. So in Smith's nomenclature, each of our voltage measurements are a consequence of his original three capacitances

(C_0, C_1, C_2) as they effect one of the receive electrodes as well as three more
 (C'_0, C'_1, C'_2) as they effect its neighboring electrode.

Part II

Biomimetic Electrostatics

CHAPTER 4

Early Experiments with Oceanic Electrostatics

4.1 Overview

Our early investigations into oceanic electrostatics met with limited success. A combination of experimental and conceptual shortcomings resulted in mixed-mode propagation wherein some signal energy was able to travel through the air above the water surface and couple into our instrumentation. These preliminary experiments are described here to provide context to the more successful evolved designs of later chapters.

4.2 Towards Artificial Ampullae of Lorenzini

The Ampullae of Lorenzini organ system in an adult male Oman shark consists of 1400 ampullae per shark with 8 vesicles per ampulla. Each vesicle has 790 sensory neurons. Therefore, the entire animal has 9.3 million receptive terminal axons [FB99]. This equates to just under 6.45cm^2 of base area which was used twice to represent the terminus at the pore.

The Artificial Ampullae (AA) were constructed from an FR-4 fiberglass substrate with a 1oz. copper lamination milled to dimension (figure 4.1). The copper was then coated with a room-temperature vulcanizing silicon sealant to insulate it from direct conduction into the seawater emulating the glycoprotein gel of the natural organ [Eva97].




Figure 4.1 – *Our artificial Ampullae of Lorenzini electrode system inspired by the dimensions of an adult male Oman Shark.*

4.3 Initial Test Configuration

A Software Defined Radio (SDR) development tool, GNU Radio [Blo09], was used to achieve full control over the signals transmitted and received via the electrodes. The lab setup consisted of two Universal Software Radio Peripherals (USRP) connected to two different computers. The computers run from independent batteries isolating them from each other. The USRPs were each fitted with an LFRX and LFTX daughter card. These receiver and transmitter front-ends operate from DC to almost 30 MHz providing the broad spectral coverage needed to study bandwidth.

4.4 Metal Plate Transconduction

When the Artificial Ampullae (AA) are submerged in the ocean, the structure approximates two series capacitors – the electrode-to-water interface of the first AA being connected to the second through a conductive channel (salt-water) modeled as a wire. Accordingly, the output should express electrical transconduction (I_{out}/V_{in}). The electrodes of AA were placed flat against a metal sheet that was coated

with an insulating material (a latex paint) to make it non-contact conducting. The remainder of the experimental setup is unchanged from the water trials. Photographs of the experimental configuration and selected results appear in figure 4.2.

As expected, the behavior appears to be that of two capacitors in series where the transmitting electrode plate couples through electrostatic (e.g. capacitive) action into the conducting body that conducts the momentary current, like a wire, to the receiving capacitively-coupled electrode plate. As the electrode plates are slid apart, the attenuation decreases as a linear function of the distance. This was observed in the lower left panel of figure 4.2.

Below each photo is a sequence of constellation plots for the GMSK modulation waveforms as detected at the receiver and plotted without clock recovery. Because there is no synchronization over time as the data is captured, the constellation rotates about the origin converting the discrete communication quadrants into an apparently continuous circular shape. In doing so, the figures readily illustrate two quantities: (1) the outer radius is an expression of the received signal power, while (2) the inner radius expresses the degree to which the link is open (essentially, the noise margin). As the interior area of the circle decreases, the probability of a bit error increases until the circle fills in and the link closes (BER too high to enable communication).

In the left column of the figure the plates are separated horizontally and in the right column vertically, as shown in the photographs.

If a plate is separated vertically from the conducting body, as was achieved here by stacking sheets of paper beneath it, the signal strength degrades rapidly. Equation (4.1) is the transfer function for a load AC-coupled through a parallel-plate capacitor, where the ϵ terms represent the dielectric, R_L represents the load, A is the area of each electrode, and d is the separating distance between the corresponding electrodes or, in this case, between each electrode and the




Figure 4.2 – *Electrostatic transduction demonstrated in a metallic plate. (left) Horizontal translation along the plate and (right) vertical translation above the plate.*

conducting body.

$$H(j2\pi\omega) = \frac{R_L \epsilon_r \epsilon_o A 2\pi f j}{d + R_L \epsilon_r \epsilon_o A 2\pi f j} \quad (4.1)$$

From inspection we expect the signal to decay at a rate on the order of $1/d$ with respect to vertical separation – the rate at which the effective source impedance increases. This was observed in the bottom right panel of figure 4.2.

4.5 Tank Attenuation Experiments

The plate experiment was followed by series of experiments taken in a controlled environment. A regular acryl glass fish tank (91cm x 38cm x 50cm) was filled with 114 liters of water. Each of the AA were mounted to a coaxial BNC bulkhead connector through a twisted pair feed line inside of an unsealed 2.5cm diameter PVC pipe. This structure allows ease of testing while minimizing RF leakage and ensuring complete dielectric submergence.

A gantry was mounted above the fish tank in order to precisely control the angle and position of one of the probes as shown in figure 4.3. The gantry can achieve and hold a linear position accuracy of $\pm 1mm$ due to an optical encoding strip along the bar (360dpi decoded in quadrature = 1440 counts per inch). For angular positioning it can achieve $\pm 2.5^\circ$ limited mostly by the accuracy of a servo motor. The gantry travel is from 25cm to 65cm in separation distance and a full 360° in angle.

4.5.1 SNR and BER

We conducted all experiments in three different mediums: fresh water, salt water of salinity 20ppt¹ (the world’s oceans lie approximately between 30ppt and 40ppt

¹ppt – Parts Per Thousand, by mass




Figure 4.3 – *Experiment configuration for signal loss over distance measurements.*

[Lab05]), and free air. The lack of free air propagation should imply that there is no direct transmission (RF leakage) between the transmitter and receiver. Figure 4.4 shows the measurements in the three mediums.

No communication is apparent in free air. The SNR is around the receiver sensitivity of 6.5dB, whereas in salt water, the average SNR is >12 dB, and in fresh water at around 11.5dB. Of note, the signal to noise ratio is decreasing with increasing distance in fresh water, whereas, in salt water it remains almost constant.

Figure 4.5 illustrates this further by plotting the Bit Error Rate (BER) for the measurements in salt and fresh water. It is self-evident that in fresh water, the BER clearly increases, whereas in salt water it remains constant. The standard model for electromagnetic RF propagation in salt water predicts a 10x reduction in transmission range for constant power with respect to fresh water [Rho07]. This was clearly not observed. In contrast, the electrostatic (transconduction) model predicts relatively little change in BER over distance given that the link remains

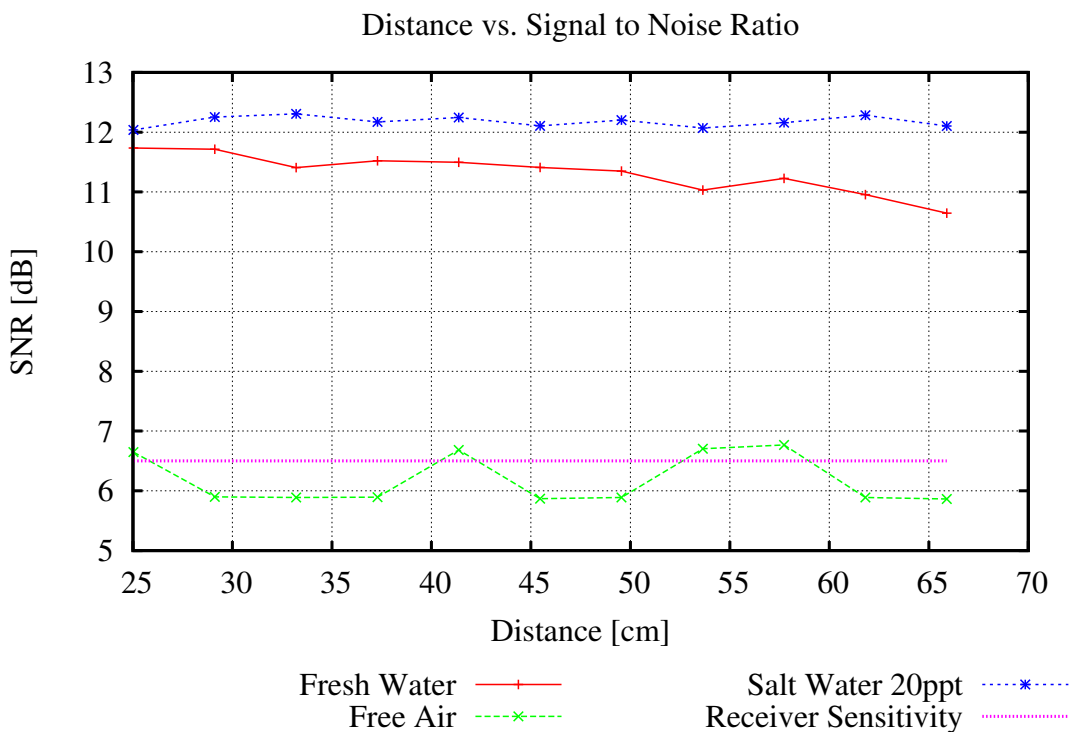


Figure 4.4 – *Distance vs. signal to noise ratio of a BPSK signal. We can observe that the signal is about 0.5 dB stronger in salt water of 20ppt salinity, as compared with fresh water. There is no propagation if the water is removed from the tank, e.g., the SNR stays at the level of the receiver noise floor. The transmitted signal has a 20mV peak-to-peak voltage swing.*

open, which agrees with our results.

4.5.2 Angular Dependency

The Radiation Power Density (RPD) distribution, more informally referred to as an “antenna pattern”, is another important indicator of propagation mode. Antenna aeriels produce characteristic RPD patterns that are principally a function of antenna geometry, *ceteris paribus*. Given our AA structure, that pattern should somewhat resemble the common linear dipole antenna if the dominant propagation mode is airborne electromagnetic waves.

The RPD pattern of an ideal generic dipole antenna lying along the Z axis and

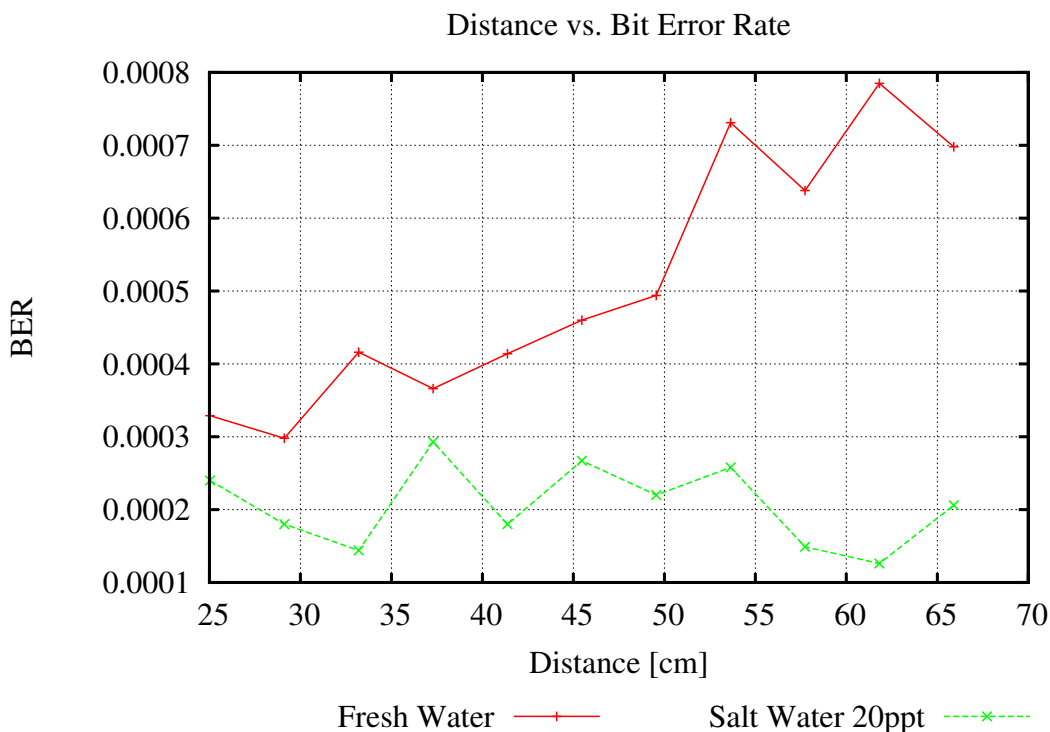


Figure 4.5 – *Variation in the SNR vs. distance is difficult to visualize, while the bit error rate (BER) clearly changes with increasing distance in fresh water and remains constant in salt water.*

centered at the origin is shown in figure 4.6. The radiation power density pattern of a linear dipole antenna is isotropic about the axis of the dipole, but experiences strong nulls along the collinear bearings. The plane of the antenna, as properly applied to our in-tank studies, is indicated by the green (shaded) plane. The 2D pattern evident on this plane is “figure eight” in form. We would expect to see this style of RPD in the AA tank study measurements – although that was not observed (figure 4.7).

Directed propagation is not unique to antenna geometry, the same is valid for acoustic hydrophones which always have an angular dependency in terms of signal strength. In our RPD investigation experiment, the two electrodes are held a fixed distance from each other, while the transmitting AA’s electrodes are rotated through a complete 360°. The two AA’s electrodes (two each, four total) are




Figure 4.6 – *The radiation power density pattern of a generic ideal linear dipole antenna is isotropic about the axis of the dipole (in this case, Z), but experiences strong nulls along the collinear bearings. The plane of the antenna, as properly applied to our in-tank studies, is indicated by the shaded plane.*

aligned and facing each other at the 0° position.

Figure 4.7 shows the signal to noise ratio of the electrode for each angle. We observe that there is a slight advantage when the probes are aligned, but nothing of remarkable significance. This finding is illuminating in that it neither agrees strongly with the multi-lobe electrostatic RPD expectation (figure 7.4), nor the electromagnetic one (figure 4.6).

4.6 Ocean Deployment

After the successful metal plate experiment and laboratory tank testing an expedition was undertaken to evaluate the artificial ampulla design in the Pacific

Angular Dependency of the SNR while Rotating the Transmitter

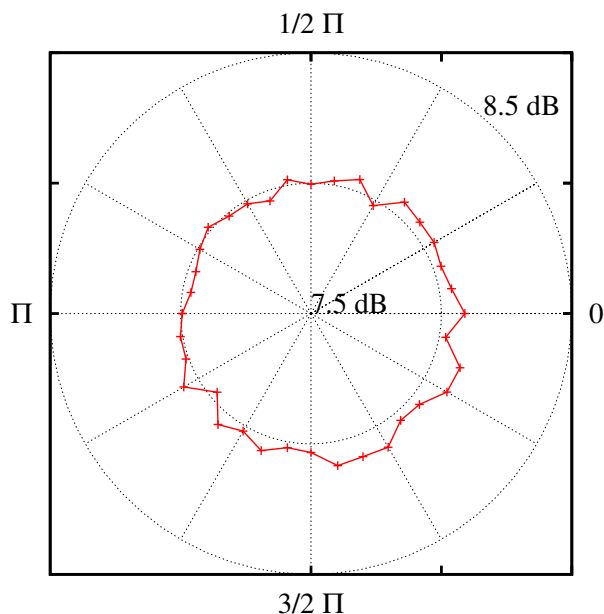


Figure 4.7 – *The angular dependence of the received signal to noise ratio given a fixed receiver and rotated transmitter. The transmitter and receiver probes are aligned at 0° . This test was performed in fresh water.*




Figure 4.8 – *Deployment test in Will Roger's State Park, California.*

Ocean. Two sites were chosen, one in Will Rogers State Park, near Santa Monica, California (figure 4.8), and one in Marina Del Rey, California (figure 4.10), for their proximity to infrastructure and vehicular accessibility.

For the duration of the expedition to Will Rogers state park (figure 4.8), we switched the modulation scheme to GMSK at a data rate of 250kbps, using a




Figure 4.9 – *Digital communication using GMSK at 8MHz and 250kbps data-rate at only 23mW of transmission power. Transmission with 100 bytes per packet (left). Transmission with 10 bytes per packet (right). A reduction in the packet length helps re-open the link.*

center frequency of 8MHz. In addition, we set the USRP to its full transmit power of 23mW. Figure 4.9 illustrates our findings.

Our ocean expeditions continue the disturbing trend of better-than-predicted performance, yet non-compliance to either linear radiator electromagnetic or electrostatic theories. In the ocean, we successfully transmitted digital data packets over distances of up to 5 meters using GMSK at 8MHz and 250kbps data-rate on only 23mW of transmission power with the AA’s submerged. The packet length was 100 bytes per packet (figure 4.9 (*left trace*)). Reducing the packet length helps re-open the link. Decreasing the packet length, given constant BER, increases the number of successful packet transmissions on average. At only 10 bytes per packet, we demonstrated successful packet transmissions out to 9

meters (figure 4.9 (*right trace*)). Communication range might be further extended given higher output power.

While the electrodes were completely submerged in these experiments, feed lines were used to return the signals to on-shore acquisition and processing equipment as visible in figure 4.8. In order to completely isolate and submerge the transmitting and receiving systems, we designed and implemented a waterproof computer system (figure 4.10) which encapsulated the prior experimental setup hardware (USRP, TXLF, RXLF, Battery, Hard disk, et. al.) replacing the computer with a single-board variant based around the Intel Atom dual-core 1.6GHz 330 model CPU.

The Artificial Ampulla (AA) was attached to this computer through a watertight bulkhead via a twisted-pair feed line to minimize RF leakage and ensure complete dielectric submergence. This feed line is visible in figure 4.10 just above the enclosure. A $20MHz$ wide chirp signal was transmitted from a completely submerged unit through its attached AA to an identical receiver AA located just below the surface and then fed out to a computer system positioned dockside (not submerged to allow for experimental control and monitoring). Figure 4.11 is a representative finding from the expedition with no detectable chirp signal above the natural background and self-induced noise floor.

4.7 Conclusion

The designs and experiments of this chapter have demonstrated the promise of submerged electrostatics, yet ultimately met with the dissatisfaction of the lack of reproduction once complete, isolated, submergence was attempted. This is not to discount the complete promise of the technology, rather that the approach itself must be reconsidered. In the next chapter, we endeavor to explain via modeling and follow-on experimentation what has occurred and how the design of the AA




Figure 4.10 – *Our next generation underwater Software Defined Radio System (WURI). It is completely submersible and while powered by one multi-cell LiPoly battery pack. The computer is built around a dual core 1.6GHz Intel Atom 330 CPU using the CISC x86 architecture.*

may be better adapted to serve its original purpose. In so doing, we will return to our biological model, the electroreceptive fish, and further seek its functional design.




Figure 4.11 – *Quasi-static electric field spectrogram taken in the Pacific Ocean in Marina Del Rey, California. Only noise is visible. The injected 20 MHz wide sinusoidal sweep signal is not detectable.*

CHAPTER 5

The Folded Unipolar Model

5.1 Overview

The preceding experimental findings offer mixed results. No clear conclusion emerges. At issue is the mode of propagation and the path (medium) in which it travels. From equation (2.17), the question may be posed somewhat more formally as a question of the relationship between $\epsilon\omega$ and σ . The two pathways (through the air, through the water) and the two modes (electrostatic conduction current, electromagnetic displacement current) form four possible hypotheses.

Airborne electrostatics is the least compelling of the four. The expectation for electrostatic energy transfer through low-density media from Coulomb's Law is abysmally low. The metal plate transconduction experiment (section 4.4) confirmed that there is basically no measurable transmission through non-permissible media (ex. air) and that the inverse is also true – that transmission improves dramatically as permittivity increases. The results from the tank experiments of section 4.5.1 offer further argument against the hypothesis of airborne electrostatics.

Submerged electrostatics is far more agreeable. From the preceding tank experiments it was observed that the signal strength is greater in fresh water than in air and greater still in salt water than in fresh water. That is in agreement with submerged electrostatics, which predicts increases with increasing conductivity/permittivity.

Submerged electromagnetics, like airborne electrostatics, is also a universally

poor fit. Its theoretical expectations, deliberated at length in section 2.4, are for huge submerged losses that increase with increasing salinity¹. The exact opposite trend was observed with respect to medium.

Attenuation with distance was also better than the quadratic decay expected by isotropic airborne or submerged electromagnetics. Given that submerged electrostatics predicts less decay per unit distance, it is a better fit. The electrode/antenna performance is quite good even though the emissive structure is planar, smaller than 2% of the wavelength, and fed from a twisted pair balanced line. Electromagnetic predictions for such a structure indicate minimal, in any, radiation under these conditions.

Unfortunately, submerged electrostatics does not fit all of the experimental data. The directivity experiments of section 4.5.2 show findings that are directionally invariant. While different patterns are expected, both submerged electromagnetics and submerged electrostatics expect directional behavior. Lastly, the lack of any observable signal during ocean deployments with WURI (section 4.6) raises obvious concerns.

Airborne electromagnetics seems a poor place to turn for an explanation since every experimental precaution was taken to intentionally eliminate EM propagation through the air. Indeed, the same arguments that may be made against submerged electromagnetics also rule out airborne electromagnetics. That said, airborne EM agrees with the observed non-directional behavior if the assumed radiator (antenna) is the linear structure feeding the artificial ampullae themselves.

5.2 Objectives

To provide a complete explanation of the observations, we must endeavor to explain (1) how the balanced twisted-pair feed line could, somehow, be able to radiate

¹Increased salt, sodium chloride, concentration is correlated with increased conductivity.




Figure 5.1 – (left) *Basic schematic representation of the Folded Unipole Antenna (FUA) showing its four principal lumped elements.* (right) *The bottom of an FUA in New Jersey, USA operating at 1.4MHz and 1kW output power [Haw11].*

energy into the air, (2) prove that propagation through multiple modes can and does exist simultaneously, and (3) show that such *mixed-mode propagation* occurs in significant ratio. The remainder of this chapter answers these questions in order.

5.3 The Folded Unipole Antenna

The question of how a balanced differential pair feed-line could possibly radiate is best answered by modeling it as a Folded Unipole Antenna (FUA). Broadly, antennas fit into two groups: linear and aperture. The FUA is in many ways the superset of all linear antennas. The FUA has been well enough understood to be practically employed for over 50 years, yet, has not been studied or optimized

sufficiently to enjoy widespread use [Rai07]. It is the integration of linear radiators with transmission lines and can take on near-limitless forms.

Electrically short monopole antennas have low resistance and high reactance. These antennas, therefore, require a capacitive matching network to tune out the natural reactance present in the short structure, but this matching network results in narrow bandwidth, relatively small aperture, and high losses. The cause of these losses is primarily the relatively high currents flowing at the feed-point.

The FUA reduces these losses because it reduces these currents. FUA's consist of no less than six independent design variables and compositing them into arrays offers five more. This creates an enormous design space and searching for optimal solutions using a field solver package is not viable (even with the best available computing power). In the course of this dissertation I developed a new antenna design tool, based on Raines' Method, specifically aimed at high speed analysis of FUAs.

The FUA consists of endless combinations and possibilities. Here, we restrict our investigation to the 1-segment folded unipole antenna – the so called, “simplest configuration”. The reasons for doing this are three fold: (a) It best highlights the capabilities of our analysis showing key concepts without involving large numbers of variables (as appear in the multi-segment cases). (b) It is sufficient to build advanced and novel antenna structures. (c) It allows for significantly faster analysis, in turn, allowing us to demonstrate a broader search of the design space.

The FUA in its simplest configuration is shown in figure 5.1. It consists of an excited linear element, optional top termination impedance (Z_T), a non-excited return element, optional base impedance (Z_B), and a ground plane. The U-shaped pattern in figure 5.1 is illustrative.

5.3.1 Numerical Analysis of FUEs

FUEs are difficult to simulate using pure numerical methods (such as the Method of current Moments “MoM” - which involves dividing arbitrary 3D structures into tetrahedra composed from 6 wires and looking at the field contribution from currents on each of the wires). As noted, FUEs possess a number of exceptionally useful qualities including the ability to perform efficiently, well below $1/4$ wavelength (where linear antennas are typically ideally sized). FUEs combine attributes from antenna design and transmission line theory to create a complex structure that has both strong differential (transmission line, balanced) and common-mode (radiative) currents. Because neither component vanishes in the structure, the numerical analysis must accurately follow and determine current flow in both domains. Since differential current elements are under-determined at wire junctions MoM adds the Kerchkoff Voltage Law constraint for each junction. For even simple top-loaded FUEs this can produce matrices with thousands of elements. When this happens, the probability for an ill-conditioned matrix improves dramatically – as is typically reported in the literature [Rai07].

Ill-conditioned matrices occur when two constraining equations are only barely independent of each other. For example, consider:

$$x + y = 1 \quad \text{and} \quad (x + u) + y(1 - e) = 1 \quad (5.1)$$

for u, e many orders of magnitude smaller than x, y . The solution for this system is:

$$x = 1 - \frac{u}{e} \quad \text{and} \quad y = \frac{u}{e} \quad (5.2)$$

If u and e represented the uncertainty in a current measurement by an instrument or a stray current traveling transverse to the primary common-mode current wavefront, it is easy to see how extremely sensitive the resulting solution equations

are - this is called a numerically noisy solution and it has low accuracy confidence. This difficulty analyzing and optimizing FUA structures has prevented their wider adoption despite tremendously attractive size and efficiency characteristics.

5.3.2 History of the Folded Unipole Antenna

Although Roberts in 1947 [Rob47] was the first to describe experimental observations of the FUA he could not provide a complete mathematical model to explain them. Leonhard, Mattuck, and Pot [LMP55] were the first to explain the FUA's remarkable performance by attributing it to a combination of two factors: (a) the division of the radiating current between the driven and undriven portion of the antenna and (b) the flow of current in a non-radiating transmission line mode. However, their analysis could not explain FUA's longer than $1/4$ wavelength or with top-termination. King and Harrison in 1960 [JK61] provided the first coherent model for the simple FUA configuration with both top and base impedances but their model was integral in nature and could not be solved analytically for arbitrary cases. It was not until recently that a clear analytical model for the arbitrary FUA has emerged. Jeremy Raines seminal work [Rai07] proposes a boundary condition approach which is analytically based on King yet resolves the complex integration inherent in King's model. Raines' model opens the door to rapid analysis of the design space allowing greedy searches over many variables in a matter of hours.

5.3.3 The Raines Model of the FUA

The Raines model of the FUA consists of four equations – the result of evaluating common and differential modes at both the top and bottom (base) of the antenna.

5.3.3.1 The Top of the FUA

At the top of the FUA (figure 5.1) the current in the common-mode must vanish as it theoretically has no where to go:

$$I_c(h) = 0 \tag{5.3}$$

In practice, however, the non-infinitesimal diameter of the antenna wires requires that some capacitance exists between the antenna tip ($z = h$ for $h = \text{height of the unipole}$) and the ground plane which is capable of storing charge. This can slightly invalidate equation (5.3) in which $I_c(z)$ is the common-mode current's initial phase as a function of height. This is known as the antenna end-effect, but in this analysis we ignore it sacrificing a comparatively tiny amount of accuracy for a large increase in performance and the ability to ignore the precise geometry of the antenna and instead focus on lumped parameters.

The differential mode current in figure 5.1 proceeds from the signal generator V_s upward through the driven element across Z_T and down through the undriven element. Consequently the differential voltage, $V_d(z)$, as a function of its distance along the height axis, z , should be defined as $V_{driven} > V_{undriven}$. Applying Ohm's law to this differential mode at the top of the unipole yields:

$$V_d(h) = I_d(h) \cdot Z_T \tag{5.4}$$

5.3.3.2 The Base of the FUA

At the base of the folded unipole, the common mode voltage is just the signal generator:

$$V_c(0) = V_s \tag{5.5}$$

To derive a differential mode expression for the base of the unipole we apply Kirchhoff's Voltage Law around the loop beginning at the negative terminal of the signal generator:

$$-V_s + Z_B [I_d(0) - I_c(0)] + V_d(0) = 0 \quad (5.6)$$

5.3.3.3 The Determination of the Current Moments

Excitation of a folded unipole sets up a series of travelling waves akin to a stretched string fixed at only one end. In the common-mode this point of fixation is at the top of the antenna. In the differential-mode it is at the ground plane. In either case the wave travels forward until reaching the point of fixation and then reflects travelling in reverse until reaching the originating node. Standing waves may setup when the string is fixed at both ends. This is the case with series-fed monopoles and dipoles. Unlike standing waves which may reflect multiple times the traveling wave is back in phase after its initial reflection. Using these phases, we can define the common (equation (5.7)) and differential (equation (5.8)) currents, with respect to height as:

$$I_c(z) = I_{cF} e^{-jkz} + I_{cR} e^{jkz} \quad (5.7)$$

$$I_d(z) = I_{dF} e^{-jkz} + I_{dR} e^{jkz} \quad (5.8)$$

This yields the common and differential mode voltage equations:

$$V_c(z) = Z_c \cdot I_c(z) = Z_c (I_{cF} e^{-jkz} + I_{cR} e^{jkz}) \quad (5.9)$$

$$V_d(z) = Z_d \cdot I_d(z) = Z_d (I_{dF} e^{-jkz} + I_{dR} e^{jkz}) \quad (5.10)$$

These equations can be combined to give the matrix expression:

$$\begin{bmatrix} I_{c_F} \\ I_{c_R} \\ I_{d_F} \\ I_{d_R} \end{bmatrix} = M^{-1} \cdot \begin{bmatrix} V_s \\ V_s \\ 0 \\ 0 \end{bmatrix} \quad (5.11)$$

where M is defined as:

$$M = \begin{bmatrix} -Z_{base} & -Z_{base} & Z_{base} + Z_{diff} & Z_{base} - Z_{diff} \\ Z_{comm} & -Z_{comm} & 0 & 0 \\ e^{-jkh} & e^{jkh} & 0 & 0 \\ 0 & 0 & (Z_{diff} - Z_{top}) \cdot e^{-jkh} & (-Z_{diff} - Z_{top}) \cdot e^{jkh} \end{bmatrix} \quad (5.12)$$

5.3.4 The Electromagnetic Field

Using the model of equation (5.11), the closed form solution for a single-segment folded unipole (1SFU) above a perfect, electrically conducting, infinite ground plane is:

$$E_\theta = \frac{(1 + \nu)\eta k \sin(\theta)}{4\pi R} e^{-jkR} \frac{1}{jk} \left(\begin{array}{l} \frac{-I_{fc} e^{-jkh(1-\cos(\theta))} - 1}{1 - \cos(\theta)} \\ + \frac{I_{rc} e^{jkh(1+\cos(\theta))} - 1}{1 + \cos(\theta)} \\ + \frac{I_{rc} e^{jkh(1-\cos(\theta))} - 1}{1 - \cos(\theta)} \\ + \frac{-I_{fc} e^{-jkh(1+\cos(\theta))} - 1}{1 + \cos(\theta)} \end{array} \right) \quad (5.13)$$

5.3.5 Power and Input Impedance

From the field, the power may be computed as:

$$P = 2\pi \int_0^{\frac{\pi}{2}} \frac{R^2 \sin \theta |E_\theta|^2}{\eta} d\theta \quad (5.14)$$

and the current at the feed point is:

$$I_{in} = \nu (I_{fc} + I_{rc}) + (I_{fd} + I_{rd}) \quad (5.15)$$

The ratio of power to input current is the input resistance:

$$R_{in} = \frac{P}{I_{in}} = \frac{P}{|\nu (I_{fc} + I_{rc}) + (I_{fd} + I_{rd})|^2} \quad (5.16)$$

The ratio of the drive voltage to the input current is the input reactance:

$$X_{in} = \frac{V}{\nu (I_{fc} + I_{rc}) + (I_{fd} + I_{rd})} \quad (5.17)$$

5.3.6 FUA Analysis

With the advent of equations (5.16) and (5.17), we can now perform Raines analysis on FUA structures, provided we have the five requisite parameters to describe them:

1. Frequency – easily determined as it is under experimental control.
2. Common-mode impedance of the folded wire pair (Z_C) – determined from the cross-sectional geometry or by measurement.
3. Differential-mode impedance of the folded wire pair (Z_D) – also determined from the cross-sectional geometry or by measurement.
4. Top terminating impedance (Z_{top}) – this may be calculated or measured for the device explicitly added to the structure to serve in the role of top termination. As it is provided by design, it usually has a target value. It will be demonstrated shortly that, in the case of our experiments, the ocean in concert with the artificial ampullae served in this role.

5. Base terminating impedance (Z_{base}) – similar to the top termination, the base termination is typically a specific designed-in value achieved through the addition of discrete elements such as a capacitors or by structures such as shunt transmission lines or environmental factors effectively serving in that capacity.

5.4 Applicability to our Ocean Studies

Working clockwise around the schematic of figure 5.1, we may describe how the FUA model is mapped to our experimental configuration. It should be noted that the experimental configuration is rotated 180° from the figure, its feed point being generally considered to enter from above, given that the other end of the cable reaches the Artificial Ampulla (AA) which may be submerged in the ocean. It may be simpler to perform the rotation the other way and imagine the ocean as being at the top of figure 5.1. In either case, the signal generator, being connected to the driven arm of the AA's feedline, forms the first leg of the FUA. The impedance of the AA itself provides the top termination and the undriven arm of the feedline forms the return leg of the unipole. No base termination is applied ($Z_{base} = 0 + 0j$, it is a perfect short circuit), consistent with the feedline's direct ground connection to the frame of the signal generator.

Given that the geometry of the feedline readily yields values for Z_c and Z_d , Z_{base} is shorted, and a typical current transformation ratio value of $1/2$ is reasonable, the only missing value necessary for a Raines' analysis of the FUA model is the top termination. Consequently, the impedance of the Artificial Ampullae must be determined.

5.4.1 The Impedance of the Artificial Ampullae

In recognizing that, mechanically, the Artificial Ampulla (AA) is a large flat structure with little inductance to speak of, it may be well modeled as a parallel plate capacitor whose capacitance, C_{AA} , may be estimated by:

$$C_{AA} = \frac{\epsilon_r \epsilon_0 A}{d} \quad (5.18)$$

where $\epsilon_0 \approx 8.854 \times 10^{-12} \text{ F/m}$, ϵ_r is the relative permittivity of the medium, A is the effective overlap among the plates, and d is the separation between them.

Air ($\epsilon_r = 1$) is a poor conductor and does not carry the field lines effectively. Consequently, $A \approx 0.9\text{mm}^2$ and $d = 25.4\text{mm}$ as only the closest interior edges of the electrodes patterned on the fiberglass body of the AA play much of a role. The DC resistance of the AA is practically infinite and, solving equation (5.18), the capacitance is quite negligible ($C_{AA} < 0.35 \times 10^{-15}\text{F}$).

The behavior changes dramatically when the Artificial Ampulla (AA) are submerged. The dense conductive medium provided by the oceanic waters ($\epsilon_r \approx 81$) readily extends the field lines to the entire electrode area. Accordingly, A improves to $\approx 6.5 \text{ mm}^2$ and d reduces to the AA's ocean insulation thickness ($\approx 0.5\text{mm}$). Re-solving, $C_{AA} \approx 1 \text{ nF}$.

The impedance of the AA electrodes, Z_C , when modeled as a capacitor,

$$Z_C = \frac{1}{2\pi j f C_{AA}} \quad (5.19)$$

results in a reactance of $-22j\Omega$ at our test frequency of $f = 8 \text{ MHz}$.

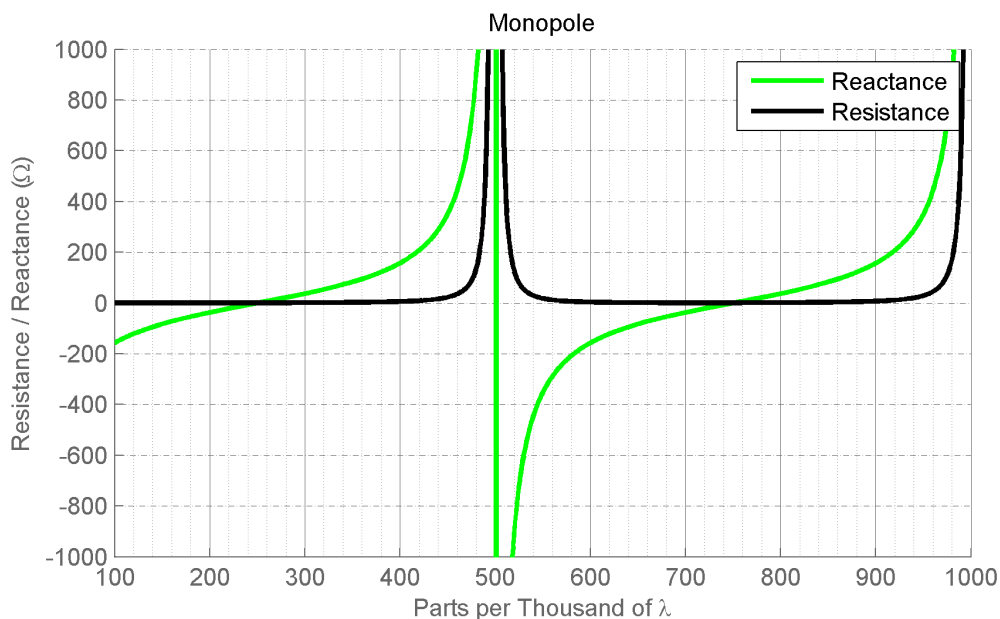


Figure 5.2 – *The Raines approach to folded unipole antenna analysis applied to the coaxial feedline connecting the Artificial Ampulla (AA) to the Universal Software Radio Peripheral (USR) when the AA are not submerged. Note the incredibly poor antenna behavior. This structure will not radiate efficiently.*

5.4.2 The Effect of Submergence

When the AA are not submerged, they offer little capacitance and almost infinite resistance (a near-perfect open-circuit). Using $Z_{top} = 10^{10}$, we get Figure 5.2 – the results from Raines’ analysis of the FUA model of the AA when not submerged.

The real (resistance) and imaginary (reactance) parts of the impedance are plotted separately. The structure is resonant, and therefore efficiently accepting of incident energy, when the reactance is zero. Reactive structures respond to input power by storing energy locally in fields, rather than letting that energy result in charge movement and, therefore, radiation. It is interesting to note that there are no values of the impedance for which radiation is efficiently achieved. When the reactance in figure 5.2 is zero the resistance is either extremely high (near-infinite) or extremely low (virtually zero).

When the resistance is extremely high, it is very difficult to inject current

into the antenna/electrode structure due to Ohm's law. The voltage required to overcome the extremely high resistance and develop a large current is also extremely high. Large voltages are expensive, complicated, and inefficient to generate – if a high enough voltage could possibly be generated at all.

On the other hand, when the resistance is extremely low the structure appears like a short circuit. It is difficult to inject large time-varying currents into the structure because the generating power supply and signal amplifier output are essentially clamped to ground and unable to oscillate.

When the AA are submerged in the ocean, the entire situation changes – *even if the AA are barely submerged*. The coaxial feedline connecting the AA to the USRP is now top-terminated by a much larger capacitive component of the impedance as explained in section 5.4.1. Now, the ocean and AA act collectively as a capacitive reactance.

The results of applying the Raines approach for folded unipole antenna analysis *with the AA submerged* in the ocean, appears in figure 5.3. The zero points in the reactance have shifted to much smaller fractions of the wavelength, repeat with greater frequency, and occur simultaneously with usable values of resistance. The resonances no longer occur at near-zero or infinite resistance. Not only will the retuned structure radiate very efficiently at several points, but the probability of operating at, or near, one of those points has increased almost 66% over the air-terminated monopole.

5.4.3 An Explanation for Prior Discrepancy

We are now ready to provide an explanation for the concerns raised at the start of this chapter. The angular invariability may be explained by the fact that the energy is really travelling from one AA feed line to the other, where each feed line is an ocean-terminated folded unipole antenna. Because the feedlines are not

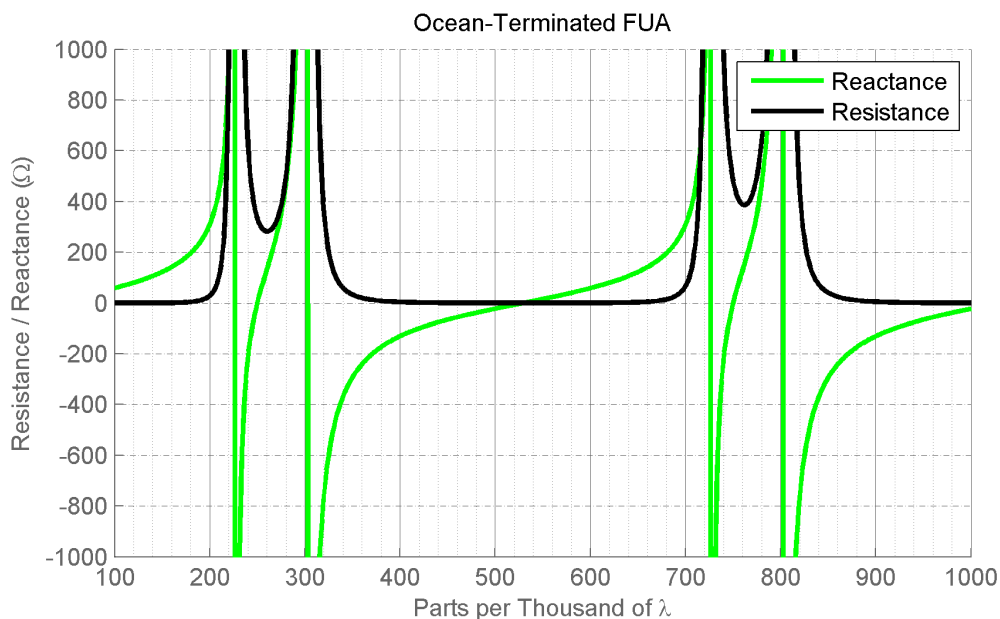


Figure 5.3 – *The Raines approach to folded unipole antenna analysis applied to the coaxial feedline connecting the Artificial Ampulla (AA) to the Universal Software Radio Peripheral (USR) when the AA are submerged in the ocean. The ocean and AA act collectively as a capacitive top-termination. The retuned structure will radiate very efficiently at several points.*

appreciably moving relative to each other, and coaxial FUA's are not directional with respect to azimuth, the rotational pattern (figure 4.7) shows a near constant output.

When the feedlines are substantially shorter, as was the case with WURI (figure 4.10), the number of resonance length possibilities is also dramatically reduced. As WURI was submerged, the Z_c and Z_d terms of the feedline-as-FUA structure were altered (reducing the radiation resistance and, thus, the effectiveness as an antenna), while the RF absorbing nature of the ocean waters eliminates virtually all of the output power as SNR climbs too high to detect.

5.4.4 A Note of Caution Regarding FUA Analysis

It should be noted that there remains some points of disagreement between the work of this author and that of the author of [Rai07], Dr. Jeremy Raines. For illustration, consider:

1. A one segment folded unipole (*the simplest configuration*)
2. Swept over height (holding all other parameters constant)
3. Evaluated at the first resonance occurring lower in height than $1/4$ wavelength

| [Rai07]'s Fig.# | Raines | Friedman |
|-----------------|--------|----------|
| 2-8 | 222 ft | 215 ft |
| 2-9 | 215 ft | 201 ft |
| 2-10 | 210 ft | 152 ft |

Table 5.1 – *Comparison of Raines and my models for the resonance point below $1/4$ wavelength*

Table 5.1 summarizes the discrepancy. My implementation of the Raines model produces a model with too much sensitivity to the current transformation ratio. For clarity, lets look exclusively at the input current.

The magnitude of the input current is at a local minimum when the sub- $1/4$ wave resonance occurs. In figure 5.4 this is shown. If the conditions are altered to the conditions of figure 2-10 in Raines' text (current transformation ratio= 5, $Z_c = 1200$, $Z_d = 120$, $f = 984kHz$, $Z_t = Z_b = 0$), figure 5.5 results. The resonance has shifted to 152 feet, which is, again, inconsistent with [Rai07].

Input current is the first term derived in the analysis and, therefore, foreshadows trouble to come. My implementation solves for I_{in} by working equation 5.12 through equation 5.11. The input current is merely

$$I_{in} = \nu (I_{fc} + I_{rc}) + (I_{fd} + I_{rd}) \quad (5.20)$$

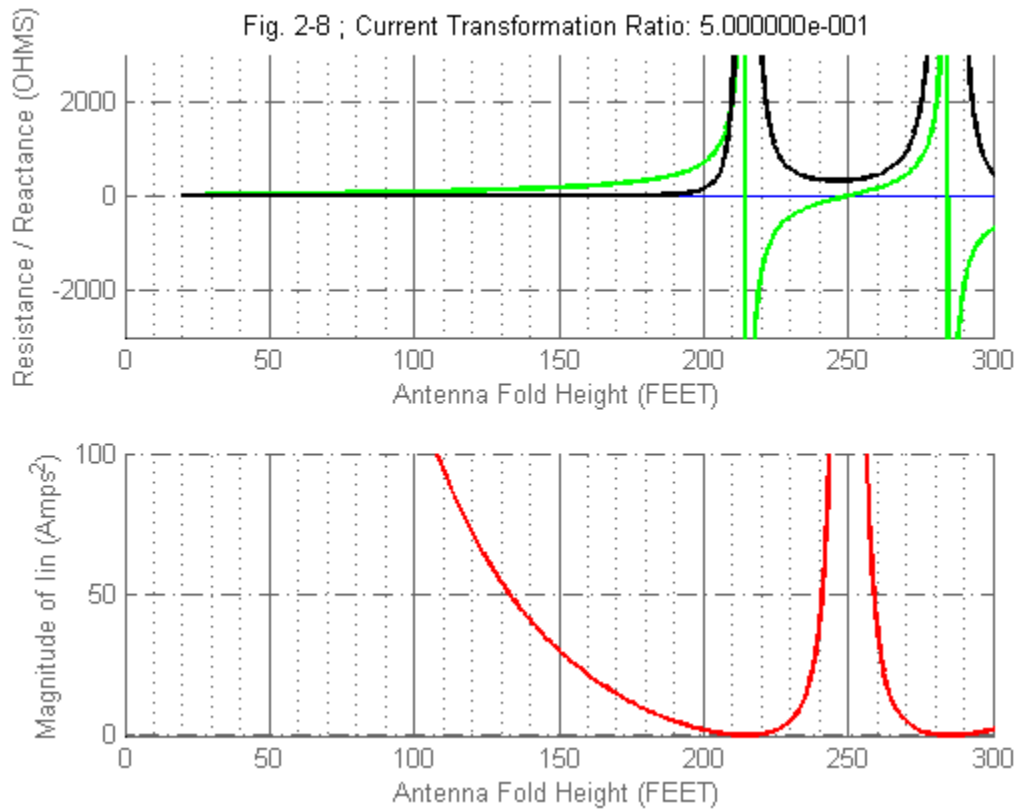


Figure 5.4 – *Sweep of Impedance and Input Current vs. height. This is my recreation of Fig. 2-8 from [Rai07], which shows some disagreement with the results presented here.*

given that the antenna is a one-segment folded unipole.

Correspondence with Dr. Raines over many weeks has confirmed correctness in general approach and verisimilitude in many of the equations. At the time of printing, the precise source of the variance had yet to be located. With that cautionary note aside, the exact behavior of the structure is largely immaterial. The overarching conclusion is that the folded unipole antenna model is apt in explaining observed results and proffers a model of great utility for resuming our quest for biomimetic electrostatics. Those improvements are discussed in detail in the coming chapters.

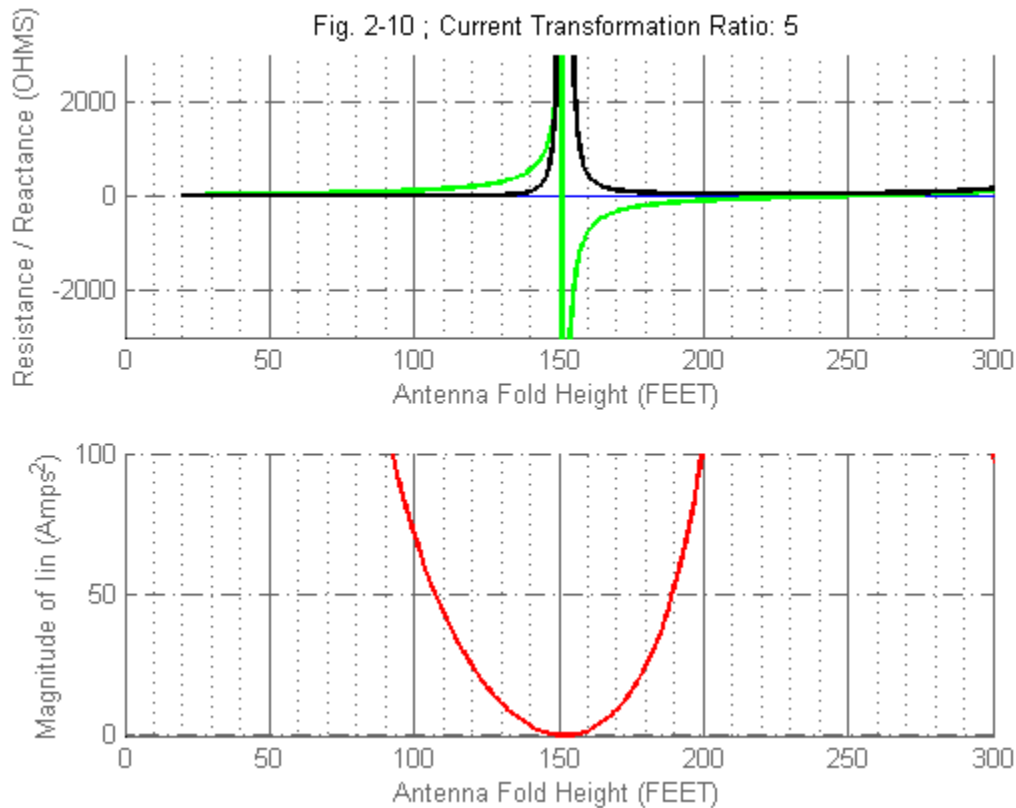


Figure 5.5 – Sweep of Impedance and Input Current vs. height. This is my recreation of Fig. 2-10 from [Rai07], which shows some disagreement with the results presented here.

5.5 Mixed-Mode Propagation

For the Artificial Ampulae (AA), as originally designed, airborne electromagnetics clearly dominates in some experiments. While in others, results unexplainable by airborne electromagnetics were observed. The interaction of the modes of propagation deserve explicit attention. In so doing, the last two questions posed in the introduction to this chapter will be addressed. Namely, it will be shown that both static and dynamic means of propagation contribute to the total energy transmitted between our AA, that the distribution of those contributions follows expectation, and that the AA design may be optimized to improved submerged performance by reducing the output impedance.




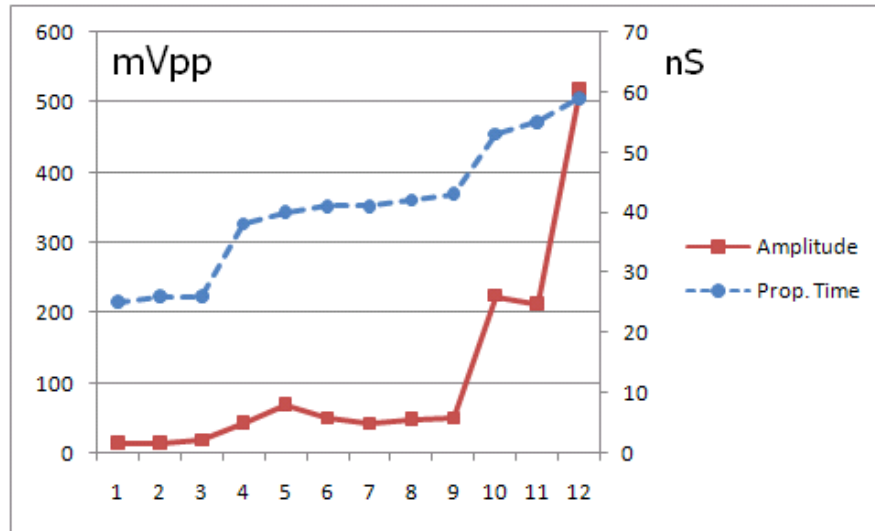
Figure 5.6 – *Experimental configurations for investigating mixed-mode propagation through water and air*

5.5.1 Experimental Configuration

For the experiment, three tank configurations are considered, as indicated diagrammatically in figure 5.6. In configuration (b) of the figure, each of the probes are either submerged in, or surfaced above, the *same tank* and separated by a distance of one meter. In configuration (a), the probes are introduced to *separate adjacent tanks*, while still separated by one meter. In (c), the two *tanks are moved apart* until the separation distance has increased to two meters.

For each tank configuration the electrodes of the transmitter may be either submerged or held slightly above the surface, the same is true of the receiver. This leads to four possible combinations. In total, there are three tank configurations and four electrode configurations – resulting in the twelve experiment configurations of table 5.2.

Each of the tanks was filled with fresh (tap) water for the duration of these trials. The attenuation of electromagnetic radiation is much lower in fresh water



| Point | Dist | Tank | TX | RX | Tau (ns) | Vpp (mV) |
|-------|------|------|-----|-----|----------|----------|
| 1 | 1M | Diff | Out | Out | 25 | 14 |
| 2 | 1M | Same | Out | Out | 26 | 14 |
| 3 | 2M | Diff | Out | Out | 26 | 18 |
| 4 | 1M | Diff | In | Out | 38 | 43 |
| 5 | 1M | Same | In | Out | 40 | 69 |
| 6 | 1M | Diff | Out | In | 41 | 50 |
| 7 | 2M | Diff | In | Out | 41 | 42 |
| 8 | 2M | Diff | Out | In | 42 | 48 |
| 9 | 1M | Same | Out | In | 43 | 50 |
| 10 | 1M | Diff | In | In | 53 | 224 |
| 11 | 2M | Diff | In | In | 55 | 213 |
| 12 | 1M | Same | In | In | 59 | 518 |

Table 5.2 – *Propagation through air and water*

than sea water, but fresh water has a similar permittivity. Permittivity describes the readiness with which a material electrically polarizes – e.g. the material’s ability to transmit (or “permit”) an electric field [HRW97]. Water, with a relative permittivity near 80, is among the most permeable materials in common occurrence [SK73] [ASS04]. The permittivity of air is extremely low, suggesting that, at the surface interface, refraction from an underwater-to-surface EM wave would bend severely (close to a right angle) and propagate above (and almost parallel to) the water [Rho07]. When above the water, the propagation velocity of an EM wave in air is close to the speed of light, while in water (although frequency dependent) it

is on the order of 100 times slower [ASS04].

For each of the combinations: \langle *Same tank/Different tanks*, *Separation distance*, *Transmitter submergence*, *Receiver submergence* \rangle , the signal strength and propagation time were recorded. It is important to note that the entire configuration including cabling, connectors, and instrumentation was undisturbed between trials and that all of the reported data of figure 5.2 was recorded continuously in one session to minimize the introduction of parasitic variance and relative error among data points. A LeCroy Waverunner 64Xi digital sampling oscilloscope running at 10 billion samples per second was used to digitize the receiver signal and an Agilent 33220A arbitrary waveform generator running at 5MHz and 4Vpp was used as the transmitter. The transmitted signal consisted of a single sinc function pulse transmitted on a manual trigger.

5.5.2 Experimental Findings

The results of the experimental trials appear in table 5.2 where they are sorted by propagation time. The “Point” field refers to the x-axis location and implies the configuration used for that trial.

The propagation times group into three categories: (1) The fastest group, which propagates through the air (Points 1-3), (2) a middle group, which propagates through both the air and the water, and (3) a slowest group, which propagates almost exclusively through the water.

The amplitude and propagation time are correlated, which is encouraging. The feedline, even with top-termination (AA submergence), while efficient enough to radiate, is still a poor antenna when compared with an intentional radiator. As the AA were designed to effectively mimic the electrostatic means of their biological archetype, they are substantially better at transmission in this mode. This is reflected in the dataset with submerged electrodes outperforming airborne ones

consistently.

When neither AA is submerged, the only relevant factor is the separation distance. Comparing points 1 and 2, practically no difference is detected. This is logical given that the airborne signal is not dependent on subtle material variation beneath it. The electrodes are held above the water's surface. Whether that surface is composed of one tank or divided into two adjacent tanks makes little difference – the signal still propagates above a consistent homogeneous aqueous medium. Increasing the distance also has little effect as airborne propagation occurs at the speed of light in air². Comparing points 2 and 3 almost no measurable difference in the propagation time is detectable.

There is a clear performance benefit for submerged electrodes. The strongest recorded signal occurred when both electrodes were submerged in the same tank (point 12). If the distance between the electrodes is held constant, but the electrodes are moved into two separate tanks, the signal received decreases dramatically – $\approx 59\%$ between points 11 and 12.

Three general trends with respect to signal strength are clear: (1) submerged performance always exceeds dry performance, (2) more distance results in more attenuation, and (3) electrostatic propagation results in less attenuation than electromagnetic propagation over short tank distances.

5.6 Conclusion

As demonstrated in this chapter, the originally designed artificial ampulla were a good imitation of the electroreceptive organ (receive), but were a poor emulation of the electric organ (transmit). That is to say, while our input impedance was sufficiently high, the output impedance was not sufficiently low. Due to the Ohmic

²The refractive index of air, for visible light, is about 1.0003, so the speed of light in air is about 90 km/s slower than c , the speed of light in a vacuum [HRW97].

nature of the conduction current and the naturally low resistance of the ocean water, establishing a large I_{out} implies providing a large V_{out} , which in turn requires a very small source output impedance, R_s (so named, because it is just resistance in the case of conduction currents). In terms of the FUA model, we must endeavour to eliminate the reactive component of the top-termination, which will reduce the airborne electromagnetic component and simultaneously increase the conduction current responsible for the electrostatic component.

CHAPTER 6

Electrostatic Transceiver Design

6.1 Overview

To incorporate the design revisions suggested in the previous chapter, all manner of insulation was removed and the electrodes were exposed to the ocean simulant. Despite contact, the input impedance remained unchanged as it was instead provided by the voltage measurement equipment. In addition to the reduction in output impedance, the receiving electrodes were reduced in size from those of figure 4.1 to those of figure 9.2 to better approximate the point detectors of the reference adult male Oman shark. This will permit imaging, in addition to communication and other goals.

Having established confidence in our electrode configuration through the work of prior chapters, this chapter covers the design, development, optimization, and in-ocean validation of the individual channel element that is replicated to form our 16-channel and larger Biomimetic Electrostatic Imaging (BEI) system arrays. The contributions include design considerations for in-ocean operation of BEI devices with specific attention paid to the attendant chemical and physical issues, design and analysis of an optimized BEI analog front end, construction and validation of a channel prototype, and the highest-resolution electrostatic field recordings published to date.


6.2 Design Considerations

The design of each BEI channel is driven by material, environmental, and electrical design considerations.

6.2.1 Material Considerations

The material design of the electrodes is critical as it establishes a practical upper bound on achievable performance. Prior work with steel electrodes [FTC10] revealed a number of operational problems due to its iron content [FTS10]. Reactions with the ocean water can result in electrical noise (unstable capacitance, long-term ionization, etc), sensor degradation (ex. oxidation, producing a substantial increase in resistance at the interface), and even device failure (the electrode dissolving in its entirety).

In the work of this chapter, the steel content of [FTC10] was replaced with 99.9% pure platinum. The reduction potential for platinum is much lower than that for iron and, thus, can endure a much higher voltage without reacting with the salt water. Among all of the oxidizing agents in the solution, oxygen has the highest reduction potential, therefore oxidation carried out by oxygen is most likely. The whole redox reaction is equation 6.1:




To make a quantitative determination, we apply the Nernst equation:

$$E = E^\circ - \frac{RT}{nF} \ln Q \quad (6.2)$$

where E is the non standard reaction potential, E° is the standard reaction potential, R is the gas constant, T is the absolute temperature, F is Faraday's

constant, n is the number of moles of electrons transferred per mole of the reaction, and Q is the reaction quotient. The reaction's potential depends on both the chemical properties and concentrations of the species involved. In ocean water, where the concentration of Pt is less than 4×10^{-7} ppm [Tur68], the modified reaction potential, from equation (6.2), is: $-0.080V$.

Given that $E < 0$, the reaction between oxygen, water, and platinum is not spontaneous. Conversely, the oxygen, water, and iron reaction in sea-water is quite spontaneous. The complete iron reaction is (6.3) and substituting $E^o = 0.84V$ and $n = 4$ into (6.2) yields $E = 1.41V$. E is greater than zero, validating the conclusion.



The magnitude of the modified reaction potential sets a conservative upper bound for the strength of the applied excitation field for our imaging system¹. Positive reaction potentials represent spontaneous reactions. Driving the imaging system with waveform amplitudes in excess of the absolute value of E results in a positive overall value of E : the applied external field is sufficient to overcome the inherent non-spontaneity of the reaction. The experiments of section 6.2.3, with an amplitude of 50mV (e.g. $< E$), fall squarely in the intrinsically safe category.

When more substantial drive strength is required (the experiments of section 7.1.2 used an amplitude of 4.80V), careful consideration of the drive parameters can substantially mitigate the likelihood of significant electro-chemical involvement and degradation of the electrodes. The overpotential of the platinum reaction in sea-water is significant and may add as much as several hundred millivolts to the safe working range.

Given that the attendant reactions are based on mass transfer (platinum

¹Over-potential and other modifiers have not been expressly incorporated into our calculations, but their net effect serves only to increase the magnitude of the modified reaction potential adding to the safety margin of our original conclusion.

atoms leaving the electrode), the reaction requires excitation beyond the order of milliseconds. In all of our experimentation the drive waveform was sinusoidal with a comparatively high frequency ($f \approx 1kHz$). Even in the case of section 7.1.2, the amplitude of the drive waveform exceeds E for no more than $495\mu s$ without rest in any contiguous timespan.

The forward reaction of platinum with oxygen and water is not spontaneous, which implies that the reverse reaction *is* spontaneous under the same conditions. To the extent that we forward bias the reaction through excessive excitation in the positive drive phase, the negative drive phase will help the inherently spontaneous reverse reaction to occur resulting in minimal long-term change to our electrode.

The electrodes are cylindrical in shape and made as small as possible to best emulate a point source in the water. For practical reasons of handling, durability, and ease of construction, the electrodes were fashioned from 30 gauge (AWG) 99.9% pure platinum wire with final assembled dimensions of $267\mu M$ in diameter and $3600\mu M$ in length. The top surface was sealed to the feed system resulting in an electrode working surface area of $3mm^2$.

6.2.2 Salinity and the Electric Field

Biomimetic Electrostatic Imaging begins with creating an electric field in the salty aqueous environment of the ocean. On average, the ocean is 99.3% water (H_2O) by mass with the remainder composed mostly² of salts – chief among them Sodium Chloride ($NaCl$). The actual salinity of the ocean varies with position, depth, temperature, and season over a range of 28ppt³ to 35ppt [Mar07] in open ocean and as high as an average of 39ppt in the Mediterranean and Persian Gulf [Lab05].

To establish and verify oceanic verisimilitude, the experiments of section 6.2 were conducted in a 55-gallon (208 liter) capacity tank filled with 50 gallons (189

²The balance is composed of trace amounts of minerals and other metals[Tur68]

³Parts Per Thousand (ppt) by mass ratio with respect to water.




Figure 6.1 – *The salt transfer function. Experimental (blue, solid) and ideal (red, dashed). Ideally, a straight line, the good linearity indicates good relative instrumentation accuracy. The more positive slope of the experimental trace indicates a deficit in absolute accuracy.*

liters) of drinking water (0ppt salinity). Salt in the form of Sodium Chloride was added in measured amounts to vary the salinity of the test vessel. A Marineland Labs, inc. Instant Ocean Hydrometer [Mar07] was used to record the salinity and the dry-weight labeling of the salt storage containers was used to estimate the mass of the sodium-chloride added. The results appear in figure 6.1, where the dashed red line indicates the expected result.

The dataset traces the salt transfer function, that is, the relationship between salt input (added to the tank) and salt output (dissolved into the tank solution). Accordingly, the more positive slope of the experimental curve suggests that our system has somehow provided an amplification function (more mass appearing in




Figure 6.2 – *The electrical resistance of the water path between the transmitter electrodes as their separation distance is increased. Resistance varies almost linearly with separation and exponentially with decreasing salinity.*

solution than we actually added). Since this cannot physically occur, the error can only be attributed to inaccuracy on the part of our salinity measuring equipment.

6.2.3 Conductivity

As the salinity of water increases its conductivity increases, as is evident in figure 6.2. This data was captured by employing an Agilent 33220A Arbitrary Waveform Generator (AWG) configured to output a 1.000kHz, 1.000Vpp, sine waveform through a 50.0Ω source resistance. The driven waveform was observed with a LeCroy Waverunner 64Xi Oscilloscope connected to the electrode feed lines near their point of contact with the tank’s water using a 10:1 passive probe with 600MHz

of analog bandwidth and with the instrument's 20MHz Low-Pass Filter (LPF) enabled. The two transmit electrodes, as described above, were manually held in place against a submerged ruler 2-4cm below the surface – the two electrodes presenting the driven and reference ends of the sine wave signal to the tank's water solution respectively.

While the dataset begins with fresh water (0ppt), it is important to note that the actual resistance of fresh water is much greater than reported. This is a consequence of the extremely small currents able to flow in the highly resistive solution. After several recordings, it was apparent that the signal was lost in the noise floor of the measurement instrument and further recordings from fresh water were abandoned. For completeness, the recorded values are presented in between the dashed black horizontal bands of the figure.

In figure 6.2, two general trends are readily apparent: (1) observed resistance is proportional to the distance between the electrodes (transmitter separation) used to introduce the field and (2) resistance is inversely proportional to salinity. To understand these trends analytically consider:

$$G = \kappa \frac{A}{L} \tag{6.4}$$

where G is electrical conductance in Siemens ($1S = 1/\Omega$), A is the working area of each electrode (in units of m^2), L is the distance (in meters) between the electrodes, and κ is the conductivity of the solution when two flat electrodes are separated by 1m and the area of each electrode is $1m^2$. It is defined more directly as the product of the molar conductivity (Λ_m) in units of Siemens square-meters per mole (Sm^2/mol) and molar concentration (mol/m^3).

$$\kappa = \Lambda_m c \tag{6.5}$$

Combining (6.4) and (6.5) gives a complete analytic expression for the expected values of figure 6.2.

$$G = \frac{\Lambda_m c A}{L} \quad (6.6)$$

Given that conductivity (G) is the inverse of resistance ($R = 1/G$), the two experimental trends previously introduced are clearly reflected in (6.6). Namely, salinity (c) and conductivity (G) are linearly related and positively correlated, whereas salinity and electrode separation (L) are inversely related and negatively correlated.

6.2.4 Power

From the standpoint of establishing the field, it becomes easier (increased output power for constant electrode potential) to inject power into the ocean as salinity is increased. This is consistent with biologists, who have observed that electroreception in salt-water is more efficient due to the higher conductivity of the medium. Freshwater catfish have a best reported sensitivity of $1.0\mu V/cm$, whereas marine sharks achieve better than $5nV/cm$ [Zel11] – an improvement by three orders of magnitude. Fish not only find generating the fields more metabolically rewarding, but detecting them is more efficient as well.

Given that Ohm's law equates electrical output power (P) with the product of current (I) and voltage (V) and voltage, itself, being the product of current and resistance (R), it may be stated:

$$P_{rms} = \frac{V_{mpp}(V_{spp} - V_{mpp})}{8R_s} \quad (6.7)$$

where P_{rms} is the DC equivalent transmit power, V_{mpp} is the observed voltage, V_{spp} is the sourced voltage (1Vpp), and R_s is the output resistance of the voltage




Figure 6.3 – Effective transmit power into the water as both the distance between the transmit electrodes and salinity of the water is increased. The conductivity of the water plays the greatest role in determining output power.

source (50.0Ω).

The DC equivalent power of an AC sine wave is obtained from the root-mean-square (rms) operation. It is well approximated by the peak amplitude divided by the square-root of 2. The peak amplitude is one-half of the peak-to-peak amplitude resulting in a V_{pp} -to- V_{rms} conversion factor of $1/2\sqrt{2}$. Collecting these terms as they develop results in a constant value of 8 in the denominator of (6.7).

The experimental findings appear in figure 6.3. The bottom-most trace is for fresh water. It represents the smallest observable output power due to the measurement instrument's noise floor. The tiny currents in the highly resistive

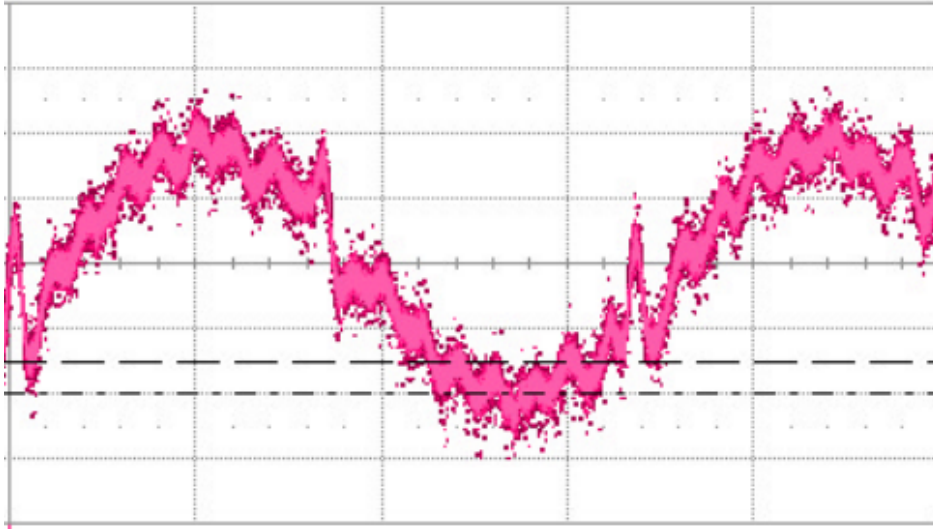


Figure 6.4 – *An amplified electrostatic signal in the Pacific Ocean. Sampled after differential amplification across a pair of receive electrodes spaced 5cm apart. The actual signal is the 1kHz waveform riding on a 60Hz noise-form with high-frequency power-supply transients interspersed every 10ms.*

fresh water, and their corresponding powers, were likely substantially lower than reported here.

In the more conductive waters, the output power increases with salinity and decreases in accordance with an increase in resistance as the electrodes are separated (see figure 6.2).

6.2.5 Electrical Noise

Powerline, power supply, nearby processors, airborne radio signals, and other sources of electrical noise can severely hamper the high-impedance, high-gain measurements required for artificial electroreception. Figure 6.4 shows a representative amplified electrostatic signal recorded in the Pacific Ocean during one of our early expeditions. It is shown digitized after differential amplification across a pair of receive electrodes spaced 5cm apart. The actual signal is the 1kHz waveform riding on a 60Hz noise-form with high-frequency power-supply transients interspersed

every 10ms. The captured signal is DC coupled with a 1Mega Ω input impedance, 200mV/division on the y-axis, no offset, 5ms/division on the x-axis, with data captured at a rate of 5 million samples per second.

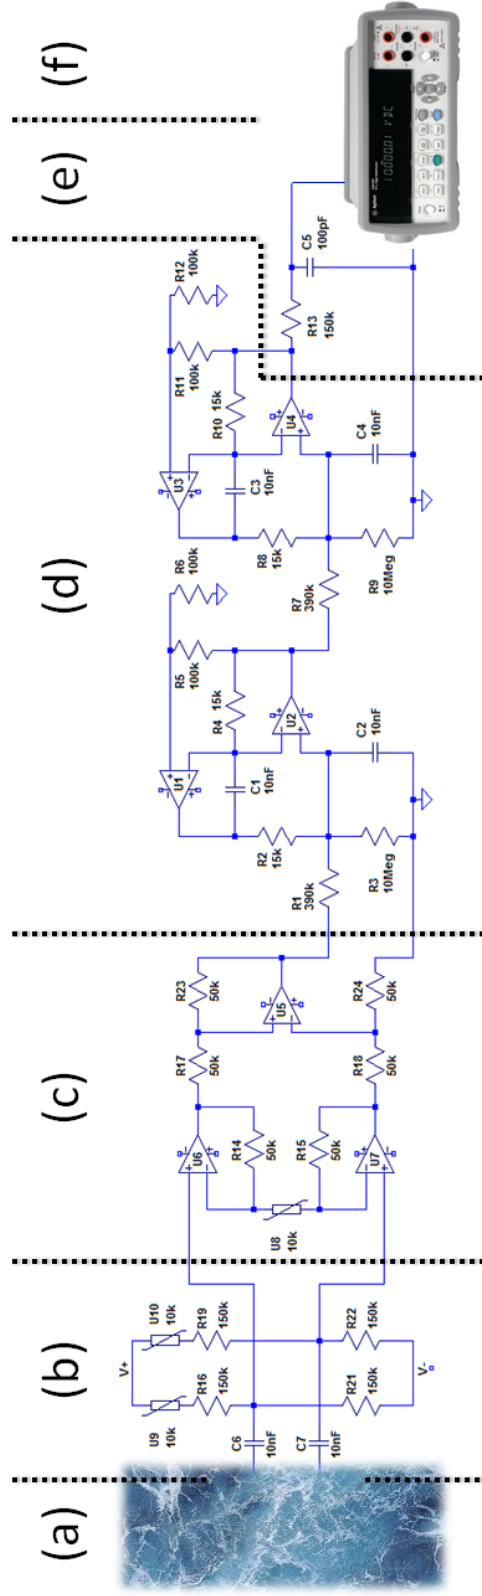


Figure 6.5 – The six stages of the Analog Front End (AFE) designed for our Biomimetic Imaging System: (a) platinum electrodes in the ocean, (b) bridge bias balance network, (c) instrumentation amplifier, (d) high-quality active BPF, (e) passive LPF, and (f) data acquisition – Agilent 33401A.

6.3 The Single Differential Channel Analog Signal Path

To minimize noise and resolve the electrostatic component of the received signal, a multi-stage Analog Front-End (AFE) was designed and employed as indicated in figure 6.5. Within the figure, from left:

- (a) *Electrodes*. Each channel consists of two platinum electrodes inserted into the ocean to form an electrode pair. The probe, as deployed in our validation experiments, consisted of two such orthogonal channels with the electrodes located collinearly $\pm 2.54\text{cm}$ from the center point. An acrylic frame was fabricated to fix the relative positions of the electrodes (figure 6.7).
- (b) *Bias Restoration*. The electrodes are AC-coupled into the bias restoration unit which consists of a manually-tunable Wheatstone bridge. The bridge allows the bias currents flowing in the bridge network and the subsequent stage input to be balanced out. AC-coupling into the bridge input impedance creates a first-order passive High-Pass Filter (HPF) with a corner frequency of approximately 100Hz. This serves to eliminate DC-drift from long-term ionization as well as interference from the power distribution system⁴.
- (c) *Preamplifier*. Balancing the input currents via the bias unit eliminates the majority of the electrical noise that originates in the power supply as the output signal from that stage is fed to an instrumentation amplifier which amplifies only the difference between the potential experienced by the electrodes and expresses that difference as a voltage with respect to our instrument's reference level (ground). Although indicated with discrete components, the actual realization of (c) used an integrated circuit – an Analog Devices AD623.
- (d) *High-Quality Band-Pass Filter*. The conversion to a single-ended reference allows the signal to continue to a High-Quality Band-Pass Filter (HQ-BPF).

⁴In the United States, where data was recorded, the public electrical utility grid is operated at a frequency of 60Hz

Implemented as a fourth-order Fliege BPF, the filter is extremely tight, passing less than 40Hz of bandwidth and transitioning between passband and stopband amplitude faster than $15V/Hz$.

The transfer function is of type I Chebyshev form, designed to sacrifice passband flatness (not material with our tiny passband) for steeper roll-off. In our application, the steeper roll-off results in greater effective sensitivity (from greater available dynamic range) as well as greater accuracy (from the higher isolation of our signal from stopband energy).

Our Fliege BPF is actually a two-stage Fliege cascade where each stage is a two pole (second order) filter topology with its attendant fixed gain of 2 (6 dB). Given two Operational Amplifiers (Op-Amps) per stage, the Fliege architecture is the lowest component count topology possible for achieving our design parameters. As a capacitor is connected from the Op-Amp output to inverting input, the Fliege topology cannot be used with fully differential amplifiers – an additional motivation for the presence of (c).

The HQ-BPF was implemented with Texas Instruments TLC25M2 Op-Amps. The actual design objectives were for a unity-gain active filter with a quality factor of at least 10 and a center frequency very near to 1kHz. The filter was designed and then its component values adjusted to fit parts availability, minimize the number of unique parts, and improve performance. Two identical filters were built with the results from one of them (the left channel when viewed as mounted in figure 7.1) compared against the simulated expectation. The results appear in table 6.1.

The measurements were made with an Agilent 33220A Arbitrary Waveform Generator (AWG) as the signal source ($1V_{pp}$ ⁵, Sine wave) with the output assessed by a Rigol DS1102E 2-Channel 100MHz oscilloscope. Readings were taken using the measurement processor with the scale set to the smallest range

⁵Volts peak-to-peak




Figure 6.6 – *Expectation and achieved performance of the HQ-BPF. The dotted line represents the simulated high quality band-pass filter while the solid blue line shows the measured filter output. Despite lower gain, the filter achieves the desired $Q > 25$ criterion.*

that the offset could still compensate to keep the top peek of the signal in view of the scope. Both the analog bandwidth limiter (20MHz LPF) and a digital filter (2.5kHz LPF) were applied. Then 4 wave crests' maximum values were averaged to get each amplitude reading. The offset was measured post-sweep and added to the recorded values. The observed amplitude values were doubled prior to comparison with the simulation as the simulation used a 2Vpp signal as the source.

All of the critical design objectives were easily met. The center of the passband hits the mark to within about 1% error. The passband is extremely narrow resulting in a quality factor that is well above our objective of 10. The gain

meets our original objective of unity, but falls far short of the simulated expectation. In practice, this was fortunate as our field trial pushed the preamplifier gain quite high and an additional $> 3X$ gain would have exceeded the available dynamic range, saturating the filter’s component amplifiers. The large gain error could be explained by a missing modeled load. A missing load of $240k\Omega$ would bring the simulation and build to parity. Other parasitics, missing from the model, could also be the culprit.

The performance of the second filter (the right channel) is extremely similar to the first with the center frequency differing by less than 10Hz ($f_0 = 1.080\text{kHz}$). Its assessment was omitted for clarity.

Table 6.1 – *The active bandpass filter*

| Parameter | Simulation | Measured | Units | Error |
|-----------|------------|----------|-------|--------|
| f_0 | 1.058 | 1.071 | kHz | 1.2% |
| Bandwidth | 27 | 37 | Hz | 37.0% |
| Quality | 39 | 29 | – | -26.1% |
| Offset | 0 | -19 | mV | – |
| Gain | 3.646 | 1.891 | X | -48.1% |

- (e) *Post-Filter*. A passive first-order Low-Pass Filter (LPF) is used to clean-up the high frequency ripple introduced by (d). It has a corner-frequency of approximately 10.5kHz.
- (f) *Signal Capture*. Digital acquisition is handled by a pair of Agilent 34410A digital multimeters. Their read significance is 6.5 digits, which is achieved internally through the use of additional signal amplification. This internal range handling serves as the final gain stage in our imaging platform reducing the need for gain in the preamplifier, where providing gain can compromise the dynamic range available.




Figure 6.7 – *The receive electrodes and mounting fixture shown just prior to submergence. The electrodes are arranged into two differential channels on orthogonal axes.*

6.4 Conclusion

The design proposed in this chapter for a single differential channel electrostatic receiver is the terminal evolved product of this dissertation. While the immediately subsequent chapter is concerned with performance expectations and validation, later chapters will explore compositing these channels into larger arrays, shaping signals for transmission and detection through them, and encoding information such that receivers of this type may recover it.

CHAPTER 7

Modeling Differential Electrostatics

7.1 Overview

Predicting ocean performance is an involved task. BEI leverages the quasi-static condition to greatly simplify Maxwell’s equations, but relies upon a narrow-band signal propagating in a homogeneous environment to do so. Validating our design requires that we demonstrate the degree to which these assumptions hold when deployed – namely, that our element captures the self-generated field (to the exclusion of ambient energy) and that we are not suffering internal or airborne leakage (whereby signal is merely cross-coupled from our transmit hardware to our receive hardware – bypassing the ocean entirely).

To achieve this degree of confidence, a model of the electric field in a dielectric will be developed and employed to predict the nature of the voltaic response expected. When our transmit electrodes are excited and detected by a pair of differentially configured electrodes, a characteristic field pattern results. If we observe this same pattern in-situ with actual hardware, we may be reasonably assured that the system is indeed utilizing quasi-static electric fields and no significant leakage has occurred.

7.1.1 The Electric Field in Theory

Let us consider a Cartesian space in two dimensions representing a submerged plane of constant depth. $R = \langle p_n, p_{n+1}, \dots, p_m \rangle$ where $p_{x+1} = p_x + i_R$ defines the

extent of the square space¹ from diagonally opposing corners (p_n, p_n) to (p_m, p_m) with a quantized spatial resolution of i_R .

We define the unitary matrix

$$\mathbb{U}_{i,j} = \begin{cases} 1 & (i,j) \in R \end{cases} \quad (7.1)$$

such that $\mathbb{M}_x = \mathbb{U} \cdot R$ and $\mathbb{M}_y = R^T \cdot \mathbb{U}$ and enclose element-wise matrix operations in a circle. Therefore, \odot is equivalent to the Hadamard product:

$$(\mathbb{A} \odot \mathbb{B})_{i,j} = (\mathbb{A})_{i,j} \cdot (\mathbb{B})_{i,j} \quad (7.2)$$

Similarly, \otimes indicates raising each element to a power and \oslash indicates element-wise division.

The distance matrix \mathbb{D}_k ,

$$\mathbb{D}_k = ((\mathbb{M}_x - \mathbb{U}x_{q_k}) \otimes^2 + (\mathbb{M}_y - \mathbb{U}y_{q_k}) \otimes^2) \otimes^{\frac{1}{2}} \quad (7.3)$$

whose i, j^{th} element represents the Euclidean distance between the point in space $(\mathbb{M}_x_{i,j}, \mathbb{M}_y_{i,j})$ and the location of electrode k is calculated for each emitting electrode in the system. For our current BEI, there is only one pair of transmitting electrodes. Those two electrodes, $k \in \{0, 1\}$, are located at coordinates (x_{q_k}, y_{q_k}) and each inject a net charge of q_e opposite in polarity from each other.

Accordingly, the voltage experienced by a unit positive test charge in the field of the BEI transmitter is \mathbb{V}_{SE} ,

$$\mathbb{V}_{SE} = \frac{q_e}{4\pi\epsilon_0\epsilon_r} ((\mathbb{D}_1 - \mathbb{D}_0) \oslash (\mathbb{D}_0 \odot \mathbb{D}_1)) \quad (7.4)$$

obtained through applying Coulomb's law given ϵ_0 is the permittivity of free space

¹The extension to rectangular spaces is trivial – select n, m separately for each axis.

and $\epsilon_r \approx 81$ is the relative permittivity of ocean water.

We define the transmitter-parallel shift operation (\triangleright) as follows:

$$\mathbb{B} = \mathbb{A} \triangleright c_s \quad \Rightarrow \quad \mathbb{B}_{i,j} = \begin{cases} NaN & j < c_s \\ \mathbb{A}_{i,(j-c_s)} & \text{otherwise} \end{cases} \quad (7.5)$$

where NaN indicates an undefined (incalculable, unplotable) value and c_s is the shift constant. The shift operation occurs only in an axis parallel to the axis of the transmit electrodes mirroring our experimental configuration.

It is now possible to predict the differential voltages that will be observed when our given BEI imaging element configuration is deployed in the ocean by:

$$\mathbb{V}_{Diff} = \left| \left[\mathbb{V}_{SE} - \left(\mathbb{V}_{SE} \triangleright \left[\frac{d_{TX}}{i_R} \right] \right) \right] \triangleright - \left[\frac{d_{TX}}{2i_R} \right] \right| \quad (7.6)$$

Of note, the absolute value models our AC-coupled input stage and the second shift operation is necessary to recenter the matrix on its ground truth location since our gantry coordinates refer to the center of the differential channel. The predictions of the model, compared against the actual values retrieved during an expedition, appear at the end of the chapter (figure 7.4).

7.1.2 The Electric Field in the Ocean

For ocean validation, several expeditions were made to the Pacific Ocean. A site in the coastal city of Marina Del Rey², California (United States) was chosen for its proximity to infrastructure and vehicular accessibility.

Access to the site, vacant boat slip G-2823 on Pier 44 of the Mindinao Road marina, was made possible by the University of Southern California's Information Sciences Institute (USC ISI) who also has instruments deployed at this location.

²Marina Del Rey is approximately 30km west of downtown Los Angeles, California




Figure 7.1 – *Instrumentation and measurement equipment, shown on location, includes: an oscilloscope, data processing workstation, analog front ends, bipolar power supplies, and Agilent digital multimeters.*

The purpose of the expeditions was to map the differential electric field in a salt-water tank with an effectively infinite far-field (e.g. the ocean). In order to collect field measurements from precise positions, a purpose-built motion-controlled gantry system consisting of a two axis carriage with a 78.75cm (31") x-axis range of motion and a 83.80cm (33") y-axis range of motion was designed and constructed

(figure 7.2). The position resolution of the gantry system is 2000 steps per inch (787.4 steps per cm) and repeatable to better than $5/1000$ of an inch. Naturally, this far exceeded the environmental variation imposed on our electrode positions by the wind, waves, and human on-dock passersby.



Figure 7.2 – *A custom-built motion control platform and cantilever frame was constructed to gather positionally precise high density electric field measurements at depth in the Pacific Ocean.*

The gantry was placed atop a custom purpose-built cantilever frame and suspended above the water as shown in the figure. 180 kilograms (400 pounds) of ballast, in the form of sand bags, were used to secure the cantilevered gantry. During safety testing in dry dock, the cantilever frame was demonstrated to safely support more than 73kg (160lbs) applied at its maximum extent, albeit with significant bending.

The receive (measurement) electrodes were submerged to a depth of 104cm (3'5") using a 1.9cm (3/4") diameter acrylic dowel, the submerged-end appearing in figure 6.7. The transmit electrodes were submerged to a depth of 120cm (3'11") using a second acrylic dowel of identical dimension. In figure 7.2, the two acrylic rods are visible above the waterline with the receive electrodes mounted to the left-most rod. The electrode positioning rods were suspended from a floating dock. The mean lower low tide water depth (MLLW) for this location is 3.05 meters[Boa]. During recording, the water depth below the electrodes (themselves $\approx 1m$ down) varied from $\approx 2m$ to as much as $\approx 5m$ during the 6.5 hours required to collect the dataset.

As we previously demonstrated [FTS10], submerged non-conductive bodies may be used to shield or focus the field. Since, in this case, that behavior is undesired, the receive electrodes were kept slightly above the transmit electrodes³. The transmit power, from the method of equation (6.7), was $250mW$.

At the time of the final expedition, the ocean water salinity was 31.5ppt as recorded with the same hydrometer used in section 6.2.3.

The results from the differential receive channel parallel to the transmit electrodes appear in figure 7.4. The dataset consists of 10,610 points sampled on a linear spatial grid with intra-point spacing of just 0.66cm. All of the points were captured at a uniform depth. The capture time for a single datapoint (including

³The receive electrodes protrude from the bottom of their mounting frame (figure 6.7). Placing the field below the mounting frame results in less disturbed measurements. Up to 6dB of attenuation was observed for signals originating above the mounting frame.




Figure 7.3 – *The electrostatic field of a BEI transmitter electrode pair as calculated at depth in the Pacific Ocean.*

transit time) is about 5 seconds.

The hyperbolic shaped nulls are clearly visible and a strong position and polarity dependence is indicated – agreeing well with the expected pattern of an electric field generated by an electric dipole and measured differentially. Three




Figure 7.4 – *The electrostatic field of a BEI transmitter electrode pair as observed at depth in the Pacific Ocean gathered during an expedition to Marina Del Rey, California. Note the strong agreement with the simulated data of figure 7.3.*

lobes are visible. The middle (primary) lobe is detectable out to 70cm (figure 7.4 (bottom)). The smaller lobes reach to only 23cm. This data, the first of its kind in publication, verifies the existence and utility of near-field electric phenomenon for




Figure 7.5 – A model of two sensor channels (cyan and blue) being disturbed by an induced dipole. The red X's are the the transmitting electrodes.

in-ocean application.

7.2 Modeling the Differential Disturbance Field

Electrostatics, as the name implies, involves the use of electric (E) fields which are, traditionally, invariant with time [FTS10]. For our purposes it is illustrative to consider a sequence of time-invariant fields with each successive field having more (then less, then more, etc) strength than its predecessor. When this approximation is valid, the field is said to be *quasi-static* [MT76].

7.2.1 Modeling Electrostatics

Having established the operating bounds for the validity of the electrostatic assumption, we may now introduce what is effectively a static (DC) field into the ocean environment. An electric field in a conductive environment will necessarily result in current flow. We refer to the electrode pair utilized to establish this field




Figure 7.6 – *A detail view of the receiver channel located at the midpoint between the two transmitting electrodes. Only the center part of the response is shown.*

as the *transmit electrodes*. Detection is performed differentially through adjacent equally-spaced electrodes arranged collinearly with the transmitters (the *receive electrodes*).

In [FTS10], we introduced our purpose-built physical modeling engine, built in Matlab, used to visualize the electric fields and disturbance that occurs with environmental objects. The work demonstrated the effect of a conductive object (which presents as a dipole) on a universally referenced field (a single-ended measurement). However, when the sensor was actually constructed [FTC10], differential measurements were used. In this work we have extended our model to evaluate field disturbances differentially.

The model works by ignoring the homogenous background environment and restricting analysis to only the relevant net charges (with respect to the background

ion concentration). Direct analysis then becomes computationally practical via⁴:

$$Voltage = \sum_{n=0}^{N_c} \frac{1}{4\pi\epsilon_0} \frac{q_n}{r_n} \quad (7.7)$$

Where N_c is the number of net charges in the environment to consider and q_n expresses the net magnitude of each charge. r_n is the distance between the spatial location under consideration and each specific net charge. ϵ_0 is the permittivity of a vacuum. The equation may be scaled by the relative permittivity of other media for consideration there, but as this has no spatial consequences (only influences magnitude), here it is ignored.

Figure 7.7, is an example field analysis showing the transmit electrodes as black X's, the receiving electrode array as red squares, and the poles of the induced dipole as blue circles. The actual units and magnitudes are unimportant, the purpose is to illustrate the isovoltic features at scale. A dotted line is drawn horizontally through the figure to help visualize the observations made by a linear array in the field which is parallel to the transmit axis.

7.2.2 Basic Assumptions of BEI

In the specific case of making our electric field disturbance detections, some additional constraints are required:

The diameter of the dipole must be much smaller than the transmit basis (the minimum distance between two transmit electrodes). The end effect is that the solution space must be sparse. It may contain multiple small-diameter objects so long as they are unclustered. If the dipole were modeled as a substantial percentage of the transmit-basis then the receiver array would experience the dipole as a set of smaller dipoles all interacting (as we are in the near-field).

⁴This model is only partially correct as the presence of the induced dipole will, in turn, affect the other charges in the vicinity, but we may ignore these higher-order effects for now as we concentrate on establishing the basic BEI image formation theory.




Figure 7.7 – *Simulated field analysis showing the transmit electrodes as black X's, the receiving electrode array as red squares, and the poles of the induced dipole as blue circles. The actual units and magnitudes are unimportant, the purpose is to illustrate the isovoltic features to scale.*

The transmit-basis must remain much larger than the separation between receive electrodes used in the differential measurement. Increasing the receiver separation reduces the spatial resolution of the measurement and in the limiting case (electrodes very far apart when compared with the transmit electrodes), the differential measurement devolves to the single-ended case.

7.2.3 BEI Expectations

With these restrictions in place, a series of analyses were performed to identify the expected behavior of a differential field disturbance probe (a single channel composed of two individual electrodes over which a potential difference is recorded). In each of the figures, the red X's indicate the location of the transmitting electrodes,

the colored circles indicate the center location of the receiving electrodes (which were located equidistant to either side of the marker), and the colored curve families represent the observations of the correspondingly colored receive channel. The X and Y axes form a plane in real space where the dipole is physically located at the time the Z value for field potential is observed. The spatial coordinates are units of relative distance normalized to: transmit-basis = 0.4 units. The amplitude values are relative to their far-field zero, but lack significance as the actual magnitude will depend on the strength of the induced dipole.

Figure 7.5, shows the results from two receive channels, one located at the origin and one to its left. For each channel, three curves are plotted with the decrease in amplitude corresponding to the increasing distance of the dipole away from the transmit-axis (into the page). Several observations should be readily apparent from the figure: (1) The disturbance field is limited in extent and will not influence all of the receive channels if the array is large. (2) The disturbance field is centered on the receive channel which observed it. Put more simply, the peak sensitivity of the channel is to dipoles located directly in front of it. (3) This is a near-field phenomenon. The disturbance quickly returns to baseline with increasing distance between the dipole and the sensor. (4) There is a baseline value unique to the channel which is determined by the transmit electrode spacing, receiving electrode spacing and position⁵, output power, and medium.

Figure 7.6 explores the sensing range in greater detail. The parabolic near-field behavior results from the dipole's diameter being equal to the receiver electrodes' separation and vanishes with increasing distance (the dipole diameter appearing smaller to the receivers). The signal strength vanishes inversely proportional to the separation (e.g. $\Delta V \propto 1/d$).

⁵relative to the transmitting electrodes

7.3 Conclusion

This chapter concludes our discussion on the fundamentals of electrostatics. It has culminated in a differential quasi-static model of the disturbance field and has validated it against data collected from ocean deployments. It has been shown that the electrostatic channel exists and is accessible. The following chapters will utilize electrostatic channels and explore optimizations in terms of channel access as well as best practices in encoding location specific information.

Part III

Waveform Design

CHAPTER 8

Waveform Shape

8.1 Overview

With the completion of the BEI single channel sensor hardware in chapter 6, we can now move to the design of the transmission waveform. There are three principal aspects to consider: (1) the shape of the waveform, (2) the emission pattern into which energy is focused, and (3) the information encoded in the transmission to improve image formation performance. This chapter is devoted to the first design point. Chapter 9 is devoted to the second and chapters 12 and 13 are devoted to the third.

8.2 Waveform Classification

Sorted by spectrum, the electric-field excitation waveforms used by electric fish fall into two broad categories: pulse-type (broadband) and wave-type (narrowband). Most readers are familiar with the popularized “electric” eel (*Electrophorus electricus*) so referred for its ability to generate powerful electric shocks (in excess of 500 Volts!) which it uses for both hunting and self-defense. These emissions are pulsatile, repetitive, and uniform in polarity.

In contrast, wave emissions are bipolar alternating the current direction between head-to-tail and tail-to-head. The *Apteronotus albifrons* (black ghost knifefish) is an example of a wave-type electric fish. It is a member of the Gymnotiform order,

just like the electric eel¹.

Knifefish generate electric fields using a specialized electric organ located in the tail region of the fish and detect disturbances to the generated field with receptors along the body. If the animal were to swim in the conventional manner, by deforming a caudal (tail) fin, the attendant body oscillation would distort the sensor array superimposing external field disturbances (prey and predators) on self-induced swimming oscillations. Consequently, Gymnotiforme bodies remain rigid. Propulsion, instead, comes from a long, almost transparent, undulating fin running beneath the body. The rigid body approach makes the fish move like a knife through the water, hence the name.

8.2.1 Tank Setup

To compare the waveforms, a 55-gallon (≈ 208 liter) acrylic fish aquarium ($91\text{cm} \times 51\text{cm} \times 38\text{cm}$) was modified to support further study. The tank configuration is indicated in figure 8.1. It consists of four electrodes. As the electrodes used in these experiments were intentionally subjected to extreme environments, the electrodes used in these experiments were *steel* due to their greater availability and substantially lower cost. Doing so allowed the superior platinum electrodes to be reserved for actual ocean deployments.

The four electrodes used in these experiments are an excitation electrode, P_T , a neutral electrode, P_0 , a reference electrode, P_S , and an electrode, P_G , mounted on a movable motion-controlled gantry above the tank – as labeled in figure 8.1. All of the electrodes are submerged to an equal depth in the tank of approximately 7.5cm , while the tank itself is filled near capacity. The current axis runs between electrodes P_T and P_0 which are separated by 16cm . Subsequently, when reporting Voltage measurements, it is to be understood that all measurements are taken at

¹Despite its name, the electric eel is not closely related to true eels (Anguilliformes) but is a member of the Neotropical knifefishes (Gymnotiformes)

the indicated electrode with respect to $P_0 = 0Volts$.




Figure 8.1 – *A diagrammatic view of the tank electrode configuration.*

The filled water was an ocean simulant constructed from tap-water and iodized salt. The salt concentration was increased until a Marineland Labs, inc. Instant Ocean Hydrometer reported a salt concentration of 28ppt and a relative density (pure water reference) of 1.0205 at room temperature. This is the extreme low-end of the normal range for the Pacific Ocean [Mar07], which creates the intended worst-case scenario for testing.




Figure 8.2 – *The experimental salt-water tank and supporting test, excitation, and measurement equipment.*

8.2.2 Pulse-type EOD

Pulsatile transmission is typically unipolar and consists of sharp rise and fall times. To achieve emulation, excitation using an Agilent E3631A configured as a 10mA current source and operated under computer control was applied. The E3631A was powered on for one minute and then switched off for five minutes moved to the next longitudinal position and the cycle repeated.

The turn-on time of the E3631A was measured into a 50Ω load targeting 20mA of load current with 19.2mA average achieved during testing. Under these conditions the average rise time was 4.271ms with a standard deviation of 1.246ms varying over 3.403ms to 6.738ms. The sample size was 100 trials. This closely approximates the rise time of the Brown Ghost Knifefish [TZ05].




Figure 8.3 – *Conductivity study of constant polarity excitation in our ocean-simulant tank.*

ions are generated around the electrode). Without sufficient means to facilitate the rapid diffusion of Fe^{2+} ions, the concentration near the surface becomes larger than the that in the bulk solution – resulting in a Voltage between the surface and the solution [Chi71]. This voltage is called the *concentration polarization potential* and acts to *reduce* current flow (increase resistance, reduce conductivity).

Alternatively, consider that for the first 50 seconds, the Fe in the electrode is oxidized but remains in the electrode’s lattice structure as the potential is insufficient to overcome the work function and drive the ions from the surface. The positively charged ions collect on the surface layer by layer while negatively charged ions from the solution surround the anode. As the potential difference grows it becomes sufficient to break the lattice and the ions release into the solution. The sudden addition of ions to transfer charge produces a dramatic drop in resistance.

8.2.3 Wave-type EOD

The exact chemistry of the tank system was not explored in detail experimentally, but the outcome is indicative. Constant unipolar currents result in unintended disturbances to the background conductivity of the sea-water. For this reason, both nature² and experimental measurements (conductimetry) employ oscillating field potentials to drive these chemical reactions and polarization potentials in both polarities and, therefore, average out their effects.

Our experiment was repeated replacing the E3631A with an Agilent Arbitrary Waveform Generator (AWG) 33220A. The 33220A produces a sine wave at 200Hz and a 10Vpp output through a 50Ω output impedance which results in an average of $\approx 225mVAC$ in the water. The results of the experiment appear in figure 8.5 in which, again, the lower surface represents P_S , the middle, P_G , and the top, $P_G + P_S$. Notice that all of the surfaces are now time-invariant. The lower density of points in time, as compared to figure 8.3, is a product of the discrete windowed

²While not truly oscillatory, even pulse-type fishes do not maintain DC level for very long.




Figure 8.5 – *Conductivity study of alternating polarity excitation in our ocean-simulant tank.*

nature of AC measurement with digital sampling instruments.

8.3 Conclusion

The electric organ discharge (EOD) of most species of the knifefishes (Gymnotiformes) of South America is of the wave, not the pulse, type [Kra99]. This is true for the broader ocean as the majority of electric fish species are weakly electric (amplitude < 1 Volt) wave-type [Eva97]. Wave EODs are usually of constant frequency and amplitude and show a bewildering multitude of species-characteristic waveforms. The EOD of *Eigenmannia* is sexually dimorphic in waveform and in the intensity of its higher harmonics [Kra99].

Our experimental findings help motivate this biological reality. As shown in this

chapter, unipolar transmissions, repeated frequently over enough time, will disturb the electrical environment around the fish precluding background subtraction and reducing sensitivity severely. Electric fish species find that it is more advantageous to use phase, frequency, and amplitude modulations to a continuous, bipolar, wave output rather than attempting time-based modulations to a pulsed or DC-type output. Accordingly, it is this behavior that we will endeavor to emulate.

CHAPTER 9

Waveform Foveation of a Single Dipole

9.1 Overview

To enhance signal-to-noise ratio (SNR) and, correspondingly, range, fish attempt to make their EOD's directional by focusing the current into the region they are most interested in. This is possible because the fish is covered in a skin system that is very electrically resistive. Internally, its physiology maintains a chemistry that is consistent in concentration and uniform in distribution. That, in turn, creates a reference potential that is isolated by the skin from the outside world. The Voltage measurements taken by the sensory organs are a transdermal potential – the Voltage across the skin. This is indicated in figure 9.1. Because current follows the path of least resistance, if distributed from head-to-tail, it will flow outside around the fish rather than through the fish. By deforming its body the fish can bias the flow of current more heavily to one side or the other.

9.2 Directivity

To evaluate the possibly of mimicking this behavior in our sensor, a Delrin (non-conductive) plate was installed behind the plane of the transmission electrodes at various distances (figure 9.2). The strength of the resulting field was measured along the long-axis of the tank (x direction) and across it (y direction).

The results appear in figure 9.3. Within 2cm, the effect of the Delrin plate




Figure 9.1 – *Electric fish have thousands of voltage sensors arrayed along their body and concentrated around their head. During prey capture, they curl their bodies to focus the electric field toward their prey. Figure derivative of [Nel11b].*

appears to saturate with little additional benefit observed on a linear scale. However, when viewed on a logarithmic scale (figure 9.4), the continual improvement in gain with proximity is more readily seen. The plotted ratio is the gain in detected signal level ($V_{ac,rms}$) versus the absence of the plate. At its closest, the gain was 0.452dB.

With the reflector closest to the plane of the transmit electrodes a gain of 15.7% in signal power (as detected at a reference point located along the line normal to the transmit axis in the same depth plane) occurred. However, transmit power at a reference point does not adequately convey the significance of reflector-based focusing because the reflector changes the shape of the decay function. Due to this phenomenon, transmit range improved as much as 11-fold at the peak of the




Figure 9.2 – A Delrin (non-conductive) plate installed behind the plane of the transmission electrodes. The gantry, visible above the tank, is installed in its y-axis configuration (across the short axis of the tank)




Figure 9.3 – Directivity along the long dimension of the tank (top) and across the short dimension of the tank (bottom)

focused field.




Figure 9.4 – *Signal gain across the short-axis of the tank as a result of foveation.*

9.3 The Differential Field Pattern

To this point, the experimental measurements of this chapter have all been single-ended. The zero voltage electrode of the transmit pair is used to reference a single measurement electrode moved about the space. To improve noise performance in our actual deployed systems, it becomes necessary to take differential measurements in the space. Two measurement electrodes are employed and their values subtracted to yield a more sensitive signal that makes better use of the available dynamic range. This hardware was detailed in chapter 6.

Recalling the ocean expedition data from section 7.1.2, when viewed end-on, three of the four¹ lobes become quite evident (figure 9.5). While maintaining this

¹The fourth lobe mirrors the primary lobe in the opposite direction. In the context of figure 9.5, the fourth lobe exits the page coming straight out at the reader, while the primary lobe it mirrors goes straight into the page.




Figure 9.5 – *The Pacific ocean recording*

viewpoint, if the lower section of the dataset is enlarged figure 9.6 results. Clearly evident are four thresholds defined from top as: (1) the limit above which the SINR is clearly sufficient for detection ($> 30mV$), (2) a reasonable guess as to the limits of detectability for our current receiver hardware ($> 20mV$), (3) the expected limit of detection as the receiver hardware is improved, and (4) the absolute zero (noise floor) of our measurement equipment.

Plotting contours between the first and second thresholds gives figure 9.7. The central red region indicates the area of assured signal reach, while the multi-colored transitional regions indicate areas of intermediate detection probability. The white regions are areas where the electric field has virtually no influence. Evident is the inherently directional nature of the differential electric field pattern.




Figure 9.6 – *The expeditionary dataset with four thresholds overlaid. From top: (1) the limit above which the SINR is clearly sufficient for detection ($> 30\text{mV}$), (2) a reasonable guess as to the limits of detectability for our current receiver hardware ($> 20\text{mV}$), (3) the expected limit of detection as the receiver hardware is improved, and (4) the absolute zero (noise floor) of our measurement equipment.*

9.4 Foveation Through Selection

Foveation is the biological concept of focusing [Eva97]. The human eye is heavily foveated [LF92]. It is why you can only read fine print when you are looking directly at it and not with your peripheral vision. Structurally, this manifests from the non-uniform distribution of light sensitive cells in the eye ball. The highest concentration of photoreceptors are located dead center at the back of the eye in a region known as the macula [LF92]. The macula is at the center, or fovea², of the light sensitive surface called the retina. The fovea (or fovea centralis) is responsible for the sharp central vision (also called foveal vision [LF92]), necessary in reading, driving, playing ball, and any other activity where a high degree of visual acuity is advantageous.

Electric fish foveate their electroreceptive sense as well, demonstrating greater

²The term fovea comes from the Latin, meaning pit or pitfall






Figure 9.7 – *The inherently directional nature of the differential electric field pattern. The central red region indicates the area of assured signal reach, while the multi-colored transitional regions indicate areas of intermediate detection probability. The white regions are areas where the electric field has virtually no influence.*

sensitivity near the prey capture area (head/mouth) than along their bodies [MSM01]. In analogous behavior to a human turning their head to put an object of interest in the center of their visual field of view, the electrostatic emission may be pointed as well. As evident in figure 9.7, differential, dipole-generated, electric fields are inherently lobular and, therefore, directional. To demonstrate this, consider figure 9.8.



(a) A differential electrostatic model of a single dipole transmitter operating at in the ocean using parameters tuned to match an actual deployment.



(b) The dipole transmitter of 9.8(a) rotated a few degrees clockwise. Note the corresponding rotation in the field pattern.

Figure 9.8 – *The effect of transmitter location on the differential electric field pattern*

Figure 9.8(b), in comparison to figure 9.8(a), shows the effect of “turning your head”. The pattern remains the same, but is pointed in a new direction. The figures are calculated from our differential electrostatic model with its parameters tuned to match the configuration of our ocean expedition experiments (figure 9.7). The red contour regions use the same threshold, but the model generated figures use zero for the lowest threshold resulting in a deep blue coloration for the signal loss regions – rather than the white coloring used in the experimental data. The black X’s in the figures indicate the location of the transmitting electrodes.




Figure 9.9 – *A concept biomimetic jellyfish robotic vehicle with a circular array of electrodes around its bell. The pattern direction may be chosen simply by selecting the appropriate electrode pair.*

“Turning” via mechanical actuation is slow and costly (in terms of size, energy, and financial cost of the actuator) in comparison to electrical switching. If the system can accommodate an array of transmit electrodes, you can steer the beam direction simply by selecting an appropriate pair of electrodes to energize. A concept jellyfish-like robot with a ring of electrodes around its bell appears in figure 9.9 illustrating this approach.

9.5 Conclusion

Achieving directivity is of vital importance in sensing or communicating in an unstructured environment. Learning the position from an assumption of direction of arrival/origin is implicit in all biological notions of localization. In this chapter two biomimetic mechanisms were explored. The first employed the use of body structure to inherently eliminate the back lobe of the transmitted field. The second looked at physical pointing as a means of directing the transmission in an appropriate direction. In the chapters that follow, higher level concepts of sensory integration (combining individual emitters/detectors) will be explored to extend our mimicry to the central nervous system.

CHAPTER 10

Waveform Foveation through Array Synthesis

10.1 Overview

Given the work of the prior chapter, further focusing is possible. While the physical construction and position of the electrodes may focus a single sensing or emitting element (section 9.2), a typical adult male gymnotiform has over 15,000 of these electroreceptor sensing elements arranged in what approximates a rectangular array [Nel11b]. Utilizing the array nature of the organ, a more focused (or otherwise beneficially selective) beam may be synthesized.

We are now in a position to discuss the full range of foveation. To review, electric fish employ three mechanisms to foveate. In order of decreasing generality, they are:

1. Differential, dipole-generated, electric fields are inherently lobular and, therefore, directional. This is clearly evident in figure 9.7, which is an actual recording of a differential electric field at depth in the Pacific Ocean.
2. The fish's body is employed to give direction and focus to the emitted field.
3. The electric fish has thousands of receptors along its body among which differences may be combined and weighted to produce areas of increased sensitivity and reduce the sensitivity to other areas not of interest.

The remainder of this chapter (and the next few) will look at this final foveation

mechanism, synthesis¹, to improve the position sensitivity of a biomimetic electrostatic system.

10.2 Drawing Parallels with RADAR




Figure 10.1 – *Static electric fields generated from a dipole produce curved voltaic field lines as shown in the model. Unlike isotropic radiation, traveling along the horizontal dashed line results in non-linear changes in voltage.*

The problem of biomimetic electrostatic array synthesis has many parallels in RADAR and other antenna engineering problems. For the remainder of this chapter and the two that follow it, I will present the work in the context of RADAR

¹Array synthesis – the act of beam pattern construction from the interaction of multiple emitters/detectors

systems design. There are several advantages to making the transition at this point:

The method of propagation in a medium is only fundamental to the transit of energy from the transmitter to the intended receiver. Given that no energy (zero output power) can always be “sent” from the transmitter to the intended receiver, the existence of an energy pathway between the two, by its very existence, implies that a range of amplitudes may be sent in finite time. Given this reality, the encoding of information in a physical channel becomes largely independent of the channel propagation mechanism itself – up to the limits in amplitude slew rate, attenuation, and chromatic dispersion imposed by the channel.

Beyond the physical channel, information encoded in the signal, and other higher layer concepts, remain the same. To the extent that we can establish verisimilitude, I am able to retain compatibility with (and therefore leverage) existing RADAR literature and demonstrate the cross-applicability of my work with electrostatic physical channels to those of airborne electromagnetic ones.

RADAR makes one additional assumption – straight line propagation. In order to determine bearings in mono static pulsed RADAR, the antenna is given a very narrow beamwidth and a single primary lobe. The angle to a reflective target is then assumed to be the direction in which the antenna was pointed. In the case of electric fields generated by a dipole transmitter, the field lines are curved (figure 10.1). Despite the curvature of the field itself, *changes* to the field propagate along the straight line path since changes in the electric field result from changes in the conduction current which will follow the shortest path in order to minimize resistive losses².

Ironically, electromagnetic radiation (which proceeds outwards in all directions from an isotropic transmitter) will likely not propagate in a straight-line path between the transmitter and receiver in marine waters. It will experience too severe

²Assuming the quasi-static condition, derived in section 2.5, has not been violated.




Figure 10.2 – Propagation mechanisms for electric charges in the ocean.

an attenuation along that direct line. The permittivity of the air (≈ 1) and sea (≈ 100) are so different that the refraction at the ocean surface (water-air) interface is effectively a right angle. Electromagnetic energy detected at the receiver will be mostly composed of those EM waves which travelled the path of lowest loss, *not necessarily the path of least distance*. It is often most advantageous for EM radiation to travel indirectly, proceeding straight up to the surface (suffering loss) traveling along and just above the surface in a low-loss medium (air) and then returning straight down to the receiver. This process is illustrated in figure 10.2.

In this chapter and those that follow it, the terms *antenna* (electromagnetics) and *electrode* (electrostatics) will be used interchangeably. This is not to imply that they are structurally equivalent or even offer similar sensitivity patterns, rather it is to imply that at the system level they are *functional* equivalents and serve the

same purposes in their respective domains. Likewise, it is more conventional to discuss RADAR transmissions in air, rather than water, and doing so will aide the reader in making comparisons with alternate RADAR architectures. It is to be understood that ranges (d) and wavelengths (λ) may be converted from air ($_a$) to water ($_w$) or vice versa via the following relationship:

$$\lambda_a \approx \lambda_w \times 81 \quad (10.1)$$

$$d_a \approx d_w \times 81 \quad (10.2)$$

10.3 Numerical Analysis of Electromagnetic Structures

The problem of array synthesis, traditionally, has been one of numerical analysis of electromagnetic structures for the purpose of building novel airborne antenna systems³. Approaching the design problem from the numerical analysis standpoint is bottoms-up – the simulation must be repeated for each parameter in the array antenna design space over which one iterates. Even for mild design spaces (four elements, uniform type) this can produce enormous computational burdens to search even small fractions of the possibilities. Array systems offer no less than four independent configuration control parameters for each element in the array: (a) position, (b) phase, (c) amplitude, and (d) antenna pattern. Further, one can choose the number of elements in the array independently.

Available analysis tools divide into two broad categories: (a) closed form solutions and (b) Maxwell-based field solvers. The former category is used to handle homogeneous arrays (arrays composed of identical individual elements) in topologies that consist of uniform element spacing and uniform phase progress between

³The most common example is RADAR system design.

successive elements. In these limited situations the symmetry and uniformity of the design permits significant simplification of Maxwell's equations and a closed form solution is relatively easily obtained. When those constraints are relaxed, one needs to turn to completely general field solver packages (like Ansoft's HFSS) which solve Maxwell's field equations (usually approximating them using a numerical method, like Method of Moments – MoM). HFSS can take all night to simulate a simple array due in large part because HFSS must derive every performance nuance from the specifics of the antenna geometry described to it.

In this work we have combined a number of existing analysis and optimization techniques into a unique approach toward solving BEI electrode (or antenna) networks. We use an alternative analysis approach based on estimating these complex and arbitrary array compositions quickly and realistically by skipping the intense process of generating each element's radiation pattern from individual EM moments. Instead we reconstruct (e.g. estimate) these values from a set of radial-uniform radiation power density measurements - identical to the type of measurements reported in commercial antenna datasheets and those provided by rigs in anechoic antenna test chambers or the gantry setup we built for electrode water-immersion testing. Using real measurements sets up a hardware-in-the-loop simulation/emulation model offering faster development conversion to real hardware and potentially better accuracy all at a faster processing rate (current results take Matlab less than two minutes to render a on desktop PC). Further this process is of fixed complexity where complexity becomes largely independent of every design space parameter except the number of antenna elements in the array. This approach, in our current implementation, can explore 625 unique array configurations, for an array of two elements, in under two seconds.

The limitation to this approach is that, without an analysis of the antenna geometry, near-field coupling can not be accounted. However, the phase velocity of electromagnetic phenomenon – of which, quasi-static electric field conduction

currents are a part – decreases as a function of the permeability. For seawater, with a permeability near 80 [Rho07], the propagation velocity is two orders of magnitude slower than in air (equation 10.3).

$$v_\phi = \frac{c}{n_r} \tag{10.3}$$

This lower velocity results in an equivalent reduction in the near-field radius and thus supports our assumption of independent array element consideration.

10.4 Arbitrary Array Analysis

Other difficulties occur because numerical analysis of EM structures proffers only a bottoms-up approach in which the simulation must be repeated for each parameter in the array antenna design space over which one iterates. Even for mild design spaces (four elements, uniform type) this can produce enormous computational burdens to search even small fractions of the possibilities. Array systems offer no less than four independent configuration control parameters for each element in the array: (a) position, (b) phase, (c) amplitude, and (d) antenna pattern. Further, one can choose the number of elements in the array independently.

Available analysis tools divide into two broad categories: (a) closed form solutions and (b) Maxwell-based field solvers. The former category is used to handle homogeneous arrays (arrays composed of identical individual elements) in topologies that consist of uniform element spacing and uniform phase progress between successive elements. In these limited situations the symmetry and uniformity of the design permits significant simplification of Maxwell’s equations and a closed form solution is relatively easily obtained. When those constraints are relaxed, one needs to turn to completely general field solver packages (like Ansoft’s HFSS) which solve Maxwell’s field equations (usually approximating them using a numerical method, like MoM). HFSS can take all night to simulate a simple array due in

large part because HFSS must derive every performance nuance from the specifics of the antenna geometry described to it.

In our work we have combined a number of existing analysis and optimization techniques into a unique approach toward solving FUA-based array antenna networks, vitally important to so many systems where imaging, localization, and noise/jamming rejection are required features. We use an alternative analysis approach based on estimating these complex and arbitrary array compositions quickly and realistically by skipping the intense process of generating each element's radiation pattern from individual EM moments. Instead we reconstruct (e.g. estimate) these values from a set of radial-uniform radiation power density measurements - identical to the type of measurements reported in commercial antenna datasheets and those provided by rigs in anechoic antenna test chambers. Using real measurements sets up a hardware-in-the-loop simulation/emulation model offering faster development conversion to real hardware and potentially better accuracy all at a faster processing rate (current results take Matlab less than two minutes to render a on desktop PC). Further this process is of fixed complexity where complexity becomes largely independent of every design space parameter except the number of antenna elements in the array. This approach, in our current implementation: the Far-field Arbitrary Array Antenna Simulation Toolbox (FAAAST) for Matlab, can explore 625 unique array configurations, for an array of two elements, in under two seconds.

10.4.1 From an RPD Point Cloud to an RPD Model

Radiation Power Density (RPD) measurements taken from an anechoic test chamber or an antenna datasheet consist of 3-tuples defined in spherical coordinates, $\langle \theta_i, \phi_i, P_i \rangle$, where θ_i , the elevation angle, and ϕ_i , the azimuth angle, define the look-direction from which P_i , the power, was observed. In common practice, this information is displayed on two polar plots, each in a plane orthogonal to the other,

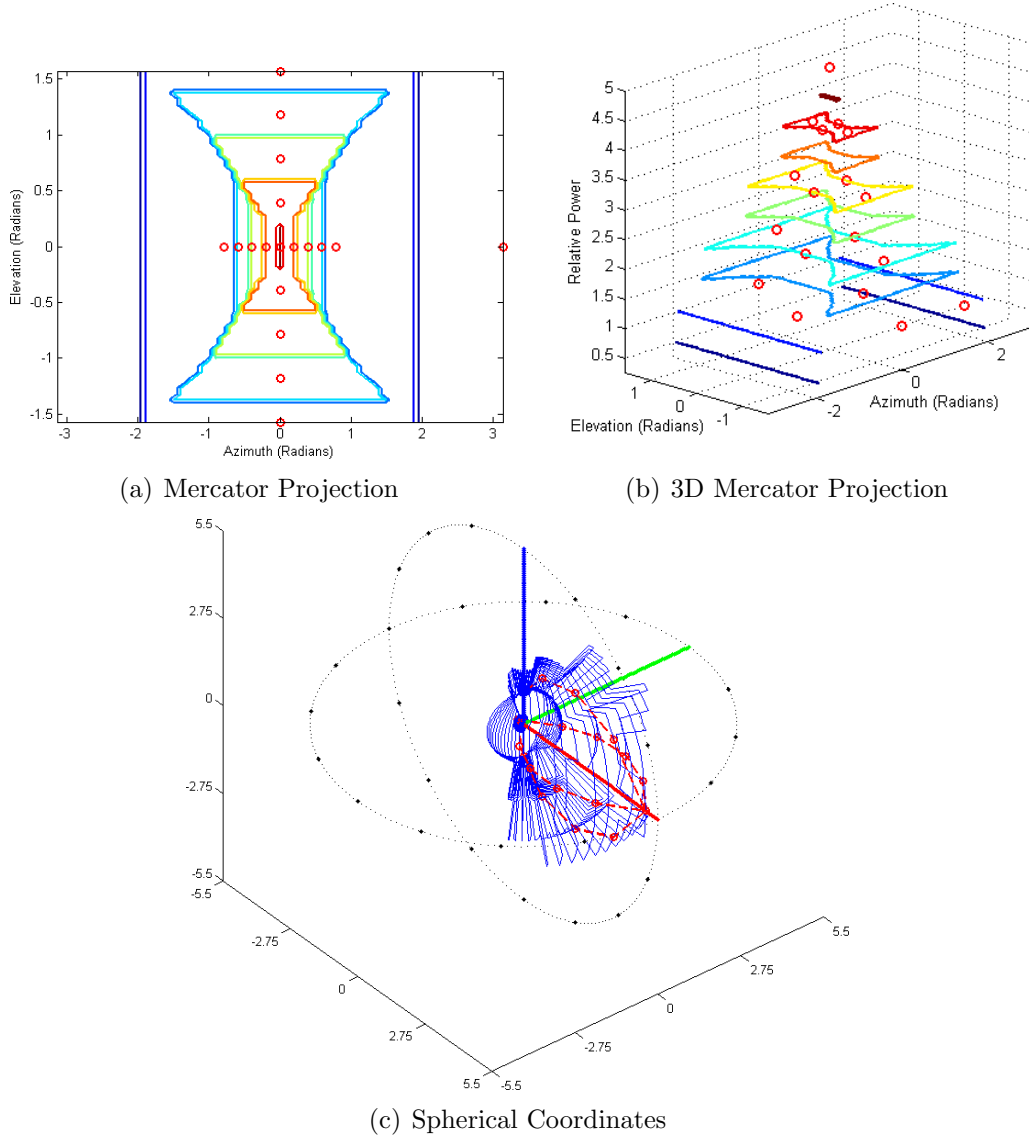


Figure 10.3 – *The MPED algorithm predicts radiation power density in an unknown direction by selecting the closest known value in a Mercator projection. Here an example dataset is shown as a flat Mercator projection (a), a Mercator projection with relief (b), and in the traditional spherical coordinates (c). Circle data markers indicate known points.*

that both include the vector extending toward $\max(P)$. Our notation demarcates each triple with a unique index i within the set of triples defining the RPD of the antenna. For n RPD measurements, we define a set $I = \{1, 2, 3, \dots, n\}$ for which $i \in I$. We can then re-express the set of RPD observations, $\mathcal{P}_{meas}(i)$, in terms of

measurements taken for each unique index i and for each parameter θ , ϕ , or P . This notation helps illustrate how we may decompose the measurements into three distinct sets interpolate among them and then recompose the measurements back into valid tuples since the index i remains in common among the elements of a given 3-tuple. It is also important to clarify that θ is defined only on the interval $[0, \pi)$, ϕ is defined only on the interval $[0, 2\pi)$, and $P \in \mathbb{R}$, $P \geq 0$.

Conceptually, we may consider \mathcal{P}_{meas} as a point cloud. As such, it may only describe the RPD for the antenna at a few discrete points. In order to achieve a complete RPD definition of the antenna – and, hence, a usable model – we must develop an algorithm to extrapolate. The approach here is to use a Mercator projection to unwrap the antenna’s far-field observation sphere onto a 2D map with accurate lines-of-bearing. Using just the angle information from the set of measured points, we can mark the locations contained in \mathcal{P}_{meas} . To determine the value of P when its angle $\langle \theta, \phi \rangle$ is not included in \mathcal{P}_{meas} , select the closest angle that is included in \mathcal{P}_{meas} and use its value of P for the unknown angle. This approach relies on minimizing the Euclidean distance between the two angles when those angles are represented in a projection as points in a Cartesian space. Accordingly, we will refer to this algorithm as the Minimized Projected Euclidean Distance (MPED). MPED can be made exceptionally fast since it relies on selection rather than interpolation to resolve unknown coordinates in the model.

To explore the MPED algorithm more formally, we will need to extrapolate from the definition of the argmax and argmin functions to define the more general arg function. In our usage, $\arg_Z(\cdot) \equiv \underset{i \in Z}{\arg}(\cdot)$ returns a set of i values (each $i \in Z$), for which each $\mathcal{P}_{meas}(i) \equiv \langle \theta_i, \phi_i, P_i \rangle$ meets the specified conditional (\cdot).

set J (equation 10.4) covers the azimuth distance relationships. Set K (equation 10.5) covers the elevation distance relationships. Finally, set L (equation 10.6) finds the optimal selection from the global minimum of both azimuth and elevation. To summarize our notation one should observe that set L is a subset of K , which

is a subset of J (and not necessarily a proper one), which, although not actually required, should be a proper subset of I . Symbolically, $L \subseteq K \subseteq J \subsetneq I$.

$$J = \left\{ \begin{array}{ll} \arg_I (\theta_i > 2\pi + \theta - \Delta_\theta) & \text{if } \theta - \Delta_\theta < 0 \\ \cup \arg_I (\theta_i < \theta + \Delta_\theta) & \\ \arg_I (\theta_i < \theta + \Delta_\theta - 2\pi) & \text{if } \theta + \Delta_\theta > 2\pi \\ \cup \arg_I (\theta_i < 2\pi) & \\ \arg_I (\theta_i > \theta - \Delta_\theta) & \text{otherwise} \\ \cap \arg_I (\theta_i < \theta + \Delta_\theta) & \end{array} \right\} \quad (10.4)$$

$$K = \left\{ \begin{array}{ll} \arg_J (\phi_i > 2\pi + \phi - \Delta_\phi) & \text{if } \phi - \Delta_\phi < 0 \\ \cup \arg_J (\phi_i < \phi + \Delta_\phi) & \\ \arg_J (\phi_i < \phi + \Delta_\phi - 2\pi) & \text{if } \phi + \Delta_\phi > 2\pi \\ \cup \arg_J (\phi_i < 2\pi) & \\ \arg_J (\phi_i > \phi - \Delta_\phi) & \text{otherwise} \\ \cap \arg_J (\phi_i < \phi + \Delta_\phi) & \end{array} \right\} \quad (10.5)$$

$$L = \{i_1, i_2, i_3, \dots\} = \underset{i \in K}{\operatorname{argmin}} ((\theta - \theta_i)^2 + (\phi - \phi_i)^2) \quad (10.6)$$

Although (10.6) treats the spherical coordinate tuples as if there were located in a 2D cartesian space, the Euclidean minimized distance approach is still valid.

Equation (10.6) will return L , a set of indices for all points matching the closest angle. Typically this just one point. However, in the case of a redundant dataset or perfectly spaced data points with respect to the desired test point, we need to settle the tie. We choose, arbitrarily, to always pick $L(1)$ – the first result.

$$P(\theta, \phi) = \mathcal{M}_P(L(1)) \quad (10.7)$$

Equation 10.8 is an expression of the Radiation Power Density (RPD) – the energy contained in an electromagnetic wave per unit area of incidence. $W(\theta, \phi)$ is the RPD, also called the Poynting vector, and takes units of $Watts\backslashmeter^2$. When we assume that the measurements are taken only in the far-field and consist of energy propagating in only one mode, the fundamental transverse electromagnetic, equation 10.8 simplifies as shown by the arrow.

$$W(\theta, \phi) = E_m(\theta, \phi) \times H_m(\theta, \phi) \Rightarrow \frac{|E_m(\theta, \phi)|^2}{\eta} \quad (10.8)$$

This process can be further optimized by sorting the matrices. In our nomenclature, the sorted versions are $\mathcal{M}_s(\theta, i)$ and $\mathcal{M}_s(\phi, i)$. Where the intuitive definition holds:

$$\mathcal{M}_{\theta,s}(i) = \text{sort}[\mathcal{M}_\theta(i)] ; \mathcal{M}_{\phi,s}(i) = \text{sort}[\mathcal{M}_\phi(i)] \quad (10.9)$$

The appeal is that searches for inequalities over ordered sets can be performed in $O(n \log N)$ time because they are fundamentally binary searches. We only have to do this once. For all future operations we assume that we have the sorted version and the original available. This means that the $O(\text{xxx})$ time for the sorting operation can amortize over every use of the model where its final cost approaches zero.

10.4.2 Optimizing RPD Models to Speed Analysis

At first glance, it seems awkward to begin the discussion of our antenna analysis tool with an algorithm that creates a set of points on the surface of a sphere. In practice, however, this operation is required multiple times in every visualization




Figure 10.4 – *Plot of the azimuthal distribution using the conventional double-integral distribution (left) and our twisted-quadratic distribution (right). The plot on the right contains 33 fewer points in total, yet looks denser. Its distribution is 320% more uniform by our D metric definition.*

procedure, in beamforming, and in array optimization. Producing a uniform distribution of points minimizes redundant analysis, or expressed conversely, for a given number of operations a uniform distribution of points yields the highest fidelity answer possible in that timeframe. Consequently, we endeavor early in this work to lay a solid foundation for efficient analysis.

It is surprisingly non-trivial to create a “uniform” distribution of points on the surface of sphere, in part, because even the definition is open to interpretation. Several valid interpretations exist and they lead to different distributions of points. There exists a set of trivial solutions to the classic geometric notion of uniform surface distribution - namely, that all distances between distinct points on the surface are equal. For a sphere you can have at most four points - the corners of the largest contained regular triangular pyramid. This definition does not produce usefully dense distributions. The most common definition for the sphere is based on the generally accepted notion that the vertices of a regular polygon are equally spaced in S^1 . One can therefore ask that every point have some property, such as being equidistant from its nearest neighbor.

Indeed, the very definition of the surface area of a sphere, the surface integral over π radians of elevation and 2π radians of azimuth is based on tracing concentric circles at different altitudes.

$$\int_0^\pi \int_0^{2\pi} r^2 \sin(\theta) d\theta d\phi \quad (10.10)$$

Distributing points in this manner results in points equally spaced along the equator and at every elevation angle prior and subsequent. While numerically convenient in spherical coordinates, such a distribution will cluster points at the poles and create longitudinal vanes concentrating points along these spars (fig.10.4, left).

It is more helpful to measure uniformity not point-by-point but by considering an aggregate measure of the dispersion. Therefore, we define a function D as the minimum distance between any pair of distinct points in the set. This provides a metric to evaluate the quality of the distribution we've achieved. Consider four points equally spaced on the equator in comparison with four points on the regular interior pyramid.

$$\begin{pmatrix} D_{\text{equator}} \\ D_{\text{pyramid}} \end{pmatrix} = \begin{pmatrix} \sqrt{2} \\ \frac{2}{3}\sqrt{6} \end{pmatrix} \approx \begin{pmatrix} 1.41 \\ 1.63 \end{pmatrix} \quad (10.11)$$

We choose this D metric since it naturally favors distributions consistent with the electrostatic potential distribution of points on a sphere were those points taken to be electrons. In general this statement of the problem guarantees a solution exists. Borrowing from an analysis by David Rusin of Northern Illinois University, since D is any continuous function of N points on the sphere and since the sphere is compact, so is the Cartesian product of N copies of it, so that D is a continuous function on a compact set, and so achieves a maximum and a minimum. Given the existence of a maximum there is the question of how one would compute

an optimal configuration. The problem can be solved with Lagrange multipliers, but this is hopelessly cumbersome for any but the smallest values of N and the simplest functions D . Moreover, some of the most natural functions D to take are not differentiable everywhere evading Lagrange analysis.

We instead have considered a novel approach that we will call the Twisted-Quadratic Distribution (TQD). We retain the computational simplicity of the conventional surface integral approach but distribute the points sinusoidally with respect to elevation. Further, we propose a rotating phase relationship between adjacent elevations.

The TQD is comprised of two equations, one to determine the number of points equally distributed for constant elevation angle theta and one to determine the azimuth for the first point in that ring.

$$Points(\theta) = \left\lceil \frac{r \sin(\theta)}{\sum_{k=0}^{N-1} r \sin(\frac{\pi}{N-1}k)} N^2 \right\rceil \quad (10.12)$$

$$Phase(\theta) = \frac{\theta}{\frac{\pi}{N-1}} \frac{\Theta}{N-1} = 2\theta \quad (10.13)$$

Where r is the radius, theta is the elevation angle measured descending from the $+z$ axis = 0° , Θ is the total angle through which the sphere elevations are twisted, and N is the total number of elevations to consider. The choice of floor or ceiling function is arbitrary where the former will produce a point cloud with a total point count slightly less than our intended N^2 and the latter being slightly over the target value.

As one can see in figure 10.4, the TQD produces a similar concentric circle layout, however the TQD is twisted completely once around when tracing from top to bottom on the sphere. Comparison of the two plots (which show the azimuthal distribution) reveals that despite the TQD containing 33 fewer points in total it

looks denser. Numerically, its distribution is 320% more uniform by our D metric definition.

Array composition analysis is performed by computing the impact of each virtual interferer on each element of the array. Therefore, a distribution which produces greater uniformity can provide the same analytical capability with fewer points in the cloud. This relationship is roughly linear allowing our TQD based analysis a better than $3\times$ runtime improvement without penalty to the fidelity of the results.

10.4.3 Determining the Array RPD Model by Compositing

Having distributed a set of points on the far-field sphere we propose the following technique to determine the array solution. However, the specifics of this step in the process are unique to the propagation method as discussed in previously. In the case of radiation (electromagnetics) phasors are applied, but in the case of electrostatics amplitudes must be considered exclusively when determining the electrode pattern.

10.4.3.1 Electromagnetics

In the case of electromagnetics, we choose one of the TQD's far-field points to represent our desired look direction and remove it from the set of remaining points. For a given test point (the origin of the interferer direction vector) located at a far-field position in 3-space determine the distance, d , and direction (in the spherical coordinate system - e.g. two angles: elevation and azimuth) between the test point and each element of the array. For each element of the array, determine $Me^{j\psi}$ by:

$$Me^{j\psi} = \frac{|E_m(\theta, \phi)|^2 A_{rx} I_m e^{j(dk + \delta_m)}}{4\pi\eta d^2} \quad (10.14)$$

where $E_m(\theta, \phi)$ is the best-fit radiation power density from the set of points provided and in the direction of the test point (θ, ϕ) , where A_{rx} is the receiver aperture (a constant scalar term to calibrate the results against a reference antenna – substitute unity for a relative power level result), where k is the wavenumber, where I_m is the gain assigned to this element by the array controller (this gain is not the gain inherent in the antenna – e.g., not a fundamental property of the antenna – rather, it is the controlled gain added by the array processor to steer the beam), and where, similarly, δ_m is the respective phase addition. This equation will produce a set of phasors composed of the interaction of the interferer vector with each element of the array. To find the cumulative effect on the array, apply superposition (phasor-addition).

To illustrate this visually consider figure 10.5. It features an arbitrary heterogeneous array composed of four elements with different radiation patterns. Each of the elements radiates with different amplitude phase and intensity for any given direction. From the left the patterns are summarized as: (1) a uniform distribution, (2) a primary lobe radiator, (3) a random distribution, and (4) a fan pattern. The elements are shown using a hybrid coordinate system consisting of cartesian element origins and spherical power densities. Note: the near-field interactivities are ignored.

The FAAAST analysis procedure can consider arrays composed of arbitrary size, placement, heterogeneous antenna types, and element weighting. Shown here (figure 10.4.3.1):

The RPD distributions of the four different elements in the example array of figure 10.5, when weighted and combined, produce the RPD for the array. When taken in far-field and shown in a 3D isometric perspective figure 10.4.3.1(a) results. The pattern is less cluttered and more easily visualized in polar form (figure 10.4.3.1(b)), in this case taken at the equator. Observe that the array pattern bears little resemblance to any of the component individual patterns – this is the




Figure 10.5 – *An arbitrary heterogenous array composed of four elements with different radiation patterns. From the left: a uniform distribution, a primary lobe radiator, a random distribution, and a fan pattern. The elements are shown using a hybrid coordinate system consisting of cartesian element origins and spherical power densities.*

power of synthesis.

10.4.3.2 Electrostatics

The use of quasi-static fields means that field considerations are made in instantaneous time intervals and, while time-varying, are time-independent in the context of field analysis. The model proposed in section 7.1 should replace that of equation 10.14 in the analysis of electrostatic arrays.

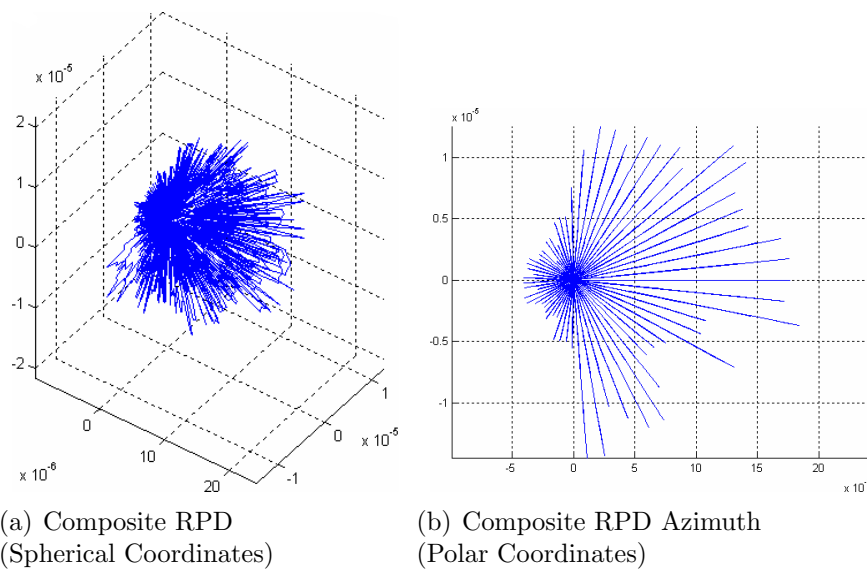


Figure 10.6 – *The FAAAST analysis procedure can consider arrays composed of arbitrary size, placement, heterogenous antenna types, and element weighting. Shown here: (a) the resulting RPD for the array taken in far-field and shown in a 3D isometric perspective, and (b) the resulting RPD in a polar azimuth perspective for clarity.*

10.5 Array Steering and Optimization

Beamforming (or beam steering) is the process of constructing interference patterns among array elements to yield a composite far-field pattern with highly directional sensitivity. It is a requisite task in array signal processing applications such as radar, sonar, acoustics, astronomy, seismology, communications, and medical-imaging. To achieve imaging of a scene, the beam must be steered, but the creation of far-field focus via interference is limited to just that – the far-field. Near-field situations, such as biomimetic electrostatics, are generally considered to be unable to benefit from array synthesis.

10.5.1 Near-field Electrostatic Beamforming

Radiative far-field synthesis works because the electromagnetic energy travels as a wave. The energy is ejected from the antenna and travels through space. As such, the point of origin of any forces exerted on charges in the environment is the current location of the energy quantum radiating through space. The origin moves because the wavefront propagates. If two coherent waves are emitted in parallel with opposite phase (180° difference) and traveling in the same direction, an object equidistant between them and collinear with the wavefront will experience zero net forces. However, if the forces were considered originating from some distance away (as opposed to collinear with the object), the net force would have a component in the direction between the wave origin and the object's location. This is the case for electrostatic charges.

All is not lost. While unable to truly cancel field energy, the number, location, polarity, and intensity of charges has an impact on the final electric force field. When evaluated differentially, as in the case of biomimetic electrostatic imaging, a surprising degree of verisimilitude with electromagnetic synthesis capabilities is achievable.

In figure 10.7, four additional charges have been added for a total of three transmit pairs. The individual transmit electrodes have a relative strength and polarity of 1, 2, -1, 1, -2, -1, left-to-right, respectively. The result is an enlarged primary lobe and reduced (relative to the primary) secondary lobes. The illustrated configuration is an example of the use of an array transmitter to increase effective range and minimize side lobes. The array transmitter provides a means to overcome the absolute amplitude limits imposed by the chemistry of the marine system as discussed in section 6.2.1 as the enlarged pattern results without increased output amplitude at any one electrode.




Figure 10.7 – *An example of the use of an array transmitter to increase effective range and minimize side lobes. The individual transmit electrodes have a relative strength and polarity of 1, 2, -1, 1, -2, -1, left-to-right, respectively. The result is an enlarged primary lobe and reduced (relative to the primary) secondary lobes.*

Array focusing of electric fields in the near-field via phasor interference is not possible. However, a more limited capability to focus the field may be achieved through field complication (adding more charges to the field). In figure 10.8, the six charges from the previous figure have been reweighted to a new relative polarity/amplitude (left-to-right) of: 0.1, 0.1, -1, 1, -0.1, -0.1.

The effect is quite encouraging. The side lobes have been reduced by approximately 50%, while the primary lobe has been reduced in amplitude by only 14.4%. As expected, unfortunately, little far-field focusing has been achieved with the reference beamwidth of 61.2 degrees declining little to 61.0 degrees when array synthesis was performed.

The effect of this particular array weighting is covered in figure 10.9 the data being the difference between the reference and focused beam patterns. The side lobes have been strongly attenuated while the close end of the hyperbolic null region has been filled in by additional field strength. Despite limited ability to reduce beamwidth, the array still provides benefit. Comparing the three figures (figures 9.8(a), 10.8, and 10.9), it is clear that the array weighting vector plays an extremely important role. Discovering the optimal weighting vector is the subject of the final component of this chapter and represents the last piece of the FAAAST tool.

10.5.2 Beam Steering

The standard data-independent beamformers include the delay-and-sum approach as well as permutations that vary amplitude in addition to the delay - or even more complex schemes to deal with the notion of delay when the signal must be phase shifted at multiple frequencies [LS05]. Collectively, this choice of a set of gains and phases to apply to the individual elements of the array is called the weighting vector.




Figure 10.8 – *Focusing of the field via complication. The six charges from the previous figure have been reweighted to a new relative polarity/amplitude (left-to-right) of: 0.1, 0.1, -1, 1, -0.1, -0.1. The side lobes have been reduced by approximately 50%, while the primary lobe has been reduced in amplitude by only 14.4%.*

Data-dependent (sometimes called adaptive) beamformers can attain better resolution and significantly better interference rejection by choosing their weighting vector in an online manner. The array is continuously or periodically re-weighted in response to the most recent inbound dataset. Traditional adaptive beamformers are prone to self-nulling as they are known to be extremely sensitive to even the slightest mismatches of the signal and interference components (perhaps due to imprecise sensor calibration). When this mismatch occurs, the array processor interprets the desired signal component in an array's observations as an additional interfering source and, accordingly, suppresses it [Ger87] [FG94].




Figure 10.9 – *The difference between the reference and focused beam patterns (previous two figures). Note the strong attenuation in the side lobes while the close end of the hyperbolic null region has been filled in by additional field strength.*

We propose to use an approach that combines the stability of data-independent beamforming with the improved resolution of adaptive beamforming through the development of a point-mainbeam constraint. We will implement this constraint by defining a set of virtual interferers scattered, initially as uniformly as possible, over the surface of a far-field sphere and let them impinge on our array during analysis. The set of virtual interferers may then be updated at any time to reflect changes in the desired pattern.

10.5.3 Array Optimization

Finally, we have arrived at a complete array analysis procedure. By impinging all of the virtual interferers on the array we can create a matrix D with each row containing the complex phasor expressing the influence of the Virtual Interferer

located in a given direction on each element of the array.

$$D_{Matrix} = \begin{bmatrix} w_{\in_0, I_0} & w_{\in_1, I_0} \\ w_{\in_0, I_1} & w_{\in_1, I_1} \\ w_{\in_0, I_2} & w_{\in_1, I_2} \\ \dots & \dots \end{bmatrix} \quad (10.15)$$

Further, we define the matrix C with each row containing a proposed weighting configuration for the array:

$$C_{Matrix} = \begin{bmatrix} w_{\in_0, A} & w_{\in_1, A} \\ w_{\in_0, B} & w_{\in_1, B} \\ w_{\in_0, C} & w_{\in_1, C} \\ \dots & \dots \end{bmatrix} \quad (10.16)$$

The analysis is now straightforward. We choose the configuration that maximizes the Signal-to-Interference-plus-Noise Ratio (SINR).

$$INR_{Vector} = \frac{1}{N_{Vector}} \cdot [C_{Matrix} \cdot D_{Matrix}^T] \quad (10.17)$$

$$SINR = \frac{S}{INR} \quad (10.18)$$

Choose the configuration (row of the CMatrix) that has the same row index as the maximum entry of the SINR vector - that is the optimal weighting vector for the array.

10.6 Conclusion

This chapter concludes my work with the physical generation and foveation of differential electrostatic fields. Its contributions have included the introduction of concepts, the application of biomimetics to field-foveation, and the development of tools to optimize and exploit the relevant mechanisms. It has been demonstrated that, despite being a near-field phenomenon, differential electrostatic fields may be guided through the use of array elements and synthesis. The following chapters proceed forward from the capabilities introduced here. Given that the field may be manipulated to produce desirable patterns, the following chapters discuss what should be done with those patterns – that is, how can additional information being encoded in the wave shape to assist in the tasks of localization, navigation, and communication so important to biological organisms and engineered oceanic applications.

Part IV

Waveform Encoding

CHAPTER 11

Bio-inspired Informatics

11.1 Overview

This chapter begins a series of chapters dedicated to the idea of learning the position of objects in our structured emitted field by encoding features into the waveform that, once disturbed by the object, impart additional knowledge as to the object's whereabouts. We will rely on the work of previous chapters to provide a foundation of capabilities to which we add this final layer. In continuing with our desire to learn from ocean-going biological organisms, broadly, I entitle our forthcoming efforts at channel encoding *bio-inspired informatics*. Within this topic area, we will develop several bio-inspired technologies that continue the discussion in RADAR systems terminology.

In the case of RADAR, objects (targets) are reflective. In the case of electrostatics, the objects are different from the background in their conductivity. For many objects, these two phenomena are correlated. We will interchange the two concepts based on the context of the application presently being discussed (RADAR or oceanic electrostatics). It is left to the reader to envision the alternate scenario when only one is presented.

11.2 Bio-inspired Informatics

In a telling experiment conducted and reported by [TZ05], two Brown Ghost Knifefish (*Apteronotus Albifrons*) were placed at opposite ends of a shared tank with two barriers placed between them – the middle section being void of either fish. In the middle section was a plexiglass tube of sufficient size and shape to be considered by both animals as a suitable shelter. *Apteronotus albifrons* is a nocturnal fish that is competitive in the case of limited resources (in this case shelter) and is inherently uncomfortable in open water [TZ05]. It is also weakly electric – transmitting a field between 500Hz - 1000Hz almost continuously for navigation and hunting purposes.




Figure 11.1 – A spectrogram from a telling experiment conducted and reported by [TZ05] in which two Brown Ghost Knifefish (*Apteronotus Albifrons*) were placed at opposite ends of a shared tank. Reproduced from [TZ05].

A recording electrode and a reference electrode were inserted into the tank at opposite ends and the voltage recorded over time. Figure 11.1 is a spectrogram recording of the experiment. With the lights off in the room and the animals safely separated, they select non-interfering frequencies. After ten minutes the barriers

were removed and the animals move into close proximity, realize the inherent contention for the mid-tank shelter, and begin to exhibit hostile behavior. In the quasi-static domain, this hostility manifests itself in a jamming behavior. The smaller animal increases its output frequency to nearly that of the larger animal. At twenty minutes the lights are switched on and the nocturnal animals withdraw and reduce their jamming behavior. There are still periods within the lights-on interval where they elect to compete. At thirty minutes, the lights are switched off again and the fight resumes. Spectrally, in the last ten minutes the smaller animal, weakened from the continuous use of its electric organ at elevated output frequency struggles to maintain its jamming behavior as its output frequency varies rapidly in its ongoing effort to approach that of the other fish. There is a general downward trend in frequency of both animals corresponding to exhaustion.

Of particular interest to us are the innate capabilities this interaction demonstrates. The electric fish appear to possess the ability to operate in the frequency domain. They can detect and isolate specific frequencies, change frequency on demand, avoid interference by changing frequencies, and jam and track frequencies when desired. To confirm this, the authors of [TZ05] conducted a follow-on experiment in which the second animal was removed from the tank and a pair of silver electrodes were inserted in its place. The electrodes emitted a quasi-static field at 920 Hz continuously. In the dark, with the second animal absent, the fish in the tank exhibited the same jamming response as in the previously described experiment electrically attacking the signal generator's electrodes.

This second experiment is shown in figure 11.2. The upper panel of the figure is the frequency domain and the lower figure the time domain. The aperiodic recurring broadband spikes in the fish output are recording artifacts introduced when the fish actually bit the silver electrodes in the tank. The frequency of biting increases as the fish brings its output frequency ever closer to the signal generator's. In the time domain the close proximity of the two frequencies manifests as an

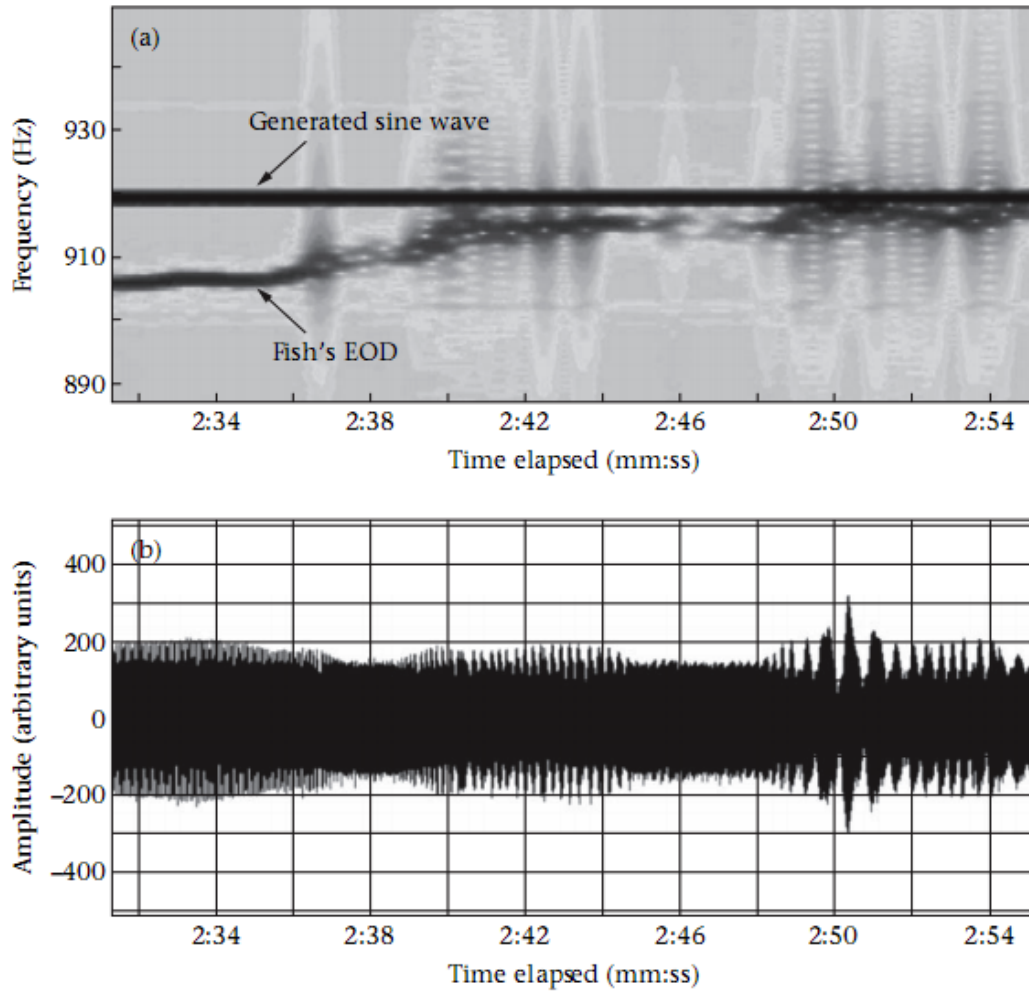


Figure 11.2 – Spectrogram (top) and time-domain (bottom) of a Brown-Ghost Knifefish interacting with an artificial electrode pair. Reproduced from [TZ05].

amplitude envelope at the frequency difference between the two – clearly visible in the top and bottom profiles of the time domain signal. This is classic interferometry: two waves at the same frequency, no matter the phase, produce a wave at the same frequency. Two waves at different frequencies, no matter the phase, produce a wave with an envelope at the frequency difference between the two. This is illustrated in figure 11.3.




Figure 11.3 – (Top) *Two waves at the same frequency, no matter the phase, produce a wave at the same frequency.* (Bottom) *Two waves at different frequencies, no matter the phase, produce a wave with an envelope at the frequency difference between the two.*

11.3 Angle-of-Arrival-Assisted Relative Interferometry

Inspired by the Brown Ghost Knifefish’s ability to foveate, interfere with, detect, and track quasi-static electric signals we propose a new signal encoding technique that makes use of Angle-of-Arrival (AoA) information (direction/foveation) in conjunction with local interferometry (jamming response) to improve target location estimation. We call this process Angle-of-arrival-assisted Relative Interferometry (ARI¹). In the context of RADAR systems design, it has a number of desirable attributes including the ability to reduce the synchronization, network, and hardware requirements when operating as the sole RADAR modality and its ability to augment existing pulsed and CW RADAR algorithms. In oceanic electrostatics it may very offer a path to higher resolution imagery and target decluttering.

¹ARI is pronounced “are-eee”.

11.4 Conclusion

In the chapters that follow the Bio-inspired waveform encoding technique of ARI will be explored in detail. Chapter 12 covers the simplest form of ARI based on a continuous-wave transmission. Having demonstrated the approach and its utility it is incumbent upon us to demonstrate its application. To this end, we will demonstrate a Software Defined Radio (SDR) approach to an Angle-of-arrival-assisted Relative Interferometric (ARI) RADAR transceiver. A software-defined model in Matlab is developed and implemented. The model is used to explore the implementation parameters and demonstrate the stability of the scheme in the face of clock jitter. At the end of the chapter a SDR-ARI transceiver is designed, simulated, and implemented utilizing USRP hardware and gnuRadio software.

In the following chapter, chapter 13, we propose a Pulse-Wave (PW) extension to the ARI technique to dramatically reduce the scan-time and the number of vantage points necessary to obtain high-fidelity target position estimation. Accordingly, we call this enhanced process PW-ARI. PW-ARI is the fusion of data from three domains: time (time-of-flight), phase (relative phase-of-arrival), and angle (direction-of-arrival). It has a number of desirable attributes. Foremost among these are its ability to rapidly image extremely large volumes with high accuracy and dense target clusters, support an infinite number of friendly submarine or surface vessels while providing stealth operational support, and present a spectrally small footprint to both hostile (evade detection and counter-measures) and friendly (maximize electromagnetic compatibility) systems.

CHAPTER 12

Continuous-Wave Angle-of-Arrival-Assisted Relative Interferometry

12.1 Overview

Broadly, the concept of target presence and position discovery is fundamental to military operations. The ability of Radio-based Detection and Ranging (RADAR) systems to provide early warning of enemy aircraft or warships has, indeed, changed the course of history [LS05]. In this chapter we propose a new technique that makes use of Angle-of-Arrival (AoA) information in conjunction with local interferometry to improve target location estimation. Accordingly, we call this process Angle-of-arrival-assisted Radio Interferometry (ARI, pronounced “Are-eee”). As we will show, it has a number of desirable attributes. Foremost among these are its ability to reduce the synchronization and hardware requirements when operating as the sole RADAR modality and its ability to augment existing pulsed and Continuous-Wave (CW) RADAR algorithms.

In our nomenclature, targets (and their corresponding applications) come in three varieties: (a) cooperative, (b) non-cooperative, and (c) hostile. Cooperative targets benefit from collaboration with the RADAR infrastructure (e.g. a GPS receiver or friendly aircraft transponder), non-cooperative targets do not assist in their own detection and localization, but do not actively attempt to prevent it (e.g. mine hunting operations), while hostile targets seek to purposefully deny their existence and location to the RADAR system (e.g. Electronic Counter-

Measures/jamming). In this work we will focus on two archetypical applications. Of type (a), a GPS-independent navigation platform is considered. ARI beacons could be deployed as needed to fill GPS coverage gaps in dense urban areas and provide an alternative when GPS is actively denied. The terrestrial nature of the envisioned ARI beacons places them substantially closer to their users than space satellites and their AoA transmission scheme limits the receiver sensitivity to interfering signals. Further consideration of the hostile aspects – type (c) – is reserved for future work. In type (b), we will consider the detection of emplaced metal objects as it may apply to Anti-Mine Warfare (AMW) when locating metal in the ground or to the emergency medical care of blast and gunshot victims when locating metal in the human body. In both situations rapid detection and localization are critical.

12.2 Angle-of-Arrival Assisted Radio Interferometry

An Angle-of-arrival-assisted Radio Interferometry (ARI) based system consists of transmitting beacons and passive receive-only objects each of which is equipped with a steerable directional antenna. Typically in class (a) applications the beacons and objects are physically separate units operating at considerable distance, whereas in class (b) applications the beacons and objects might be almost co-located and appear to an operator as a single piece of equipment. Each of the beacons knows its location and orientation. Each beacon transmits a unique pair of frequencies creating a beat envelope in the transmitted signal which is not unique to any beacon. The receiver can then identify scattering targets in the intermediate environment (see 12.4) or its own position (see 12.5) through the intersection of the vectors formed from the known beacon locations and the set of angles, look directions, of the directional antennae.

This initial AoA-only estimate is then refined by searching in the volumetric

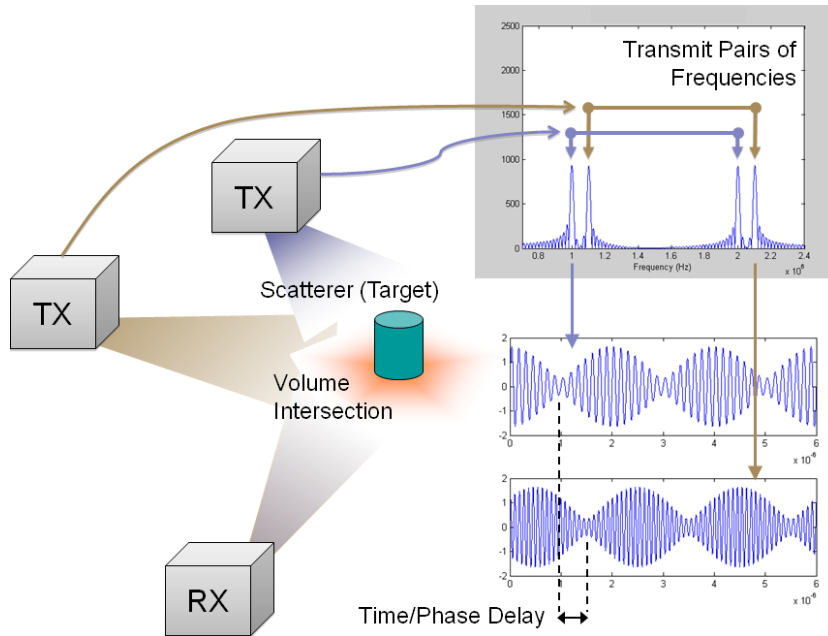


Figure 12.1 – *An Angle-of-arrival-assisted Radio Interferometry (ARI) based system consists of transmitting beacons and passive receive-only objects each of which is equipped with a steerable directional antenna. The beacons transmit unique pairs of frequencies creating a beat envelope which is not unique to any beacon.*

space defined by the intersection of the beamwidths of each beacon using the relative distances between beacons to disambiguate. This difference in path length is known through the phase difference of arrival of the envelopes from each beacon (or time difference since the envelope frequency is common to all beacons). The directionality of the receiver serves to reduce the possible search space of positions inferred from the TDOA data and to help with jamming and reflection suppression. Since the approach does not rely on signal strength, small amplitude reflections from the same direction of incidence (e.g. back reflections), differences in transmit or scattering power, and moderate frequency-specific path loss do not introduce error into the position estimation.

The ARI approach offers a number of unique advantages. Since, all of the required transmitter timing information is encoded in the signal, an object's receiver does not need to maintain phase coherence with the time source of any

beacon. Further, the beacons are immune to phase coherence errors in their carrier needing only to maintain sync amongst beacons at the modest frequency of their beat envelope. The directive nature of ARI assists in disambiguating phase-based distance relationships which are, otherwise, only disambiguous for the region $\langle -\pi/2, \pi/2 \rangle$ and, adding orientation knowledge, allows 3D localization from a mathematical minimum of just two beacons and one object. As the objects are passive, an ARI based navigation system would scale, fitting an infinite number of objects in finite bandwidth while protecting their presence and identity from observation.

12.2.1 Nomenclature

For clarity, we will use f_{x_n} to refer to individual frequencies, where x is a generic transmitting beacon, a specific transmitting beacon ($x = a, b, c\dots$), or a receiving object ($x = r$). n is the frequency designator (1, 2, or c for carrier) for the ARI component indicated by x . We do not designate objects individually since the objects (receivers) are passive and may be considered independently.

12.2.2 Transmission

The transmitted signal from each beacon consists of two frequencies – one that is unique to the beacon, f_{x_1} , and one that is shared across all beacons, f_{com} . A Hartley modulator architecture [Dar70] is used to produce the following signal over the air:

$$\begin{aligned} & \sin(2\pi f_{x_c}) \left(\sin(2\pi f_{x_1}) + \sin(2\pi (f_{x_1} + f_{com})) \right) + \\ & \cos(2\pi f_{x_c}) \left(\cos(2\pi f_{x_1}) + \cos(2\pi (f_{x_1} + f_{com})) \right) \end{aligned} \quad (12.1)$$

Algebraically we see that (12.1) resolves to just two frequencies: $f_c + f_{x_1}$ and $f_c + f_{x_1} + f_{com}$ [Dar70]. This is a Single Side Band-Suppressed Carrier (SSB-SC)




Figure 12.2 – *The proposed superheterodyne down-converter and signal-processing blocks for the ARI receiver.*

transmission. Since this is the absolute minimum amount of spectrum required to convey all of the information required for ARI we are as spectrally efficient as possible.

12.2.3 Reception

The distance information required to localize a target is encoded in the phase of the signal at f_{com} . However, two transmitters on the same frequency will interfere and obscure the identity and phase of each. ARI supports multiple simultaneous transmissions without interference, because the common frequency information is actually present on multiple unique frequencies ($f_{x_2} = f_{x_1} + f_{com}$, for $x = a, b, c...$). Accordingly, we designate f_{com} as the baseband signal and task the ARI receiver with recovering it from each beacon.

Although the transmitted signal is up-converted in quadrature pursuant to [Dar70], the received signal does not require down-conversion in quadrature – a single conversion superheterodyne receiver is sufficient (see figure 12.2). The spectrum of the resulting intermediate signal contains the identities of any beacons transmitting towards the object at its current look direction. Beacons may be identified by the presence of a signal at f_{a_1} , f_{b_1} , etc and at each of these frequencies offset by f_{com} . As f_{x_1} is unique to the transmitter and the frequency offset, f_{com} , is common to all transmitters, a second unique frequency, $f_{x_2} = f_{x_1} + f_{com}$, exists for each beacon.

Multiple beacons are detected through the use of multiple beacon-specific blocks. Each contains the appropriate channel selecting filters and a pair of Variable Gain Amplifiers (VGA) that are used to equalize the amplitudes of the two frequency components (f_{x_1}, f_{x_2}). This maximizes the amplitude of the recovered f_{com} . If multiple beacons exist, the equalized output from each beacon at f_{com} is compared in a phase detection unit. In the simplest two-beacon case, the phase detector would return the signal $\Delta\phi = \phi_a - \phi_b$, the phase difference between the signals from beacons a and b at their recovered common frequency.

$$\Delta\phi = \frac{(d \bmod \lambda)f_{com}}{c} \quad (12.2)$$

This phase difference corresponds to the relative physical distance between the paths from each beacon, through the environment (may be reflected or otherwise indirect), to the object. This relationship is presented in (12.2).

12.2.4 RADAR Classification

The most common modern day RADAR is monostatic. In a monostatic RADAR the transmitting and receiving antennae are collocated (such as in the nose of an aircraft). In a bistatic RADAR the transmitting and receiving antennae are spatially separated. A multistatic RADAR typically consists of one transmitting and several spatially distributed receiving stations with joint processing of received information [Kre]. A multisite RADAR is essentially a collection of individual RADAR systems where information from each locally detected target is fused and jointly processed to create an integrated picture of the battlespace [Che]. A multistatic RADAR requires data from all of its component transmitters and receivers in order to make a single target determination, whereas a multisite radar is essentially a collection of RADARs, each capable of making its own determinations, sharing data to improve accuracy, range, or other performance characteristics. A

passive RADAR does not involve the transmitters in the controlled aspect of the system and uses existing, typically civilian broadband FM “radio” stations, as transmitters of opportunity. The value in passive RADAR systems is that the existence of the system is hidden from observers. ARI does not fit well in any of these categories. ARI is unique in that you have a plurality of transmitters (at least two) communicating with a plurality of receivers (at least one) an architecture not commonly found in other RADAR systems.

12.2.5 Prior Art

In addition to the RADAR architectures discussed previously several similar systems deserve mention. In [MVD05] a team from Vanderbilt University used interferometry among the nodes of a wireless sensor network to determine location from multiple pairs of transmitters sending CW signals (each pair had to take turns on the channel while the receiver was stationary) and being observed by a pair of receivers. By making measurements in at least an 8-node network the positions could be determined. In contrast to ARI, which does not require any time synchronization between transmitters and receivers, [MVD05]’s entire estimate of signal arrival and, ultimately, position hinges on the ability to maintain strict time sync among nodes in the network.

Similar to ARI, in [K 98] AoA information was added to a Frequency-Modulated Continuous Wave (FMCW) RADAR. In such a system range may be determined from $R = \frac{c * f_{beat}}{2f_m}$, where c is the speed of light in the medium of propagation and f_{beat} occurs from the interference of the CW transmission and its echo. The advantage of such a system is that the hardware is arguably simpler than that proposed here. However, it suffers from poorer location accuracy since range is a function of the measurement of both f_{beat} and f_m (which is itself, $f_m = \Delta f / T$, composed from the measurement of sweep time T and sweep rate Δf) and direction estimation is limited to the steer-ability of the antenna.

In 2006, the Harris Corporation undertook development of a local interference based alternative to GPS ([Rob05]) very similar to ARI in core concept, though without any notion of directional transmission/reception. Compared to this art we achieve more efficient spectrum allocation through the use of SSB-SC and frequency interleaving. We further reduce the minimum required number of spatially distributed transmitters by fusing angle-of-arrival information.

12.3 Design Considerations

Equation (12.2) reveals that the key frequency parameter is the frequency difference f_{com} rather than the two frequencies themselves. This implies that the operating frequency is not important – only the frequency difference between the two tones. It may be noted that the frequency difference may be made arbitrarily as small or as large as desired. However, once $f_{com} < f_{x_n}$ is violated, the beat frequency aliases and simpler time-peak selection or time-peak interpolation methods fail and more complicated frequency-based approaches are necessary to extract the envelope signal.

Making the frequency difference small serves to avoid phase ambiguity (phase exceeding 180 degrees) over a long distance. To illustrate this, reformulate (12.2) in terms of Δd

$$\Delta d = \frac{\Delta\phi c}{f}; d < \lambda \tag{12.3}$$

The condition, $d < \lambda$, appears because of the modulus operation in (12.2). However, since we want to relax the synchronization requirements as much as possible, the determination of phase-lead or phase-lag between the envelopes from the two beacons is only possible when $d < \lambda/2$ otherwise it can not be determined unambiguously if beacon a is leading beacon b by $\Delta\phi > \pi/2$ or trailing

by $\Delta\phi < \pi/2$. For a frequency difference of 1kHz there will be no phase ambiguity for a distance up to 150km [Rob05]. Further, making the frequency difference larger improves the sensitivity of the phase detection and hence improves performance at a low SNR (AoA improves SINR). For a frequency difference of 1MHz a distance of 100 meters results in 120 degrees of phase shift as opposed to 0.12 degrees of phase difference for a frequency difference of 1kHz at the same distance [Rob05].

12.3.1 Modelling Phase and Frequency Error

Equation (12.2) is ideal in that all of the phase measurement may be attributed to distance. An ARI instantiation will suffer the limitations of practicable hardware. For every frequency generated there exists a possible error in both phase and frequency. Here we'll solve through the complete signal path including these errors to investigate how they aggregate.

Unfortunately, we need to define a few more terms. $f_{x_{ne}}$ is the error in generating frequency f_{x_n} . With respect to phase, we use a similar notation $\phi_{x_{n\text{type}}}$, where x is a beacon ($a, b, c\dots$) or an object (r), where n is the frequency designator (1, 2, or c for carrier), and where type is e for an error term or q for the quadrature error. The q term is the signal phase error plus the phase amount by which the quadrature term deviates from its intended $\frac{\pi}{2}$. If the phase of the quadrature signal was $\phi_{x_q\text{sig}}$, then $\phi_{x_q} = \phi_{x_q\text{sig}} - \frac{\pi}{2}$.

$$f_{X_{1\text{sig}}} = f_{r_{ce}} - f_{x_{ce}} + f_{x_1} + f_{x_{1e}} \quad (12.4)$$

$$f_{X_{1\text{err}}} = f_{x_{ce}} - f_{r_{ce}} + f_{x_1} + f_{x_{1e}} \quad (12.5)$$

$$f_{X_{2\text{sig}}} = f_{r_{ce}} - f_{x_{ce}} + f_{x_2} + f_{x_{2e}} \quad (12.6)$$

$$f_{X_{2\text{err}}} = f_{x_{ce}} - f_{r_{ce}} + f_{x_2} + f_{x_{2e}} \quad (12.7)$$

The output of two beacons as received by an object and distilled down to an intermediate frequency, but not including any path-length phase change, is


$$S_{rx,IF} = S(a, 1) + S(a, 2) + S(b, 1) + S(b, 2) \quad (12.8)$$

The solution is of the form of (12.8) where, for the simplest two beacon case,


$$\begin{aligned}
S(x, n) = & \\
& \frac{1}{4} \cos \left(2\pi f_{X_{n_{sig}}} t + (\phi_{x_{ne}} - \phi_{x_{ce}}) \right) \\
& + \frac{1}{4} \cos \left(2\pi f_{X_{n_{sig}}} t + (\phi_{x_{nq}} - \phi_{x_{cq}}) \right) \\
& + \frac{1}{4} \sin \left(2\pi f_{X_{n_{sig}}} t + (\phi_{r_{ce}} - \phi_{x_{ce}} + \phi_{x_{ne}}) \right) \\
& + \frac{1}{4} \sin \left(2\pi f_{X_{n_{sig}}} t + (\phi_{r_{ce}} - \phi_{x_{cq}} + \phi_{x_{nq}}) \right) \\
& + \frac{1}{4} \cos \left(2\pi f_{X_{n_{err}}} t + (\phi_{x_{cq}} + \phi_{x_{nq}}) \right) \\
& - \frac{1}{4} \cos \left(2\pi f_{X_{n_{err}}} t + (\phi_{x_{ce}} + \phi_{x_{ne}}) \right) \\
& + \frac{1}{4} \sin \left(2\pi f_{X_{n_{err}}} t + (\phi_{x_{ce}} - \phi_{r_{ce}} + \phi_{x_{ne}}) \right) \\
& - \frac{1}{4} \sin \left(2\pi f_{X_{n_{err}}} t + (\phi_{x_{cq}} - \phi_{r_{ce}} + \phi_{x_{nq}}) \right) \quad (12.9)
\end{aligned}$$

12.3.1.1 The Implications of Frequency Error

As is evident from 12.3.1, frequency errors add frequency offset to the returned signals without affecting their phase. This can be handled by widening the bandwidth of the channel selection filter's passbands to accommodate any frequency misalignment. In effect, the frequency error has the effect of driving the signals more spread spectrum which may be desired in some applications.



(a) In the time domain, phase shift is not apparent in the data signal despite significant shift in the baseband signal.



(b) Quadrature shifting the principal carrier has less than an 8% effect in the worst case.

Figure 12.3 – *The principal ARI data signal is relatively unaffected by initial phase error in the carrier signal or in generating a quadrature variant of it.*

12.3.1.2 The Implications to Carrier Phase Errors

The ARI data signal is carried in the envelope of a channel-selected beacon transmission as seen at the receiver after a single down-conversion. As such, the envelope shows stationary phase as the component carriers phase shift beneath it. In the time domain, as shown in the upper part of figure 12.3, phase shift is not apparent in the data signal despite significant shift in the intermediate signal. In the lower part of the figure, the phase shift is explored in detail. Quadrature shifting the principal carrier has less than an 8% effect on the phase of the data signal in the worst case. Since ARI uses a phase difference these values represent the tolerable phase difference between the carriers of two selected beacons.




Figure 12.4 – *Phase error among the generation of the signal frequencies themselves has a more direct manifestation on the final phase at f_{com} . The relative error introduced by changing the phase of f_{x_1} and f_{x_2} is shown.*

12.3.1.3 The Implications to Signal Phase Errors

The generation of the signal frequencies themselves is significant. Phase error among the generation of the signals has a more direct manifestation on the final phase at f_{com} . In figure 12.4, the relative error introduced by changing the phase of f_{x_1} and f_{x_2} is shown. Further, there exists an intermediate set where the phase offsets from each component signal cancel leaving the phase of the signal at f_{com} unchanged from the 0 error term state. This set manifests as a minimum line running diagonally through the surface of figure 12.4.

12.3.2 Frequency Interleaving

Determining the optimal value of f_{com} is not trivial. A large value of f_{com} requires a wide spectrum gap between each beacon's component frequencies. This, in turn, has significant implications for the hardware design, which is generally more difficult to implement as sampling rates go up and phase margins tighten. However, for small values of f_{com} resolution suffers as shown in figure 12.5. $f_{com} > 150MHz$ is required




Figure 12.5 – *Linear resolution as a function of $f_{com}/2\pi$, or 1 degree of phase. $f_{com} > 150MHz$ is required for most room-scale applications, while navigation might accept frequencies as low as a few MegaHertz.*

for most room-scale applications, while navigation might accept frequencies as low as a few MegaHertz. Increasing the common frequency also necessitates increasing the sampling rate to stay above the Nyquist limit and increasing the value of f_{x_1} to keep $f_{x_1} \gg f_{com}$ where the signal is readily down-converted to recover the envelope.

In order to use the available spectrum most efficiently and to optimize the hardware implementation, ARI employs a frequency interleaving scheme (see figure 12.1). The f_{x_1} frequencies are all clustered into one band and their corresponding f_{x_2} frequencies into another. The use of this dual band approach is expedient, allowing a maximum f_{com} in any finite bandwidth, but it also lends itself to the use of band-specific antenna configurations which may be necessary if f_{com} is very large and the problem of wideband interference overload is present (the directional nature of ARI receivers also helps mitigate this risk).

12.3.3 Deployed Platforms and Doppler

ARI can be deployed from moving platforms with little permutation. ARI has no effect on the Doppler returns and f_{com} , despite not being a signal existing at a real frequency, will experience a Doppler shift. We can prove this simply. Begin with the Doppler equation [Bal05]:

$$f_{Doppler} = f_0 + \frac{f_0 v_R}{c} \quad (12.10)$$

From the Doppler shift, (12.10), we can substitute to:

$$\Delta f_{Doppler} = (f_{x_1} + \Delta f_{x_1}) - (f_{x_2} + \Delta f_{x_2}) \quad (12.11)$$

and from here to:

$$\Delta f_{Doppler} = (f_{x_1} - f_{x_2}) + \frac{v_R}{c}(f_{x_1} - f_{x_2}) \quad (12.12)$$

which is in the same form as (12.10) completing the proof.

12.4 Noncooperative Targets

Decreasing the size and weight of advanced military systems while maintaining the performance and bandwidth is a complex task. In 2001, the Army Communications and Electronics Command (CECOM) put out a solicitation [Cor01] to procure an anti-mine warfare system based on Forward-Looking Ground-Penetrating RADAR (FL-GPR). The Vehicle Mounted Mine Detector (VMMD) system of figure 12.6 resulted. It consisted of 30 quad-ridged horn antennas and used FL-GPR to detect buried or surface emplaced explosive devices through the changes they introduce into the RF back-scatter. Although the VMMD FL-GPR achieved a range resolution of 5cm, look-ahead to 60 meters, and depth detection to 15cm



Figure 12.6 – *Stanford Research Institute’s forward looking ground penetrating radar system prototype (Reproduced from [SL05]) for near real-time detection of surface emplaced or buried anti-tank mines. This prototype was an outgrowth of military interest which began with [Cor01].*

[KCA02], in order to achieve this capability the VMMD FL-GPR system required large bulky antenna arrays, an ultrawide bandwidth (from 0.3 to 3.0 GHz), high RF output powers, and a mounting high above the vehicle – increasing its visual and RF cross-section and decreasing its off-road maneuverability in the event of an ambush [LWL03].

To preserve stealth, reduce power and form-factor, and increase detection range ARI beacons could be placed near vital intersections, along thoroughfares, or deployed and redeployed by recon units. The ARI receiver is then mounted to the lead convey vehicle which ascertains mine presence from the back-scatter of the beacons. As in GPS, the vehicle mounted equipment is completely passive allowing stealth operation and because the ARI beacons are forward-deployed, detection range and accuracy is improved.

The need for miniaturization and forward-deployment of such technologies is

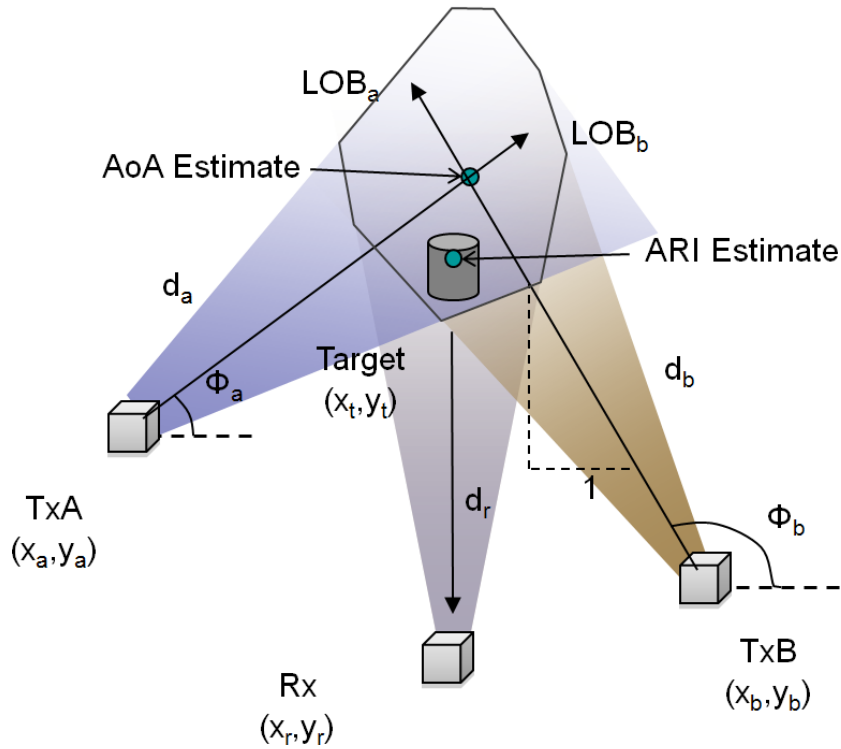


Figure 12.7 – ARI makes an initial estimate using just the intersection of the Lines-Of-Bearing (LOB) and then revises it by searching in the beam intersection volume trying to fit the position estimate to the measured phase difference of arrival.

pressing. According to [iCa08], as of March 27th, 2008, 38.7% of all US military combat fatalities in Iraq have come from Improvised Explosive Devices (IED) far exceeding the next leading cause, hostile fire, at only 13.0%. Further, an ARI-based imaging system could see use during checkpoint operations as a concealed weapon locator operating from a moderate stand-off distance.

For clarity we begin with the two dimensional case. Our two transmitters are located at $\langle x_a, y_a \rangle$ and $\langle x_b, y_b \rangle$, respectively, and they form angles ϕ_a and ϕ_b with respect to a line parallel to the x-axis. The Line Of Bearing (LOB) is the line containing $\langle x_n, y_n \rangle$ and extending in direction ϕ_n . It is determined from its slope and intercept. As shown in figure 12.7, we may define the slope using a unit distance in the x dimension and determine the corresponding y coordinate as: $\langle x_n + 1, y_n + \tan(\phi) \rangle$. The equation for an arbitrary LOB in a two dimensional

space is determined from (12.13) through (12.16).

$$y = mx + b \quad (12.13)$$

$$y = \frac{\tan(\phi) + y_n - y_n}{1 + x_n - x_n}x + b \quad (12.14)$$

$$b = y_n - \tan(\phi)x_n \quad (12.15)$$

$$y = \tan(\phi)x + (y_n - \tan(\phi)x_n) \quad (12.16)$$

This culminates in (12.16), the equation for the LOB. Setting $LOB_1 = LOB_2$ gives:

$$x_t = \frac{y_b - \tan(\phi_b)x_b - y_a + \tan(\phi_a)x_a}{\tan(\phi_a) - \tan(\phi_b)} \quad (12.17)$$

$$y_t = \tan(\phi_a)x_t + (y_a - \tan(\phi_a)x_a) \quad (12.18)$$

d_a and d_b are the line segments along LOB_a and LOB_b respectively whose length is the partial path length between the ARI transmitters and receiver. The total path length is $d_a + d_r$.

$$\Delta d = (d_a + d_r) - (d_b + d_r) = d_a - d_b \quad (12.19)$$

Given that we are in a Cartesian coordinate space, the Euclidean distance expression is valid and (12.19) may be re-expressed using (12.2) as:

$$\Delta\phi = \frac{f_{com}}{c} \left(\sqrt{(x_t - x_a)^2 + (y_t - y_a)^2} - \sqrt{(x_t - x_b)^2 + (y_t - y_b)^2} \right) \quad (12.20)$$

where $d_{max} < \lambda/2$

The system of three equations formed from (12.17), (12.18), and (12.20) is over-constrained and may be weakened by recognizing that ϕ_a and ϕ_b may, in practice, represent any angle in the current beamwidth of their respective transmitters. Figure 12.8 shows how the ARI solution, given an initial estimate from the angular intersection of arrival, can estimate target position to greater fidelity than AoA alone since for practical systems the time resolution far exceeds the angular resolution for a steerable beam antenna. In the figure, $f_{com} = 10MHz$, the beacons and objects have a beamwidth of 30 degrees, The phase detector in the ARI receiver has a ± 1 degree resolution, and the Cartesian space is in meters. The ARI receiver was able to find the target LOB to within 6.85% of its beamwidth – a more than four-fold improvement in angular resolution.

Finally, we should note that if the $\Delta\phi \Rightarrow \Delta d$ value is greater than the range of the antenna beamwidth (e.g. outside the search volume in the center of figure 12.8) then the scatter return signal is a false positive detection and may be rejected.

12.5 Cooperating Targets

The class of cooperating targets includes civilian aviation platforms like the Traffic Collision Avoidance System (TCAS) and Air Traffic Control (ATC) transponders, GPS alternative navigation systems, and mobile secure communication tracking nodes, among others.

ARI may be deployed as a GPS-free navigation system when the transmitter beacons are deployed in known locations with known attitudes and orientations (possibly derived from Inertia Navigation Sensors (INS), GPS – should it ever become available in-situ, magnetic compass, or gravimetric sensors). The receiver operates as described previously with the exception that we relax the LOS avoidance directive and instead actively seek out LOS to the beacons. The beacons may also employ isotropic transmission antennas to reduce cost, size, power, and weight.




Figure 12.8 – *The phase difference of arrival estimates the target position to within a smaller volume and a substantially smaller angle faster than traditional AoA RADAR approaches. The squares are transmitters, the dark circle is the receiver, and the central circle is a non-cooperating target.*

Otherwise it must wait for the beacon volume scan to find it. Optimizations for this process are described [Bro07]. The application tradeoff is increased secrecy (since the antennae largely do not transmit in unnecessary directions) in exchange for reduced availability of the ARI navigation service to alternate receivers (users) since the beacons are in-service tracking the current user. Although, further field testing is necessary. ARI is conceivably capable of operating when and where GPS has been denied by enemy counter-measures or environmental/urban topology.

ARI is generally advantageous in cooperative target applications because it does not require time synchronization between the transmitters and receivers since all of the phase timing information is contained in the transmitted signals. The receiver is completely passive which helps protect the location and identity of the receiver from hostile passive RADARs. ARI allows high resolution position estimation from low frequency envelopes and affords near arbitrary control over the

choice of those envelopes. This, in turn, allows tuning, given a constant noise floor, to maximize phase resolution and the unambiguous phase region. For single-target environments ARI does not require continuous angular scanning since the phase differential can find the target inside the search volume.

12.6 Software Defined Radio Approach to ARI

Having discussed the general architecture, we now endeavour to reduce the development time and implementation cost for ARI-type systems. To that end we have designed, simulated, and implemented a Software Defined Radio (SDR) Angle-of-arrival-assisted Relative Interferometry (ARI) transceiver for position estimation recovery through the use of USRP [Ett09] hardware and gnuRadio [Blo09] software.

The use of Software Defined Radio [Cab09] in our implementation of the ARI system has added several very important advantages. Unlike analog processing, Digital Signal Processing (DSP) does not accumulate appreciable noise and may be re-tuned or reprogrammed dynamically. These attributes could be exploited in an SDR-ARI system to perform Doppler tracking, dynamically select the frequency band, and perform adaptive power thresholding. A Software Defined approach also allows for a lower implementation cost due to fewer precision components.

Under ARI encoding a unique property emerges that is well-exploited by an SDR approach, namely, that the frequency chosen for electro-magnetic propagation (carrier) may be selected independently of the frequency chosen for localization. The carrier frequency may be moved at will and the localization components continue to function unaffected. This flexibility may be exploited to continue service provision in the event of hostile jamming, interference, or the detection of the primary user's denial pilot tone.

12.6.1 Transmitter

All of the transmission components have direct analogies in the digital synthesis domain. Consequently, there is little architectural difference between a hardware or software approach.

12.6.2 Receiver

Unfortunately, implementing an SDR-ARI receiver architecture is much more complex than implementing the transmitter. Given the limitations of software radios in general, SDR's use a hardware downconversion Analog Front End (AFE) to bring the signal within the sampling system's bandwidth. In this respect, it is indifferent from a pure hardware approach. In either case, our receiver consists of detection, down-conversion (mixing then filtering), and processing. For illustration, these processes are shown in figure 12.9, in the time domain, using our Matlab model.

The sampled signal prior to any further processing consists of equation 12.21:

$$S_{rx} = S_{f_{a_1}} + S_{f_{b_1}} + S_{f_{a_2}} + S_{f_{b_2}} \quad (12.21)$$

... where S_f refers to a sinusoidal signal at frequency f .

12.6.2.1 Analog Hardware Architecture

A complete hard radio approach in analog hardware is shown in figure 12.10(top). The Beacon Specific Blocks (BSB) each contains the appropriate channel selecting filters and a pair of Variable Gain Amplifiers (VGA) that are used to equalize the amplitudes of the two frequency components (f_{x_1}, f_{x_2}). This maximizes the amplitude of the recovered f_{com} . The equalized output from each beacon at f_{com} is compared in a phase detection unit for which numerous analog approaches are




Figure 12.9 – *Signal reception from analog to digital. (a) The recieved signal at the antenna; (b) The signal after mixing; (c) The signal after low-pass filter; (d) The signal after sampling.*

known.

12.6.2.2 SDR via Time Domain Filters

After down-conversion in hardware, the signal is now within the sampling bandwidth of the software radio. The USRP uses dual 12-bit 64 Megasample per Second ADCs for capture. In the time-domain a bank of band-pass filters are required to separate $S_{f_{a_1}}$, $S_{f_{b_1}}$, $S_{f_{a_2}}$, etc. from each other as shown in figure 12.11. However, the most memory and processor efficient high-quality (Q) filter architectures (such as Butterworth and Chebyshev [LS05]) sacrifice linear phase response, which

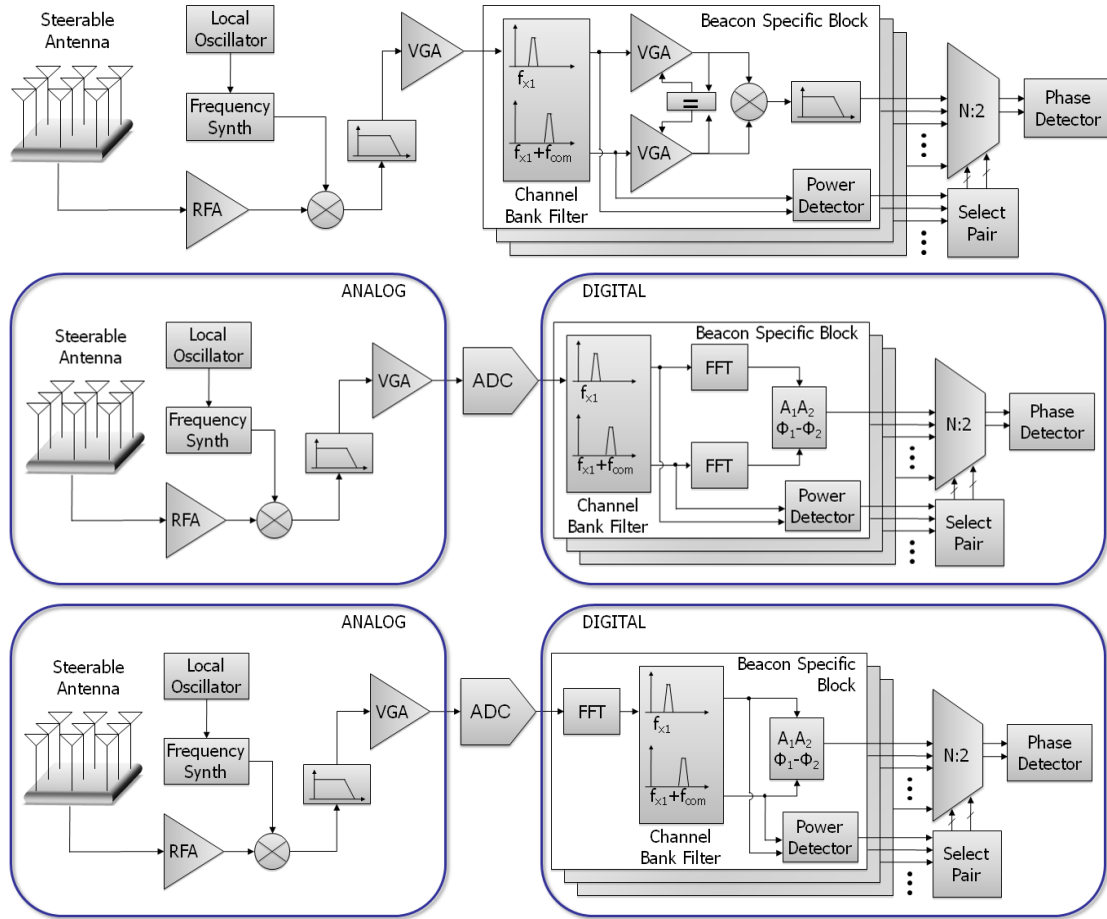


Figure 12.10 – (top) The proposed full-analog superheterodyne down-converter and signal-processing blocks for the ARI receiver [FCS08]. (middle) An SDR approach utilizing time-domain filters. (bottom) An improved SDR approach utilizing frequency-domain filters.

compromises localization performance. Alternatively, linear phase response may be achieved at sufficient Q with lengthy filter chains increasing hardware cost, delay, and scan time. Given that each component frequency must be isolated, the number of filters required is equal to the number of beacons in a simultaneous search operation multiplied by two.

Once the signal frequencies are isolated, an equivalent approach to the hardware of figure 12.10(top) is possible. The complete architecture utilizing these time-domain filters is shown in figure 12.10(middle). In which the equalizing, mixing, and phase detection are indicated by an FFT followed by a block representing

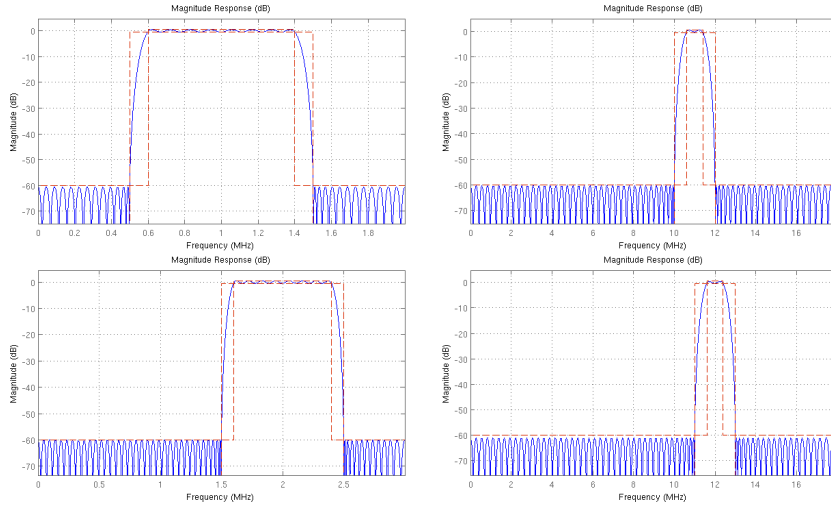


Figure 12.11 – *The Band-Pass Filter bank in the Time Domain.*

the math of equations (12.22) through (12.25). The authors recognize that time-domain approaches exist and that this is not an optimal path, but it will prove illustrative as a comparison to the frequency-domain approach which now follows.

12.6.2.3 SDR via Frequency Domain Filters

The architecture of figure 12.10(bottom) is substantially more resource and run-time efficient. Whereas the band-pass filter bank in the time-domain required long signal chains for each component frequency, frequency isolation in the frequency-domain is trivial – we simply select the frequency bin, from within a band of frequency bins around our target frequency f_n , with the largest magnitude. This process is illustrated in figure 12.12 for the four individual signals at f_{a_1} , f_{b_1} , f_{a_2} , and f_{b_2} .

12.6.3 Phase Recovery

Once we have isolated the signals at the four desired frequencies, we are able to extract the amplitude and phase information from each signal. This process is described mathematically by:




Figure 12.12 – *The Band-Pass Filter bank in the frequency domain.*

$$\sin(2\pi f_2 + \phi_2) \times \sin(2\pi f_1 + \phi_1) \quad (12.22)$$

$$\frac{1}{2} \cos(2\pi f_2 + \phi_2 - 2\pi f_1 - \phi_1) + \quad (12.23)$$

$$-\frac{1}{2} \cos(2\pi f_2 + \phi_2 + 2\pi f_1 + \phi_1)$$

$$LPF \Rightarrow \frac{1}{2} \cos(2\pi(f_2 - f_1) + \phi_2 - \phi_1) \quad (12.24)$$

$$\phi_{f_{com}} = \phi_2 - \phi_1 \quad (12.25)$$

12.6.4 SDR Feasibility

Determining the optimal value of f_{com} is not trivial. Resolution suffers as shown in figure 12.13. A large value of f_{com} requires a wide spectrum gap between each beacon's component frequencies. This, in turn, has significant implications for the hardware design, which is generally more difficult to implement as sampling rates go up and phase margins tighten. Increasing the common frequency also necessitates




Figure 12.13 – *Linear resolution as a function of $f_{com}/2\pi$, or 1 degree of phase. Values below the horizontal limit line (red) are achievable with current low-cost SDR systems.*

increasing the sampling rate to stay above the Nyquist limit. $f_{com} > 24MHz$ is presently unattainable in low-cost SDR platforms [Ett09]. This limit line is plotted in the figure.

In order to use the available spectrum most efficiently and to optimize the hardware implementation, ARI employs a frequency interleaving scheme (see figure 12.1). This allows the limited bandwidth of our SDR platform to achieve maximal values of f_{com} . The f_{x_1} frequencies are all clustered into one band and their corresponding f_{x_2} frequencies into another. The use of this dual band approach is expedient, allowing a maximum f_{com} in any finite bandwidth.

The performance of an SDR system under ideal conditions, in contrast to a hardware system under equally ideal conditions, will uniquely suffer from quantization noise introduced by its digital nature. Figure 12.14 shows the phase transfer function of the entire system end-to-end.

After a linear fit is applied the error is revealed to be quite small as indicated




Figure 12.14 – *The end-to-end phase transfer function of the SDR-ARI model including up-conversion, down-conversion, and signal processing. A linear fit is plotted in the figure*




Figure 12.15 – *The end-to-end phase transfer error with respect to a linear fit. The error comes predominantly from quantization noise during signal reconstruction.*

in figure 12.15. This error corresponds to less than 8cm of free-space length.




Figure 12.16 – *The principal ARI data signal is relatively unaffected by initial phase error in the carrier signal or in generating a quadrature variant of it. In the time domain, phase shift is not apparent in the data signal despite significant shift in the baseband signal.*

It is extremely difficult to synchronize the carriers of two independent SDR platforms [SFC08]. Fortunately, ARI encoding is substantially insensitive to phase differences in the carriers between beacons. As such, the envelope shows a near-stationary phase as the component carriers phase shift beneath it. In the time domain, as shown in figure 12.16, phase shift is not apparent in the data signal despite significant shift in the intermediate signal.

In figure 12.17, this small error is shown in greater detail. SDR-ARI is only minorly effected, $< \pm 1^\circ$, by incoherency among the beacon carriers. Consequently, time synchronization need only occur among beacons at their substantially lower data frequencies – a property which makes an SDR approach feasible.

12.7 Implementation

The SDR-ARI platform was implemented on a USRP [Ett09] using GNURadio [Blo09]. The USRP is USB-connected to a host computer system, as in figure 12.18, and used for transmitting and receiving various signals. In combination with GNURadio, we were able to receive the transmitted signal through a direct connection between the two daughterboards. As illustrated in figure 12.20, our implementation is capable of isolating received signals in the frequency domain.




Figure 12.17 – *SDR-ARI is immune to carrier incoherency between transmitters and receivers and only minorly effected, $< \pm 1^\circ$, by incoherency among the beacon carriers. Consequently, time synchronization need only occur among beacons at their substantially lower data frequencies.*




Figure 12.18 – *The Universal Software Radio Peripheral reconfigurable hardware running gnuRadio software during SDR-ARI transmit-recieve testing.*

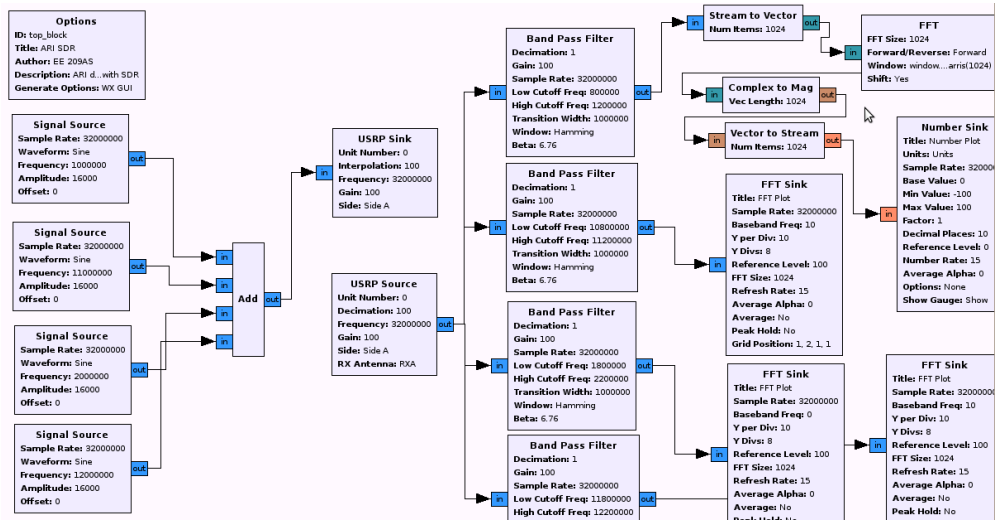


Figure 12.19 – The *gnuRadio* data flow diagram for the transmitter (left) and the receiver (right).

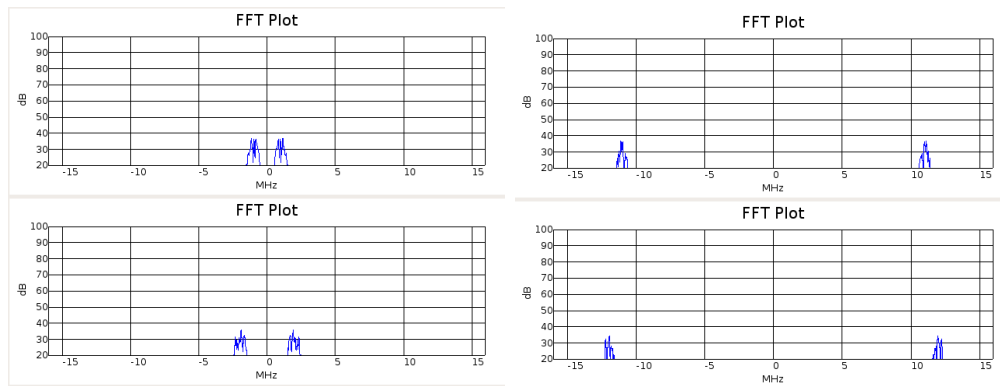


Figure 12.20 – Operation of our *gnuRadio* data flow diagram at the ARI unique (left) and common offset (right) frequencies. Here, $f_{com} = 10MHz$.

The SDR-ARI block-diagram implementation is shown in figure 12.19.

12.8 Related Work

Other indoor location estimation technologies include the use of active Radio Frequency Identification (RFID) systems, where the RFID tags are self-powered in order to identify and locate them. LANDMARC [NLL04], is a location sensing prototype system that uses RFID technology for locating objects inside buildings.

Through the use of reference tags, they are able to increase the location accuracy. However, they also face several issues including behavior variation among tags. In spotON [HWB00], another indoor 3-D location sensing technique is developed for object tagging based on RF Signal Strength. Several base stations are used to measure the signal strength that will be mapped to an approximate distance. A central sever is used to triangulate the exact object location with these measurements using an aggression algorithm.

Following the work in [MVD05], the Vanderbilt team developed inTrack, a cooperative tracking system based on radio interferometry [KBS07]. This system is developed to track a mobile sensor with high accuracy in large areas with moderate multipath tolerance. They use XSM motes by Crossbow, running TinyOS operating system. The receivers measure the frequency and phase of the interference signal, and using the interferometric ranging algorithm, the server calculates the location of the tracked node. However, in this system, again they require time synchronization between the transmitting nodes, as opposed to ARI-SDR. They are also only able to track a single target.

In [LVS08], another team from Vanderbilt University argues that although radio interferometry has proven successful outdoors [MVD05], it is still not applicable indoors due to its sensitivity to multipath without further theoretical and experimental research. They present their preliminary results, which are implemented on a Software Defined Radio, specifically, gnuRadio [Blo09] and USRP [Ett09] platform. The use of a SDR allows them to use a wide range of frequencies, powerful signal processing capabilities, and gives them flexibility as far as radio protocols. As in [MVD05] the team was facing time synchronization challenges between receivers, they decided not to put this restriction on their design. Instead, they have one of the transmitters embed an indicator in its signal, marking a common point in time for both receivers. Once this is done, the receivers have a common reference from which they can measure the phase of the signal.

12.9 Conclusion

This chapter presented early findings from an Angle-of-arrival-assisted Radio Interferometry approach to target localization, described its operation, provided an analysis of its performance limitations, and suggested future research directions. ARI's performance is expected to improve as related work [LS05, FG94, Ger87] in the field of antenna array pattern synthesis continues to advance and reduce the search volume given a fixed number of transmitters and receivers.

ARI could prove advantageous in a diverse set of military application scenarios. To preserve stealth, reduce power and form-factor, and increase detection range ARI beacons could be placed near vital intersections, along thoroughfares, or deployed and redeployed by recon units. The ARI receiver is then mounted to the lead convoy vehicle which ascertains mine presence from the back-scatter of the beacons. As in GPS, the vehicle mounted equipment is completely passive allowing stealth operation and because the ARI beacons are forward-deployed, detection range and accuracy is improved.

In this chapter we have also presented the design and implementation of an Angle-of-arrival-assisted Relative Interferometry (ARI) approach to target localization through the use of Software Defined Radio (SDR). We have described the architectural implementation of this system, simulated its operation, instantiated the transceiver signal processing chain in real hardware, and provided an analysis of its performance advantages and disadvantages.

CHAPTER 13

Pulse-Wave Angle-of-Arrival-Assisted Relative Interferometry

13.1 Overview

In the prior chapter we proposed the use of Angle-of-Arrival (AoA) information in conjunction with local interferometry to improve target location estimation in the spatial domain. In this chapter we propose a Pulse-Wave extension to our prior technique that dramatically improves the scan-time required and reduces the vantage points necessary to obtain high-fidelity target position estimation. Accordingly, we call this enhanced process Pulse-Wave Angle-of-arrival-assisted Radio Interferometry (PW-ARI, pronounced “P-W-are-eee”). As we will show, it has a number of desirable attributes. Foremost among these are its ability to rapidly image extremely large volumes with high accuracy and dense target clusters, support an infinite number of friendly aircraft while providing stealth operational support, and reducing the synchronization and hardware requirements when operating as the sole RADAR modality.

A Pulse-Wave Angle-of-arrival assisted Radio Interferometry (PW-ARI) system consists of two closely-located transmitting beacons (they might be almost co-located and appear to an operator as a single piece of equipment – as shown in figure 13.3) and any number of passive receive-only aircraft or ground-stations (assets) each of which is equipped with a sector (or steerable) directional antenna. The beacons are tightly time-synchronized. Each beacon transmits a unique pair of

frequencies, creating a beat envelope in the transmitted signal which is not unique to any beacon (the ARI payload). The receiver can then identify scattering targets in the intermediate environment through the fusion of the Time-of-Arrival (ToA) of the pulses and the Phase-Difference-of-Arrival (PDoA) of their ARI payload. Target positions are further disambiguated by the known look-directions of the antennas.

ARI encoding [FCS08] offers the unique property that the electromagnetic wavelength and the wavelength used for localization may be chosen independently. Airborne target detection might best be performed using frequencies in a particular portion of the S-Band, thereby maximizing the effective Radar Cross Section (RCS) of the expected targets. Target localization, however, is best performed at substantially lower frequencies to maximize phase disambiguity and, therefore, target bearing. ARI encoding allows both of these conditions to be met.

This independent frequency selection may be exploited to achieve the same volumetric coverage as other RADAR modalities with substantially fewer transmitters. Or this selection can allow a PW-ARI system to image the airspace in a single transmitted pulse, since PW-ARI stations do not need to mechanically scan, nor transmit one pulse per bearing (as in electronic steering). For a given number of ground stations this coverage expansion manifests as greater system redundancy, which in turn makes the network harder to attack. In an eight-node network, the system would need to experience $>85\%$ loss to enemy action before it is no longer combat effective. Conversely, repairing just one of those stations restores the network to operation. As PW-ARI transmitters cover a large volume, they may be distributed spatially to further harden the network.

The actual target detections and determinations are made by the assets which operate continuously in a passive mode. The assets must maintain time synchronization with the beacons in order to disinter ToA data, but this does not require any transmission on their part. Target detections from various assets may be fused

to improve resolution, accuracy, and coverage.

RADAR is the primary mode of aircraft detection in the battlespace as well as the principal safety and control mechanism for commercial aviation in the private sector. PW-ARI offers an inexpensive supplement or alternative to the commercial Traffic Collision Avoidance System (TCAS) because it requires few transmitters for a wide coverage area and, as the high-power transmitters are on the ground, does not stress the electrical system of the passenger aircraft (which results in greater fuel costs). PW-ARI can detect aircraft beyond those equipped with a compliant TCAS transponder (used in general aviation) and those for which the transponder is malfunctioning or otherwise disabled.

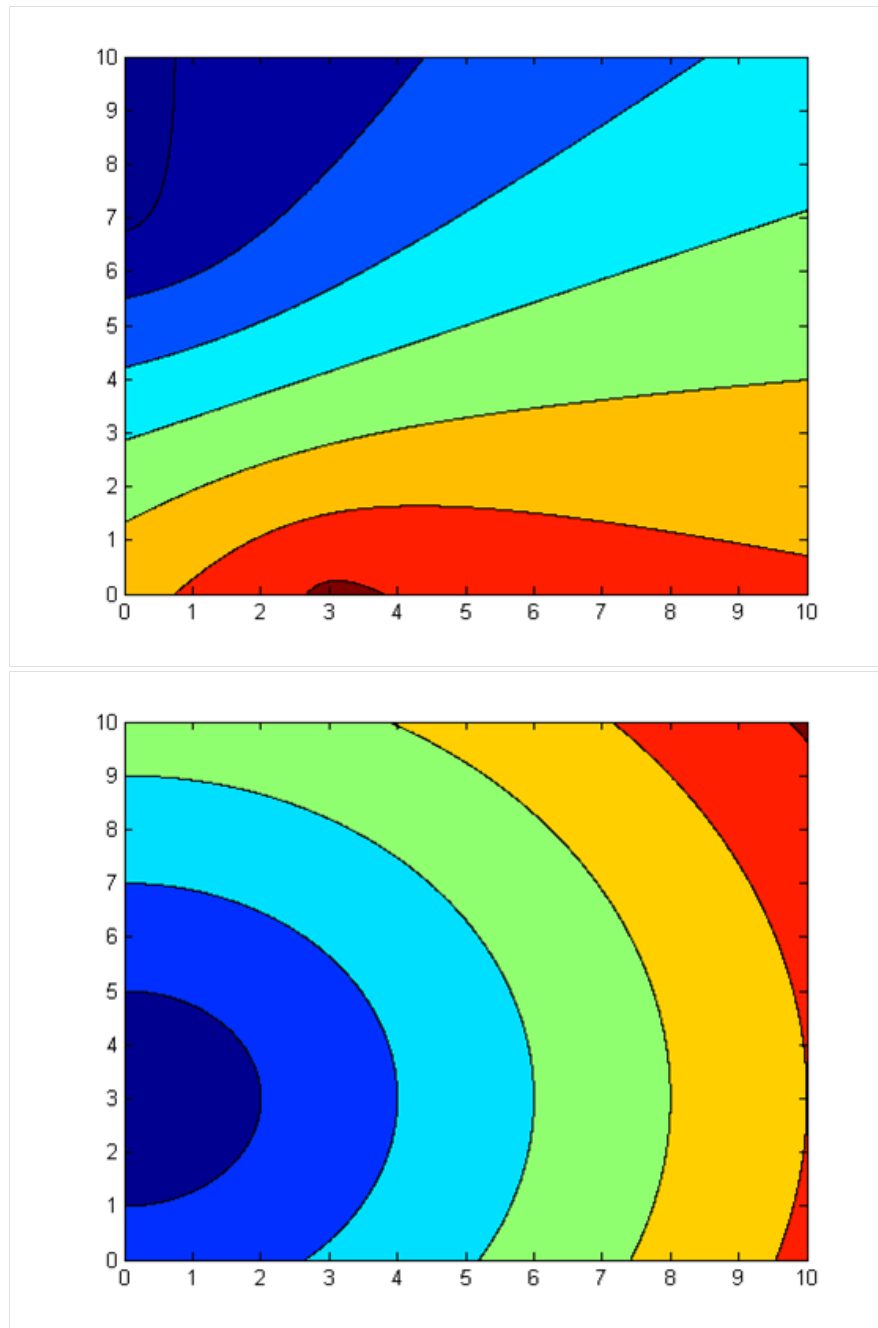


Figure 13.1 – *PW-ARI phase contours (top) and time-of-flight contours (bottom). PW-ARI recovers the phase-difference of arrival (top) and the time-of-arrival (bottom) simultaneously making fusion possible.*




Figure 13.2 – *The PW-ARI phase contours and time-of-flight contours overlaid as observed at the receiver. Note the grid-like coordinate system created from the fusion of these two contour sets. PW-ARI recovers the phase-difference of arrival and the time-of-arrival simultaneously making fusion possible.*

13.2 Pulse-Wave ARI

The fundamental contribution of this work is the addition of time-of-flight information to the ARI regime. PW-ARI is the fusion of data from three domains: time (time-of-flight), phase (relative phase-of-arrival), and angle (direction-of-arrival) – while conventional ARI includes only the latter two. Further, the relatively narrow beamwidth constraints of ARI are substantially relaxed under PW-ARI as the requisite close spacing of the transmitting beacons allows broader beamwidths to be accommodated. PW-ARI is also substantially faster at imaging large volumes replete with multiple targets as it does not need to rely on differential target illumination (sweeping) to separate targets at multiple distances (this may now be done in the time domain).

The PW-ARI performance improvements stem from the observation that the isolines created in the phase domain are independent of the isolines in the time domain. By combining the two, a coordinate grid is created which is easily mapped onto Cartesian space.

In the top frame of figure 13.1, the isophasic contours are drawn in space for a pair of PW-ARI beacons centered at $(0, 3)$. The points of each contour represent the set of locations from which the difference in the path length from the point to each of the two transmitters is the same constant. By definition, the isolines take on a hyperbolic fan-like appearance concentrating near the beacons and separating in an angular fashion.

In the middle frame of figure 13.1, the contours represent the locations from which a particular time-of-flight will be observed should a signal leaving the beacons strike a target at that location and return to an asset receiver located, in this case, at $(0, 3)$. The isolines are roughly circular in shape (with the irregularity being caused by the beacons' non-collocation with the asset) since the time-of-flight corresponds to distance travelled and the distance from a circle's perimeter to its center is constant.

Combining these two observations results in figure 13.2. In this scenario the beacons are located at $(0, 6 \pm 1)$ and the asset is at $(0, 3)$. The isolines based on phase and those based on time are overlaid and the resulting grid system becomes apparent. PW-ARI recovers the phase-difference of arrival and the time-of-arrival simultaneously, which makes accurate fusion possible. The following sections explore how this unique capability is achieved.




Figure 13.3 – A Pulse-Wave Angle-of-arrival-assisted Radio Interferometry (PW-ARI) based system consists of pairs of transmitting beacons, grouped into stations, and passive receive-only assets which make use of the transmitter timing and RADAR echo information to locate targets. The beacons transmit unique pairs of frequencies creating a beat envelope which is not unique to any beacon allowing phase comparison.

13.2.1 Transmission

The transmitted signal from each beacon consists of two frequencies – one that is unique to the beacon, f_{x_1} , and one that is shared across all beacons, f_{com} . A Hartley modulator architecture [Dar70] is used to produce the following signal over the air:

$$\begin{aligned}
 & \sin(2\pi f_{x_c} t) \left(\sin(2\pi f_{a_1} t) + \sin(2\pi (f_{a_1} + f_{com}) t) \right) + \\
 & \cos(2\pi f_{x_c} t) \left(\cos(2\pi f_{a_1} t) + \cos(2\pi (f_{a_1} + f_{com}) t) \right) + \\
 & \sin(2\pi f_{x_c} t) \left(\sin(2\pi f_{b_1} t) + \sin(2\pi (f_{b_1} + f_{com}) t) \right) + \\
 & \cos(2\pi f_{x_c} t) \left(\cos(2\pi f_{b_1} t) + \cos(2\pi (f_{b_1} + f_{com}) t) \right)
 \end{aligned} \tag{13.1}$$

$$\dots \text{for } t = [0.. \frac{1}{f_{com}}]$$

Algebraically we see that (13.1) resolves to one period of a signal with frequency f_{com} that consists of signal at just four frequencies: $f_c + f_{a_1}$, $f_c + f_{b_1}$, $f_c + f_{a_1} + f_{com}$, and $f_c + f_{b_1} + f_{com}$ [Dar70]. This is a Single Side Band-Suppressed Carrier (SSB-SC) transmission. Since this is the absolute minimum amount of spectrum required to convey all of the information required for PW-ARI, we are as spectrally efficient as possible.

13.2.2 Reception

PW-ARI signals are received as an ARI-encoded pulse train on two channels – each corresponding to a particular beacon’s unique frequency pair. The ToA on either channel implies the target range while the PDoA between the payloads of each pulse in the two channels, at f_{com} , implies the bearing. Unfortunately, two transmitters on the same frequency will interfere and obscure the identity and phase of each. ARI supports multiple simultaneous transmissions without

interference because the common frequency information is actually present on multiple unique frequencies ($f_{x_2} = f_{x_1} + f_{com}$, for $x = a, b, c...$). Accordingly, we designate f_{com} as the baseband signal and task the PW-ARI receiver with recovering it from each beacon.

Although the transmitted signal is up-converted in quadrature pursuant to [Dar70], the received signal does not require down-conversion in quadrature – a single conversion superheterodyne receiver is sufficient. The spectrum of the resulting intermediate signal contains the identities of any beacons transmitting towards any in-range assets. Beacons may be identified by the presence of a signal at f_{a_1} , f_{b_1} , etc and at each of these frequencies offset by f_{com} (a Doppler correction is required).

Multiple beacons are detected through the use of multiple beacon-specific blocks. Each contains the appropriate channel selecting filters and a pair of Variable Gain Amplifiers (VGA) that are used to equalize the amplitudes of the two frequency components (f_{x_1} , f_{x_2}) and to perform time-range compensation. This maximizes the amplitude of the recovered f_{com} . If multiple beacons exist, the equalized output from each beacon at f_{com} is compared in a phase detection unit.

The phase detection unit is constructed from a network of delay lines tapped and connected to a correlator (a summing circuit followed by a peak-detector). The tap point returning the strongest signal represents the most-likely phase of arrival value. As reported in [BOA91], such an architecture can resolve a time difference of 8 picoseconds for signals at a common frequency. Applying the error model in [FCS08] and using commercially available component performance values, this resolution degrades to ≈ 14 picoseconds – or about 0.1° at 20MHz.

In the simplest two-beacon case, the phase detector would return the signal $\Delta\phi = \phi_a - \phi_b$, the phase difference between the signals from beacons a and b at

their recovered common frequency.

$$\Delta\phi = \frac{(d \bmod \lambda)f_{com}}{c} \quad (13.2)$$

This phase difference corresponds to the relative physical distance between the paths from each beacon, through the environment (which may be reflected or otherwise indirect), to the asset. This relationship is presented in (13.2). As the beacons are closely spaced, the isophasic lines are more useful in determining target bearing with respect to the beacon pair than target range.

13.2.3 EW Considerations

PW-ARI offers a number of benefits relevant to Electronic Warfare (EW) concerns. Passive assets are intrinsically stealthy since synchronization and target detection do not require active transmission. The powerful long-range transmitters may therefore be well defended by hidden forward-positioned assets. PW-ARI transmissions may be made extremely narrow-band whereas enemy Radar Warning Receivers (RWR) must incorporate a wide-band receiver with its associated higher noise floor [Che07]. Time synchronization allows a frequency hopping scheme to further hide the transmitter's spectral location.





Figure 13.4 – *PW-ARI coverage area vs. inter-beacon distance (top). Average accuracy for a target 50km in range using 20MHz (bottom).*

13.3 Performance Expectations

13.3.1 Timing Requirements

Current time synchronization technology allows us to synchronize two systems to an accuracy well below 50ns. The Global Positioning System (GPS) is one technology which allows clocks to synchronize to Universal Time, Coordinated (UTC) to within 10ns [LNN01]. However, time synchronization is not the only thing necessary to make PW-ARI work. We also need to be able to keep the accuracy of time below 50ns over the long term, because the receivers and transmitters could be jammed or, as was recently shown in [HLP08], spoofed. Therefore, the receivers need very accurate clocks in order to keep accurate timing for the extent of the mission. The stability of the clock has to be

$$\delta f \cdot T \leq 50ns,$$

where δf is the clock stability measured in parts per million (ppm), and T is the length of time for which the clock has to perform without a resynchronization event. For example, for a mission length of 5 hours, the clock stability has to be less than 10^{-12} . Available Rubidium oscillators are near this level of frequency stability and the advent of chip scale atomic clocks [BJB06] supports our belief that, in future systems, this level of clock stability will be readily available.

13.3.2 Common Design Parameters

PW-ARI shares a number of design parameters with pulsed RADAR systems in general. As previously discussed, ARI encoding allows the independent choice of carrier frequency, which should be chosen to accommodate EW concerns (target reflection, frequency hopping hardware limitations, etc). With any pulsed system, the Pulse Repetition Frequency (PRF) is of concern although it is well understood

from PW RADAR, medical ultrasound, and other domains.

13.3.3 The Resolution-Coverage Trade-off

Of principal and unique interest in PW-ARI system design is the ARI frequency and beacon separation distance. For constant frequency, increasing the inter-beacon distance improves localization accuracy at the expense of reduced coverage volume. For constant separation, increasing the frequency reduces coverage and improves bearing resolution. Thus, the effects are largely reciprocal, suggesting an optimal configuration for each $\langle volume, accuracy \rangle$ tuple.

Given an ARI frequency, f_{com} , of $20MHz$, figure 13.4 (top) shows the effect of inter-beacon distance on the PW-ARI coverage area over a $100km \times 100km$ surface. The area-to-separation function is non-linear. In the lower plot of figure 13.4, the target location accuracy at $50km$ is shown versus the inter-beacon separation distance. Increasing the distance narrows the isophasic lines and therefore improves target bearing resolution, in turn, improving accuracy. In this example, the target bearing was directly along the boresight where ARI angular resolution is at its worst-case performance since the signal phase difference is near-zero.

To illustrate this further, consider figure 13.5. Here a central target (yellow circle) is illuminated by a ground station $50km$ due West. The receiving asset, a friendly aircraft flying patrol, is located to the South-West of the figure. The horizontal band in the figure represents the region of likely target detection based on the PDoA from the PW-ARI payload in the return signal. The diagonal band lies orthogonal to the asset's flight-path and represents the region of likelihood based on the ToA data derived from the worst case $50nS$ time synchronization between asset and beacon. The intersecting region is the target detection estimated position. Were other assets available to corroborate the detection, the data could be fused. Observed from different asset-to-target angles, the target estimation




Figure 13.5 – A PW-ARI RADAR system imaging a target at 20MHz, from 50km.

volume would reduce substantially.

Finally, PW-ARI requires hard-sector transmission – the simultaneous illumination of an entire RADAR sector, as opposed to sweeping through a sector with a narrow beam. Antennas of this type are inherently more power efficient than electronic steering and faster than mechanical scanning. However, the antenna pattern of the beacon must be limited to the sector size or returns from outside the sector will alias in the phase domain.

13.4 Conclusion

In this chapter, we have introduced a pulsed variant to the ARI architecture and discussed its possible limitations and implementations. The speed, simplicity, redundancy, and stealth offered by PW-ARI RADAR systems is extremely attrac-

tive and the technology necessary to achieve such systems is within the range of possibility at the time of authorship.

Furthermore, the lack of sophisticated antenna systems and mechanical components suggests that a PW-ARI RADAR system might be sufficiently cost-effective and reliable to suffice for civilian aviation, border patrol, and other high-quantity applications.

In the case of a biomimetic electrostatic imager, the reliance on signal characteristics instead of tightening the beam width (which, beyond some minimal degree, is not really possible) is highly advantageous.

Part V

Biomimetic Electrostatic Imaging

CHAPTER 14

3-Electrode Biomimetic Electrostatic Imager

14.1 Introduction

The sum total of the dissertation has brought us to this point. This and the three chapters that follow form the final portion. They work through the historical evolution of the system-level design and implementation of the first multichannel electroreceptive sensor array capable of creating *Biomimetic Electrostatic Images* (BEI). As such, they depend heavily on the work of prior sections especially chapter 2, which provides the fundamental boundary conditions, chapter 6, which provides the reference design for a single channel, and chapter 7, which describes a complete analytical model for this class of systems.

Each of the three described systems represents a substantial improvement over its predecessors as summarized in table 14.1.

| # Electrodes | Year of Publication | Electrical Resolution | Physical Resolution | Image Formation Time |
|--------------|---------------------|-----------------------|---------------------|----------------------|
| 3 | 2009 | 20 mV_{AC} | 12.7 cm | 2.5 hours |
| 16 | 2010 | 5 mV_{AC} | 5.1 cm | 0.5 hours |
| 17 | 2011 | 0.5 mV_{AC} | 3.0 cm | 20 ms |
| 68 | 2011 | 0.01 mV_{AC} | 1.0 cm | 20 ms |

Table 14.1 – *Improvement in the performance of our biomimetic electrostatic imaging systems over time.*

14.2 Proximity Sensing

The first step in building an imaging system was the design of a mere proximity detector. The sensor began with a purpose-built physical modeling engine we implemented in Matlab to visualize the electric fields generated and the disturbance that occurs with environmental objects. By abstracting the environment only to relevant net charges (with respect to the background ion concentration), direct analysis becomes computationally practical via:

$$Voltage = \sum_{n=0}^{N_c} \frac{1}{4\pi\epsilon_0} \frac{q_n}{r_n} \quad (14.1)$$

Where N_c is the number of net charges in the environment to consider and q_n expresses the net magnitude of each charge. r_n is the distance between the spatial location under consideration and each specific net charge. ϵ_0 is the permittivity of a vacuum. The equation may be scaled by the relative permittivity of other media for consideration there, but as this has no spatial consequences (only influences magnitude), it is ignored here.

In two dimensional space the model output given a dipole transmitter and a linear array of receivers appears in figure 14.1. The isovoltaic field lines are indicated. If the voltage field is transected along the dashed line of the figure a voltage versus position plot emerges. If you apply an electric field to a conductor, the mobile charges within it will migrate to align with the field creating a dipole with negative and positive ends and, thus, redistribute the isovoltaic field lines. When this induced dipole is introduced to the scenario, the current is redistributed by the difference in force/voltage. The revised scenario appears in figure 14.2. Note the effect of the object in distorting the field.

Recall the experimental configuration of figure 8.1. Recordings were taken across the tank for both with and without a metal pipe target present. In figure

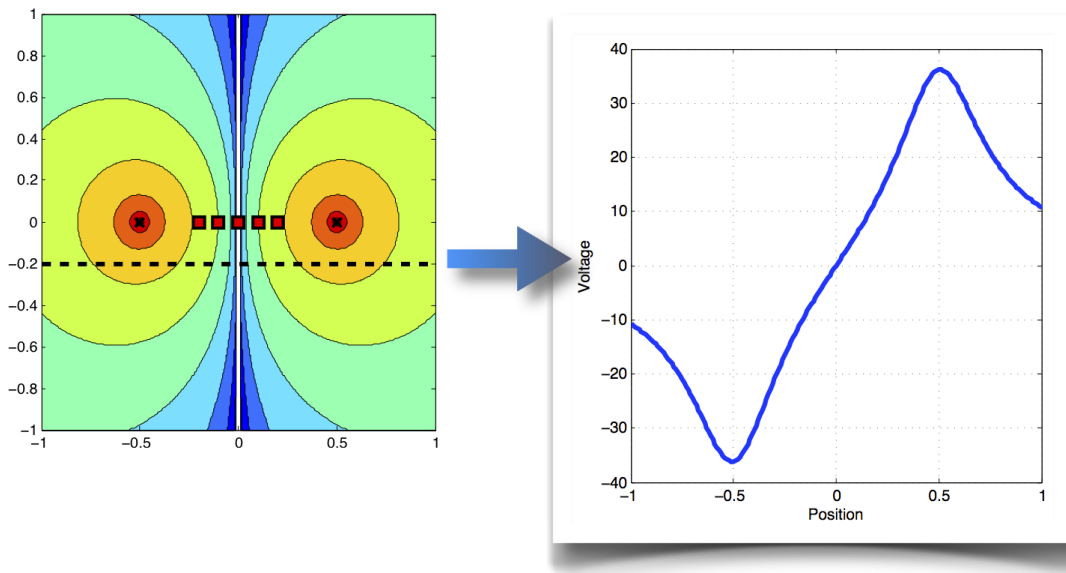


Figure 14.1 – (Left) *Output of our purpose-built analysis tool given a dipole transmitter and a linear array of receivers . The isovoltic field lines are indicated. (Right) Voltage versus position along the transect indicated by the dashed line in the left panel.*

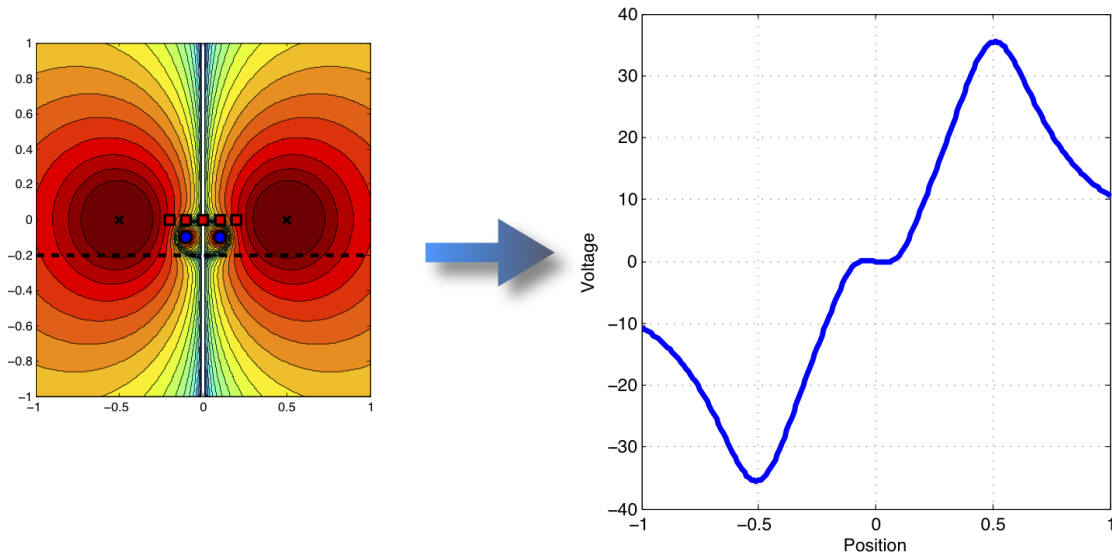


Figure 14.2 – (Left) *Output of our purpose-built analysis tool given a dipole transmitter and a linear array of receivers with an induced dipole (conductive object) present. The isovoltic field lines are indicated. (Right) Voltage versus position along the transect indicated by the dashed line in the left panel. Note the distortion.*




Figure 14.3 – A comparison of tank-acquired and simulated electric field measurements. The darker traces are the background scans. The lighter traces are with a dipole present.

14.3, we plot comparison traces from simulations and tank trials. If we transect a 2D space along a line parallel and close to the transmit axis, the resulting graph is an x-z plot indicating the Voltage per 1D position. The left pane of figure 14.3 is a dataset recorded in the actual tank, where the right pane of the figure is the output of our custom simulator which models equation (14.1). The blue (darker) traces are the tank under background conditions (no targets). The green (lighter) traces were recorded/simulated with the presence of a galvanized steel pipe (located at the position indicated by the black dash in the left pane) under otherwise identical conditions. Note the strong agreement.

14.2.1 Resolving Self-Interference

The design of a proximity detector requires the isolation of environmental dielectric disturbances (targets) from those self-induced. Electrode P_G (figure 8.1) is a source of self-interference as it is the probe tip mounted to a moving gantry. The gantry system is a poor substitute for the 15,000 simultaneous sampling channels




Figure 14.4 – *The design of a proximity detector requires the isolation of environmental dielectric disturbances (targets) from those self-induced. Subtracting the reference electrode reveals the disturbance field resulting from the target’s presence.*

of the weak electric fish [Kra94] because it physically disturbs the space and must be moved over time resulting in unintended correlation between the spatial and temporal domains.

To mitigate this effect we use the reference electrode, P_S , to measure the gantry-to-background self-disturbance and cancel it out through an additive process (the top surface of figure 8.5).

14.2.2 Results

In order to find small targets and localize them, we must subtract the excitation field by comparing object (target) scans with a background scan – this requires calibration. This is not practical for ocean deployed sensors, which is why subsequent designs (section 14.6, chapter 15, and chapter 17) transition to model based detection relying progressively less on pre-calibrated spaces until chapter 17 where

only a sensor-model (not target-model) pre-calibration is employed.




Figure 14.5 – Preliminary findings from our object proximity/imaging sensor.

Five trials were performed with just the background and five with the pipe positioned variously across the tank (the positions of the pipe are indicated by the dots in figure 14.5). Detection appears highly reliable with the pipe successfully detected in all five positive trials and successfully not detected in all five negative trials. The detection range of the study was ≈ 5 cm. The position accuracy in the middle of the transmit axis was ± 5 cm after calibration.

14.3 From Proximity Sensing to Imaging

Proximity sensing represents the simplest case of the more general problem of imaging in a non-structured (non-pre-calibrated) space. The next step was the construction of an three (receive) electrode array on our journey towards recreating the, as many as, 15,000 receptors in an adult male fish [Eva97]. To mimic this structure, transmit and receive probes were mounted to a Delrin (non-conductive) plate 18" (45.75 cm) wide and 7" (17.75 cm) tall for focusing. The Delrin plate emulates the body of the fish by preventing current flow away from the target volume [MSM01].




Figure 14.6 – A diagrammatic, target-facing, view of the tank electrode configuration. All dimensions are in inches.

The configuration consisted of five electrodes in total mounted along the horizontal centerline of the plate. Excitation was introduced between the excitation electrode, P_T , and the neutral electrode, P_0 . These electrodes were each placed 7" (17.75cm) from the vertical centerline as indicated in figure 14.6. The three sensing electrodes, P_L , P_C , P_R , were arranged left-to-right when viewed from the front of the tank (which is the reverse of the view from the target's perspective). The separation between the sensing electrodes was 4" (10.15cm) as indicated in the figure.

The current axis runs between electrodes P_T and P_0 and all of the receive electrodes share P_0 as the unipolar reference. The electrodes were shielded to the P_0 potential using twisted pair at 6 twists per 5cm (shielding terminated at the base of the feed line, above the waterline, as proper water-proofing of the shield line could not be guaranteed).

Each of the sensor electrodes was terminated into an Agilent 34410A 6.5 digit digital multimeter (three in total). 7.4mW was sourced into the tank's impedance of 307.7Ω from an Agilent 33220A arbitrary waveform generator configured for $5V_{pp}$ output at 200Hz through a 50.0Ω source resistance. These parameters were designed to approximate an electrically-active fish [Nel].

14.4 The Testing Environment

Testing was completed in a modified 55-gallon ($\approx 208\text{liter}$) acrylic fish aquarium ($91\text{cm} \times 51\text{cm} \times 38\text{cm}$). The filled water was an ocean simulant (28ppt; relative density 1.0205; at room temperature) which expresses the extreme low-end (worst case) of the Pacific Ocean's normal range [Mar07].

The tank was installed beneath a purpose-built motion-controlled gantry system (figure 14.7) consisting of a two axis carriage with a 31" (78.75cm) x-axis range of motion and a 33" (83.80cm) y-axis range of motion. All of the electrodes were submerged to an equal depth in the tank of approximately 8.9cm, while the tank itself was filled near capacity.

14.5 Waveforms and Scanning

While the sine-wave excitation was driven between electrodes P_T and P_0 , the resulting field was recorded using P_L , P_C , and P_R referenced to P_0 . This field was then disturbed by the introduction of a 1.50" (3.81cm) diameter galvanized steel




Figure 14.7 – *A model of the motion control platform with the relative positions of the tank, electrodes, and conductive pipe target indicated.*

pipe 8" (20.32cm) in length with its long axis perpendicular to, and submerged to just below, the surface of the tank. This pipe (conductive) target was then mechanically scanned around the tank in a raster-style pattern.

The entire test rig was controlled from a single PC using a Python script to interact with our custom gantry-controller and the socket-based SCPI-over-Ethernet control framework of the suite of Agilent bench-top instruments.

In each of the spatial plots to follow, the coordinate system is arranged with the Y-axis spanning the longer dimension of the tank with its zero at the midpoint of the sensor array (e.g. the location of P_C). The X-axis spans the shorter dimension of the tank and uses negative values to indicate the direction as into the tank's target volume area. The zero of the X-axis is at the plane of the sensor array. The Z-axis is the amplitude in DC-equivalent (e.g. root-mean-square) Volts AC and




Figure 14.8 – *The experimental salt-water tank and supporting test, excitation, and measurement equipment.*

not the submergence depth in the tank. All of the datasets were recorded at a depth of approximately 8.9cm below the surface.

14.6 Early Biomimetic Electrostatic Imaging

The limited volume available in the tank prevents the assumption of a clutter-free far-field testing environment. The aim of this work is to demonstrate feasibility and argue for further study. Accordingly, the interactions of the generated and disturbance fields with the tank periphery were calibrated out by performing a high-resolution spatial scan by moving the target pipe position and recording the resulting field at the electrodes rather than by developing and applying a parametric model.




Figure 14.9 – Raw sensor recordings from P_L (top, red), P_C (middle, blue), P_R (bottom, green) with respect to the position of a conductive target.

The results from a typical trial appear in figure 14.9. The recording produces three surfaces the upper-most recorded from P_L , the middle from P_C , and the lowest from P_R . The absolute amplitude of the recorded values is dependent on salt concentration, probe material, and excitation current [FTS10]. The variance may be minimized, and a parametric model simplified, by defining two differential channels (P_{dL} , P_{dR}) from our three electrodes as follows:

$$P_{dL} = P_L - P_C \quad (14.2)$$

$$P_{dR} = P_C - P_R \quad (14.3)$$

This appropriately mimics the biology, which is responsive to differential signals along the body [Nel], and portends an explanation as to why nature evolved along this course.

The presence of the conductive target in the water redistributes the current




Figure 14.10 – A spatial view of the two differential channels P_{dL} (upper) and P_{dR} (lower). The use of differential measurement improves the sensitivity of measurements as it reduces the dynamic range requirements allowing higher gain in the receiver.

which, in turn, redistributes the isovoltaic lines of the electric field. This changes the relationship between P_L , P_C , and P_R in a manner that is clear and predictable when evaluated differentially. In figure 14.10, a spatial view of the two differential channels, P_{dL} (upper) and P_{dR} (lower), is shown. The Z-axis is the amplitude of the differential signal when the conductive pipe target is located at the associated X-axis and Y-axis coordinates. The use of differential measurement improves the sensitivity of measurements as it reduces the dynamic range requirements allowing higher gain in the receiver.

To extract useful information about the surrounding environment an image must be reconstructed from the real-time probe measurements and the model known a priori. For the following analysis the model data was used directly so the same mechanical scan in the tank generated the model and the test points for the image reconstructions. This represents best case behavior and, therefore, best




Figure 14.11 – *A Biomimetic Electrostatic Image expresses the likelihood that the observed disturbances to the voltages measured originated from an object located at the pixel affiliated spatial coordinate. The images shown here are two-dimensional expressing a top-down look at a plane parallel to the surface of the tank and submerged to a depth of about 9cm.*

illustrates the remaining challenges in improving the process that we refer to as *Biomimetic Electrostatic Imaging*.

The image expresses the likelihood that the observed disturbances to the voltages measured originated from an object located at the pixel affiliated spatial coordinate. The images shown here are two-dimensional expressing a top-down look at a plane parallel to the surface submerged to about 9cm. To our knowledge, figure 14.11 represents the first published biomimetic electric field disturbance image.

The image uses the two differential channels to enhance its resolution and accuracy. The black regions are unlikely as they are neither correctly matched in either differential channel. The grey regions are matched in only one of the differential channels. In practice, these grey regions would be drawn black, but are shown here to illustrate the dramatic improvement emanating from the addition of a second channel and to suggest, by extension, how even higher channel counts contribute to image refinement. The white regions represent the location of disturbing objects in the field as these regions match in both differential channels.

14.7 Results




Figure 14.12 – *Estimated target location area (the image’s white pixels) as a percentage of the entire scan area. The X-Y coordinates indicate the location of the pipe at the time the image was constructed and analyzed.*

In figures 14.11 and 14.13, the actual location and diameter of the pipe target is indicated by a red circle. The imaging approach correctly identifies the pipe target’s location, but it also identifies several other false-positive regions. This degrades the position accuracy in aggregate, but, in certain applications, may still provide accurate position estimation since each of the white regions is distinct and one of them is correct. In our investigation of 55 images in four trials (220 images

in total) no false-negatives were recorded – a reality guaranteed by the modeling process.

The quality of the images varies dramatically in the 2-channel (3 electrode) imaging system. Poor noise isolation performance in the front-end results in a high degree of variability and limited position resolution. The various panels of figure 14.13 demonstrate the limited ability of the system to localize the metal pipe.




Figure 14.13 – *Various biomimetic electrostatic images from the 2-channel (3 electrode) imaging system. Poor noise isolation performance in the front-end results in a high degree of variability and limited position resolution.*

Images were constructed for the pipe target located at 55 unique positions.




Figure 14.14 – Mean position estimation error of the white pixels with respect to the center of the true location of the disturbing pipe target. The X-Y coordinates indicate the location of the pipe at the time the image was constructed and analyzed.

The results appear in figure 14.12. For each of the spatial coordinates indicated in the figure by the X- and Y-axes, the Z-axis represents the percentage of the image area in which the target lies as estimated by the BE Imager. In aggregate over the 55 images, that area is 91.6% smaller than the imaged space. Given the pipe’s diameter of 1.5”, the pipe quite favorably occupies at least one-third (34.9%) of that mean targeted region. Many (20%) of the images fared substantially better with the area reduction exceeding 99% corresponding to the target occupying almost three-times ($2.9\times$) the area of the target estimate indicating the high-resolution achieved with BEI.

However, the image bright spots (target estimations) are not contiguous and, therefore, the distribution of the target estimates may impact accuracy in certain applications. Generic imaging requires that these extraneous bright spots be mitigated. To quantify the current scale of the problem, the 55 images from figure 14.12 were reanalyzed such that, in figure 14.14, the Z-axis represents the mean distance from each of the white pixels in that image to the correct ground-truth

location at the center of the pipe target. This is actually pessimistic since the pipe target has a non-zero radius, but the worst-case is sufficiently illustrative. Intuitively, accuracy generally improved with decreasing distance between the pipe target and the sensor array. The maximum range of the sensor array was not determined in the confines of the tank as the pipe target was reasonably located to the maximum extend of the tested range of 8" (20.32cm).




Figure 14.15 – *The prior figure viewed end-on from the Y-axis such that it now expresses the relationship between position across the long-dimension of the tank (parallel to the sensor array) and the mean position error. The parabolic nature of this lateral error function is readily apparent.*

In figure 14.15, the prior figure is rotated so that it is viewed end-on from the Y-axis such that it now expresses the relationship between position across the long-dimension of the tank (parallel to the sensor array) and the mean position error. The parabolic nature of this lateral error function is readily apparent. The array is most sensitive, and therefore most accurate, at its center – a product of the waning influence of the further channel as the target moves farther to the extremes of the array. The vertical bands express the range of accuracies observed in measurements at different distances (X-axis values) from the sensor array.

14.8 Conclusion

This chapter reviewed our earliest work with biomimetic electrostatic imaging and presented the first images from our three electrode array as well as earlier work with proximity sensing. Given the early, yet limited, success, work with BEI continued. In the next chapter efforts to improve the hardware, expand the array, and limit the necessity of calibration are reprised.

CHAPTER 15

16-Electrode Biomimetic Electrostatic Imager

15.1 Introduction

In this chapter, the design of the previous chapter is expanded to 16 electrodes. The supporting analog electronics and front-end (AFE) have been replaced with two builds of the design described by figure 6.5. To record from all 16 channels, the two AFE's built have to be time-multiplexed. The 2-axis gantry of prior chapters is re-employed for this purpose. It's high degree of repeatability implies that the target is always in the same position, between passes. As such, it may be recorded by two different pairs of electrodes on each pass and data from those passes may be combined as if each of the sensors recorded the same pipe on the same trial. Eight differential electrode channels, requires 16 electrodes and 4 passes in the tank.

15.2 The 16 Electrode Configuration

The receive electrodes are arranged in a linear array with even $1/2''$ (1.27 cm) spacing as shown in figure 15.1. The transmitting electrodes were placed 14'' (35.6 cm) apart centered and collinear with the receive electrodes (also shown in the figure at the edges). The electrodes were originally 24 AWG tin-lead plated steel approximately 1 cm in length as shown in the figure, but were later replaced with platinum tipped electrodes to improve performance. The data appearing in this chapter was collected with these (platinum) electrodes.




Figure 15.1 – *The 16 electrode array submerged in the test tank.*

15.3 Experimental Setup

To establish and verify oceanic verisimilitude, experiments were conducted in a 55-gallon (208 liter) capacity tank filled with 50 gallons (189 liters) of drinking water (0 ppt salinity). Salt in the form of sodium chloride was added in measured amounts to vary the salinity of the test vessel. A Marineland Labs, inc. Instant Ocean Hydrometer [Mar07] was used to record the salinity. The actual salinity of the ocean varies with position, depth, temperature, and season over a range of 28 ppt¹ to 35 ppt [Mar07] in open ocean and as high as an average of 39 ppt in the Mediterranean and Persian Gulf [Lab05]. A final value of 28 ppt was achieved representing the extreme low end (worst-case) of Pacific Ocean conditions.

¹Parts Per Thousand (ppt) by mass ratio with respect to water.




Figure 15.2 – *The experimental configuration consisting of (from left) an instrument rack (sitting on the table), a custom 2-axis motion control platform with the water tank installed below, and a control station. (Inset) The stainless-steel pipe used as the target.*

15.3.1 The Connections

The experimental configuration is shown in figure 15.2 which consists of (from left) an instrument rack (sitting on the table), a custom 2-axis motion control platform (1 meter x 1 meter of travel area maximum) with the water tank installed below, and a control station. The transmit electrodes were connected to an Agilent 33220A Arbitrary Waveform Generator (AWG), which was used to deliver a 2 Vpp^2 sine wave excitation at 1,085 Hz through a 50.0Ω source resistance. The receive electrodes were connected through an Analog Front End (AFE) to an Agilent 33401A Digital Multi-Meter (DMM), which was powered from an Agilent E3631A

²Volts peak-to-peak

Power Supply delivering $\pm 2.5V$. The electrodes were submerged to a depth of 14cm ($\approx 5.5''$). During the trials the laboratory air temperature was 26.1°C .

Inset in figure 15.2, in the lower right corner, is the stainless-steel pipe used as the target. The pipe is mounted to the motion control platform and moved about the tank as recordings are taken from the receive electrode pairs. The X-Y planes of figures 7.5 and 7.6 represent the value recorded when the pipe was in the indicated position.

15.3.2 The Analog Signal Path

Powerline, power supply, nearby processors, airborne radio signals, and other sources of electrical noise can severely hamper the high-impedance, high-gain measurements required for artificial electroreception. A six stage Analog Front End (AFE) was designed for our Biomimetic Imaging System to allow it to disinterr the differential disturbance field signal. In order, they are: (a) platinum electrodes exposed to the tank's salt water, (b) a bridge bias balance network to recenter the signal after AC-coupling, (c) an instrumentation amplifier to provide gain, bandwidth limiting, and conversion from differential to single-ended signaling³, (d) an extremely high-quality ($Q > 30$) active Band-Pass Filter (BPF), (e) a passive Low-Pass Filter (LPF), and, ultimately, (f) a final gain stage provided as part of the data acquisition unit – an Agilent 33401A. For more detail, refer to chapter 6.

A complete AFE handles two electrodes as one differential pair. For the experiments 8 differential pairs (16 electrodes) were used and therefore 8 AFE's were required. Only two physical AFE's were built for the 16 electrode design. To achieve the eight required, the two AFE's were repositioned against a static environment emulating the existence of multiple simultaneous AFE's.

³Since the AFE is duplicated for each channel, minimizing the amount of hardware is important to help control costs. Processing analog signals requires about half the hardware when the signals are single-ended, rather than differential.

15.3.3 The Electrodes

In this implementation, the steel content of [FTC10] was replaced with 99.9% pure platinum. The reduction potential for platinum is much lower than that for iron and, thus, can endure a much higher voltage without reacting with the salt water.

The electrodes are cylindrical in shape and made as small as possible to best emulate a point source in the water. For practical reasons of handling, durability, and ease of construction, the electrodes were fashioned from 30 gauge (AWG) 99.9% pure platinum wire with final assembled dimensions of 267 μm in diameter and 3600 μm in length. The top surface was sealed to the feed system resulting in an electrode working surface area of 3 mm^2 .

An acrylic frame was fabricated to fix the relative positions of two electrodes. Each was placed collinearly ± 2.54 cm from the center point to form a differential receiver basis of 5.08 cm. The frame was placed flush against an acrylic plate to eliminate the back lobe (directivity is covered in detail in chapter 9). The transmit-basis was 10.16 cm.

15.4 Results

Two representative channels of the sixteen appear in figure 15.3. The upper trace is from the center channel (indicated by a lighter/yellow dot on the x-axis) and the lower trace is from one of the left-side channels (it's true position indicated by the darker/blue dot on the x-axis). The plot compares quite favorably with the expected features of figure 7.5. Most notably, the centers are in the correct locations, the pattern has the peaked oscillatory shape, the pattern features remain recognizable across channels as the center moves, the pattern is limited in extent and returns to baseline, and the baselines are unique per channel.

Looking more broadly across the array, consider the channels shown in figure




Figure 15.3 – *Two representative channels of the sixteen electrode (8 channel) imager recording a metal pipe target swept through a salt-water tank.*

15.4. They are recorded from a differential channel centered at 0 cm (upper left), one centered at 7.6 cm (upper right), one centered at 10.2 cm (lower left), and one centered at 12.7 cm (lower right). The more discerning observer may immediately recognize the twin-peaked features expected from the model (figure 7.5), but to enhance the visualization a heuristic approximation, equation (15.1), was applied.

$$\alpha \left[\frac{\sin(\beta(x - \gamma - \epsilon))}{\beta(x - \gamma - \epsilon)} + \frac{\sin(\beta(x - \gamma + \epsilon))}{\beta(x - \gamma + \epsilon)} \right] + \omega \quad (15.1)$$




Figure 15.4 – *One of the differential channels from the sixteen electrode (8 channel) imager recording a metal pipe target swept through a salt-water tank.*

Equation (15.1) consists of two superposed sinc functions that were fit to the datasets by tuning α , β , γ , ϵ , and ω to minimize the mean square error.

Retaining the fit lines and removing the datasets results in figure 15.5. The general correlation between the left-to-right channel positions and their corresponding fit lines is apparent. Also evident is a reduction in signal amplitude with outward movement from the array center.

The channel center may be estimated from the fit curves as the zero point in the first and second derivatives between the two peaks. Plotting these fit line centers results in the green (lighter) piece-wise linear curve of figure 15.6. The actual channel center locations are plotted as the blue (darker) piece-wise linear curve.

A noticeable deviation between the two occurs. The observed channel positions




Figure 15.5 – *The fit lines without their corresponding datasets. The general correlation between the left-to-right channel positions and their corresponding fit lines is apparent.*

are closer in extent than the actual positions. The model of chapter 7 makes the broader prediction that the channel response is maximal near its center based on the assumption that the unit reference dipole is parallel to the receive/transmit axis. For a permanent dipole, this is the case. However, an *induced* dipole, such as a conductive object under the influence of an electric field, merely aligns its dipole moment with the field as shown in figure 15.7.

Unlike the permanent dipoles of the mathematical model, the steel pipe target of our imaging studies is an induced dipole where its dipole moment is always aligned with the curved field lines. This discrepancy is the source of the channel position error, but, as a deterministic phenomenon, may be compensated for.




Figure 15.6 – Reducing each channel to the value detected at its center results in the green (lighter) piece-wise linear curve. The actual channel center locations are plotted as the blue (darker) piece-wise linear curve. Smooth curves are then fitted to the piece-wise linear data.

$$\frac{(\max(y) - \min(y) + \alpha) \sin(\beta(x - \gamma))}{\beta(x - \gamma)} + \omega \quad (15.2)$$

Equation (15.2) is just such a mechanism. The heuristic is a single-sided sinc function fit by tuning α , β , and γ to minimize the mean square error. The fit lines are the smooth curves plotted in figure 15.6. The x and y axis ratios of the blue-to-green smooth curves are used to scale the measured data when performing calculations to create an biomimetic electrostatic image.

Real-time sensor data, matched to the calibration data and then scaled by the fit ratio, results in an image of the space in two dimensions (figure 15.8). In the image the brighter areas represent areas of high probability in which we expect




Figure 15.7 – *Unlike the permanent dipoles of the mathematical model, the steel pipe is an induced dipole where its dipole moment is always aligned with the curved field lines. This discrepancy is the source of the channel position error, but, as a deterministic phenomenon, may be compensated for.*

the pipe target to be found. Given that the brightest areas are within the pipe’s actual boundary (the yellow circle with red x center), the calibrated system is quite accurate and effective.

Looking more broadly at imaging performance, figure 15.9 shows four example captures of the pipe target in various positions in front of the array. The sequence of capture moves from the top-left pane, left-to-right then top-to-bottom (non-raster).

The accuracy of figures 15.8 and 15.9 holds principally at the center of the array where a maximum number of channels are participating in the detection. As the target approaches the edges of the array (and moves beyond it) the accuracy degrades. It exhibits the same parabolic degradation function experienced by the 3-electrode imager in figure 14.15. Visually, this manifests as an increase in the white area of the image – figure 15.10.




Figure 15.8 – *A top-down image of the tank in two dimensions using real-time sensor data, matched to the calibration data, and scaled by the fit ratio. The brighter areas represent areas of high probability in which we expect the pipe target to be found. The pipe’s actual boundary is indicated by the yellow circle with red x center.*

15.5 Conclusion

In this chapter, we have presented our evaluation efforts regarding the construction of a 16 electrode (8 channel) biomimetic electrostatic imaging platform capable of visualizing submerged objects in a salt-water environment.

In the previous chapter, the position accuracy was no better than 10cm. The imager described here achieves 2.5cm accuracy on the center line and degrades to 5cm accuracy at the peripheral channels. The output power remains unchanged between design iterations. Curiously, this imager shows little signal degradation over the 20cm of range tested. This is likely a result of the limited volume of the water tank. The walls of the container structure serve to focus and enhance the




Figure 15.9 – *Various top-down images of the tank in two dimensions using real-time sensor data, matched to the calibration data, and scaled by the fit ratio. The brighter areas represent areas of high probability in which we expect the pipe target to be found. The pipe’s actual boundary is indicated by the yellow circle with red x center. The imager tracks the target as it traverses the tank.*

field – much like a fish’s body would – and the target-specific calibration improves the minimum acceptable SNR.

In the next chapter, the number of channels and electrodes will be increased for a third time. The introduction of substantial volumes of hardware will permit simultaneous sampling across tens of channels and we will retire our dependence on pre-calibrated target models, preferring, instead, to employ online calibration for targets in an unstructured space. It will complete our journey towards a truly capable “imaging” system.




Figure 15.10 – *The ability of the imager to localize the target degrades as the target moves beyond the horizontal extent of the array and, therefore, exceeds the sensor’s high confidence area.*

CHAPTER 16

17-Electrode Biomimetic Electrostatic Imager

16.1 Overview

In order to pursue the ultimate objectives of unstructured imaging, the single channel reference design of chapter 6 was modified to support digital control, replicated 16 times, and combined into one printed circuit board (PCB) assembly called the AFE17. Due to its difference-linear handling of electrodes, the 16 channels of the AFE17 process signals from 17 ocean-submerged electrodes. To illustrate this, consider a row of 17 dots. Each of the dots represents an electrode. There are 16 spaces between the 17 dots. These spaces represent the differences calculated by the AFE17.

In this chapter the design and utilization of the AFE17 is explored in detail, including a validation of its performance in the laboratory. The accompanying electrode array is similarly designed, deployed, tested, and then optimized.

16.2 The AFE17 Single Card 16 Channel Receiver

The AFE17 is a large complex piece of hardware made from 1,073 electronic components and 11 mechanical components. It is about 30 cm (\approx 12 inches) in length and under 7.6 cm (3 inches) in height. It is designed to fit into a card frame system two standard rack units 8.9 cm (3.5") in height.

A simplified schematic of the AFE17 16 channel receiver card appears in figure



Figure 16.1 – *The AFE17 as built. Shown with its shielding covers removed. Each of the square areas contain a complete analog front end. In the distance, near the edge connector fingers, the large power filtering components are visible.*

16.2. The card's inputs consist of power and digital control signals. In terms of power requirements, the analog sections are bipolar resulting in four voltage rails: ± 2.5 V and ± 5 V. The power rails are isolated from the external supply by a large second-order passive low-pass (LC) filter. The reference potential ("ground") is not isolated to minimize ground loop vulnerability and ground bounce.

The chapter 6 single element receiver (analog processing unit) was modified to allow for realtime digital control of its key parameters. The seven resistive elements highlighted in figure 16.3 are the key resistive elements in the single element receiver replaced by digitally controllable resistors. This provides the host software real-time control over the bias, gain, and center frequency of the channel.

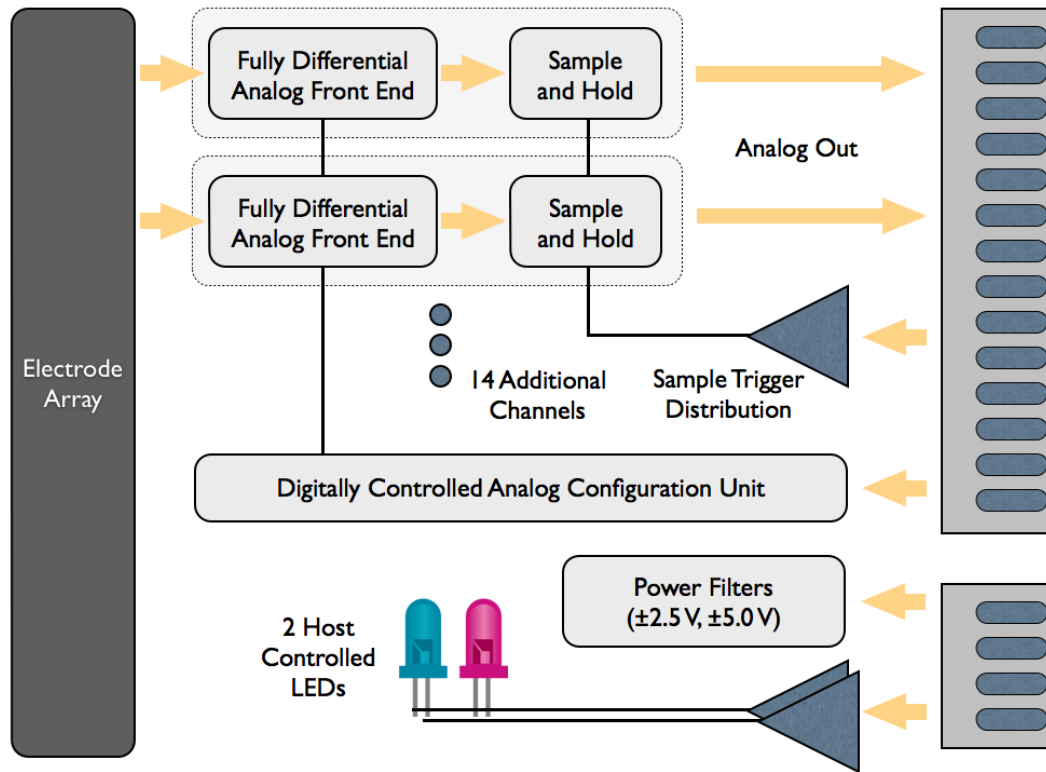


Figure 16.2 – A simplified schematic of the AFE17 16 channel receiver card. The card’s inputs consist of power and digital control signals. The analog outputs from the 16 analog processing units are available on the card’s edge connector.

16.2.1 The Control Bus

The resulting control requirements make use of eight¹ 10 bit digital resistors per analog processor. This places a large requirement for electrical output connections on the host, on the connector to the host, and on the AFE17’s internal routing. To alleviate the physical demands of these connections, the resistor configuration data is bussed to all of the digital resistors and only the chip select lines are routed point-to-point. The chip select lines are controlled from an on-board bus, designed by the author, called the CB2 (Control Bus, version 2).

¹The gain control unit uses two in series as a single 20 bit resistor

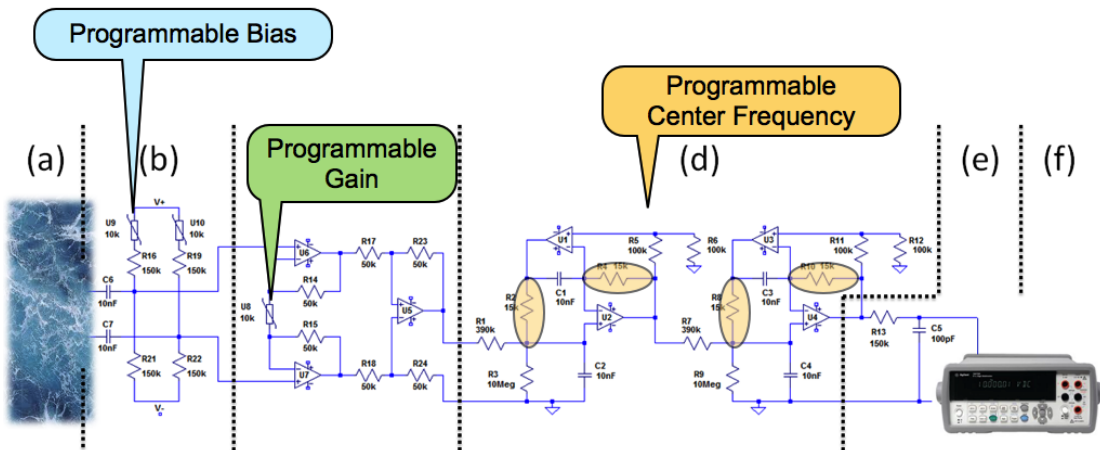


Figure 16.3 – Key resistive elements in the single element receiver were replaced by digitally controllable resistors. This provides the controlling software real-time control over the bias, gain, and center frequency of the channel.

16.2.1.1 Ragobot: the Origin of CB2

The CB2 bus was originally invented to service a similar problem on-board the Ragobot (figure 16.4). The Ragobot [Lee06a] [Fri07] was a product of the dramatic shift in wireless sensor network architecture that was occurring in the mid-to-late 2000’s. Consider the classic target application of such a network – monitoring. Whether it’s a forest fire, a platoon of soldiers, or a cosmic phenomenon, the overriding approach involved cheap nodes deployed throughout the area of interest. However, static sensor nodes suffered from numerous drawbacks. The pragmatic issues of network deployment, coverage holes, and sub-optimal density lent great credence to the addition of mobility.

Implementing these motile and mobile sensor networks required a different design philosophy from traditional robotics. By convention, robotics emphasized the capability of a single robot; in contrast, mobile sensor nets leveraged teams of coordinated entities. Estrin, et.al. concluded that two key requirements emerge: “support for very large numbers of unattended autonomous nodes and adaptivity to environment and task dynamics” [EGH99]. In designing mobile sensor network




Figure 16.4 – *The Ragobot was designed by the author and David C. Lee to enable research in mobile, motile, wireless sensor networks. This image first appeared in [Lee06a].*

devices, we discovered that the former is actually an implicit benefit while the latter imposes significant design challenges.

Ragobot was our proposed solution. Smaller than most fully-navigable robots and more heavily instrumented (video capture, audio capture, processing, and playback, IR collision avoidance, IR cliff detection, RFID read/write, inertial navigation, and more), the Ragobot could mechanically outperform (capable of 34.9 degree vertical climb) all of its known contemporaries.

Packing all of this functionality into such a small platform meant that the limited I/O capabilities of the central processor could not be wasted on low-bandwidth, latency tolerant, configuration functions. Over forty output-only control signals

were required to manage power distribution on-board the battery-powered mobile sensor platform. The solution was CB2.

16.2.1.2 CB2 Operation

Using as few as three pins on the host processor, the CB2 bus is expandable to an infinite number² of output-only pins. CB2 is a 3-wire serial bus that controls a number of daisy-chained HC595-series shift registers. The HC595 shift registers are 8-bit serial shift registers with output latches, which allows the outputs to update simultaneously once the desired configuration data has been shifted in. The ability to update simultaneously is necessary to prevent glitches in the system when the configuration is shifting into the registers. Each bit of the shift registers (the configuration) controls, on the Ragobot, either an LED or a power gating switch. In the BEI, each bit controls a select line on a digital resistor enabling them to accept the incoming resistance set value or ignore it.

The CB2 driver firmware represents the external shift registers as virtual output ports assigning each one an address in memory. The application firmware writes to these “ports” in a manner no different than its native hardware ports. When the application software is ready it executes the CB2UPDATE command (calls the CB2UPDATE() function) and the virtual ports will be turned into physical ones by distributing their state to the corresponding CB2-attached hardware registers. Application-specific higher-level Application Programming Interfaces (API) often encapsulate the CB2’s modify-then-update behavior into an atomic operation.

To use the CB2 bus, the configuration is first serially shifted through the shift registers using a clock and a data signal. Then, after the complete configuration is shifted out, the update signal changes the output of the registers to reflect the new configuration. Updates for values that do not change do not produce any glitches

²The CB2 bus expansion is limited by the amount of host memory available to the CB2 driver firmware and by the tolerable latency of attached devices.

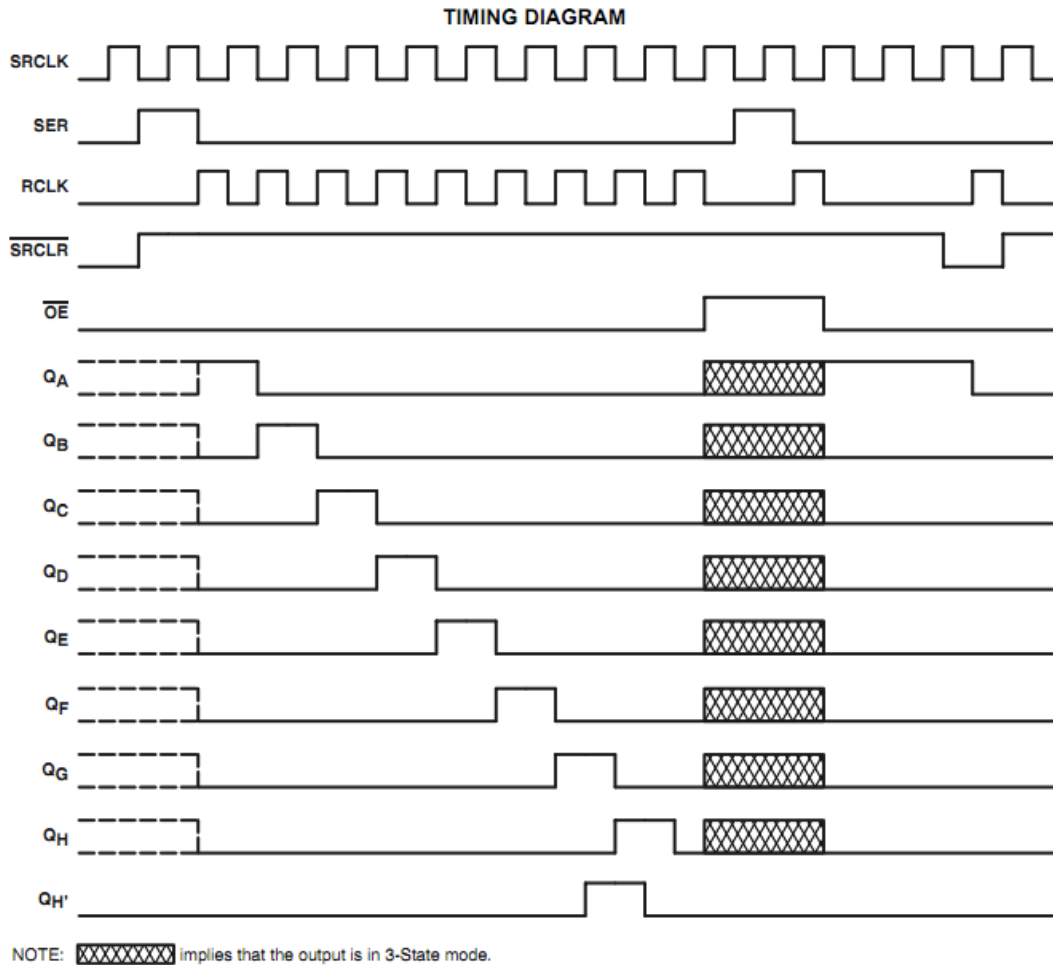


Figure 16.5 – *The CB2 bus protocol is a simplified version of the more generic 595 series 8-bit shift register with 3-state parallel output registers protocol shown here. The CB2 requires only the SER, SRCLK, and RCLK signals, which are renamed DataIn, Clock, and Update, respectively. Excerpt from [Ins09].*

or other artifacts – the value is simply held undisturbed. The CB2 bus protocol is a simplified version of the more generic 595 series 8-bit shift register protocol shown here (figure 16.5).

The CB2 requires only the SER, SRCLK, and RCLK signals, which are renamed DataIn, Clock, and Update, respectively. There is a fourth, optional, CB2 signal: Reset, which completely flushes the hardware state and asynchronously resets the CB2 hardware. Use of this command signal is unnecessary as the hardware starts

in a known deterministic state, but some hosts may experience glitching on their outputs during boot that may disturb the state of the CB2. Use of the Reset signal resolves this. The remaining 595 series signals are handled as follows: SRCLR is always held high, OE is always held low, and the last 595's Q_H' data pin is routed to an expansion interface connector (effectively left unconnected, e.g. high-Z load).

The CB2 firmware reloads the entire register chain on each update command – a data-independent deterministic behavior. As a result, the maximum speed oscillation speed of any individual control output on the CB2 bus is the same regardless of its proximity to the host (position on the bus's register chain). That frequency, f_{CB2} , is a function of the number of host instructions it takes to toggle the CB2's Clock data line (output from the host), n_{toggle} , the number of host instructions it takes to determine, set, and drive the data bit, n_{dsd} , and m , the number of shift registers in the CB2 bus. The exact formula is equation 16.1:

$$f_{CB2} = \frac{f_{cpu}}{2 n_{toggle} (8m + 1) + n_{dsd} 8m} \quad (16.1)$$

The $2n_{toggle}$ term appears because the clock edge must both rise and fall for each transmitted bit of data and the +1 term in the denominator expresses the update command at the end of the entire data shift operation. This assumes that the processor clock rate, after scaling by the instruction delays (as indicated in equation (16.1)), is less than the processor's maximum I/O clock rate and less than the CB2's maximum data clock rate.

The maximum clock rate of the specific 595 series registers used in the Ragobot and BEI is 25 MHz [Ins09], which results in an absolute maximum CB2 output pin toggle rate of 385 kHz, given an infinitely fast host processor and a bus 8 registers long. 64 control lines, each independently capable of operating at hundreds of kilohertz, is more than sufficient for all envisioned control applications.

Generically, the limit is described by equation (16.2):

$$f_{CB2,max} = \frac{25}{8m + 1} \text{ MHz} \quad (16.2)$$

16.2.2 Sample and Hold

The sixteen individual single-element receivers (chapter 6) are each followed by a sample-and-hold unit whose trigger signal is provided through the edge-card interface. The sample-and-hold unit is disabled (e.g. stays in sample mode) by default. The analog outputs from the 16 analog processing units are available on the card's edge connector. It is up to the attached host to initiate the data hold by triggering the unit through a signal pin on the connector interface. It is also a host responsibility to release the analog processing units from hold to resume sampling.

This triggering signal is bussed to all of the channels and is routed using a length-match among channels. The total variance in path length between the most extreme of the 16 channels is less than 2.5 mm (0.1"), which equates to a maximum sample aperture variation of only 13 ps. The sample-and-hold can effectively guarantee simultaneous sampling across all AFE17 channels, capturing the incoming signal state and allowing the external data converter to sequence through each of the individual channels and convert at will. The protocol consists of the host-initiated hold signal, sequential conversion across all of the channels, and then the release of the hold signal (resume sensing).

16.2.3 Bias Control

The bias unit moves the zero point in the amplitude domain. Zeroing the amplitude allows for higher gains to fit within the limited dynamic range available. High gains translate directly to longer sensing/communication range.

16.2.4 Gain Control

Achieving the highest gain without saturation requires dynamic control over the exact gain value to react to changes in the environment or in the signal itself. For this reason, gain control is a very common function in analog front ends. In our case, the gain control was implemented in the low impedance internal path of the instrumentation amplifier (section c of figure 6.5). The gain transfer function is equation (16.3)

$$\text{Gain} = \frac{2R_c}{R_g} \quad (16.3)$$

where R_c is the fixed resistance of the instrumentation amplifier and R_g is the gain controlling variable resistor. R_g is implemented as two ten-bit digitally controlled resistors in series to effect 20 bit control over the gain function.

The gain response appears in figure 16.6. It is non-linear, so the high degree of linear control (20-bits = 1,048,576 codes!) maps well enough onto the non-linear gain space enabling near-linear control of the gain multiplier³. The maximum gain is $10,241\times$ and the minimum gain is $6\times$. Our prior experimental work suggested that control around $10\times$ would likely be required, so the control range is optimized for this region. Near this gain level the step resolution is less than $1/100\times$ – a level of control unnecessarily high as it is exceeded by the noise floor.

The actual and entire gain control configuration space as measured on an AFE17 operating dry in the laboratory appears in figure 16.7. To the left edge of the figure the slight non-linearity is perfectly consistent with the modeled expectation of figure 16.6 as the output voltage is increasing at a slower rate than the input voltage (e.g. less marginal gain). To the right of the figure, the shelf like area is a result of insufficient dynamic range in the analog signal path of the AFE17 resulting in saturation and clipping.

³Effecting linear control requires a non-linear transfer function stored in a lookup table in the BEI-68's firmware.




Figure 16.6 – *The gain of an instrumentation amplifier is inversely proportional to the value of its gain control resistor and highly non-linear. The gain control unit in the BEI platform is optimized for precise control around 10×.*

16.2.5 Frequency Control

Control of the filter center frequency required the replacement of four resistive elements due to the Fliege-type architecture. In order to avoid disturbing the other aspects of the filter transfer function, the four resistive elements must be kept at the same relative value and adjusted globally to control frequency. Individual control over the resistors is provided to zero out resistive channel-to-channel variations in the factory calibration to achieve this relative accuracy. In practice, we found the laser trimming and manufacturing controls, put in place by Analog Devices, to be more than sufficient. Setting all four of the AFE channel’s digital resistors to the same code resulted in the best filter behavior. Individually adjusting the resistors’ codes relative to one another, even as little as one code, resulted in reduced performance.

Among channels, the situation is remarkably different. Figure 16.8 shows




Figure 16.7 – *The actual and entire gain control configuration space as measured on an AFE17 operating dry in the laboratory.*

all sixteen channels of one of our AFE17 cards. The diversity in the channel’s center frequencies is apparent. The mean and standard deviations are indicated. Given that the filter passband is only 37 Hz wide and the flattest part of the passband is less than 5 Hz wide, the diversity in the channel’s center frequency has a pronounced effect on imaging performance.

To correct the center frequencies to the same value, the channels must be moved. The Fliege filter transfer function solved for frequency and differenced by itself results in equation (16.4).

$$f_1 - f_2 = \frac{R_2 - R_1}{2\pi CR_1 R_2} \quad (16.4)$$




Figure 16.8 – Fresh from the factory, an AFE17’s 16 channels have a wide distribution of center-frequencies due to component tolerances. These center frequencies may be corrected to within 1 Hz using software control over the filtering hardware.

Here, R_1 and R_2 are the filter resistances, C is the filter capacitance, and f_1 and f_2 are the filter center frequencies. The equation may be simplified in analyzing the AFE17 as the term $(R_2 - R_1)$ is really an just an expression of the control resolution over resistance (R_{step}) and given that R_{step} is small compared to either R_1 or R_2 , equation (16.4) simplifies to equation (16.5):

$$\Delta f = \frac{R_{step}}{4\pi CR} \quad (16.5)$$

For the AFE17, C is 10 nF and R is 10 k Ω . Therefore the center frequency may be corrected to within 1 Hz using software control over the filtering hardware.

In simulation looks like figure 16.9 and in practice it looks like figure 16.10. Note the fine frequency resolution, high quality filter function, and non-linear change in frequency with control input – all as expected.

Although the frequency correction value may be accurately anticipated from




Figure 16.9 – *The AFE17 provides digital control over the frequency response of its contained filters. 11 possible frequency responses are shown here, distributed evenly over the 1024 possibilities. The AFE17’s 10-bit control results in a frequency control resolution of just under 1 Hz.*

measurements and the application of equation (16.5), in practice, the algorithm in the AFE17 driver operates closed loop with an applied reference signal. The channels are iteratively adjusted to the resistance value that maximizes input signal. This has the effect of not just bringing all channels to the same frequency, but compensating other manufacturing variations as well.

16.2.6 Annunciation Unit

There are two Light Emitting Diodes (LED) mounted on the AFE17’s faceplate. These diodes are buffered and defaulted (to the on state), but controlled from the external host. The design intention is that the LED’s provide a diagnostic indication that the card is properly seated and powered and will be switched off (overridden) by control signals from the host processor, once it recognizes the installed AFE17 card. In practice this might appear as a brief flash of light when the host system is started. The indicator lamps are then available for host




Figure 16.10 – *The AFE17’s frequency adjustment function validated in hardware. Note the fine frequency resolution, high quality filter function, and non-linear change in frequency with control input – all as expected.*

system functions, such as flashing briefly for each sampling event or indicating calibration-in-progress.

16.3 Evaluation

Having covered the generation and detection of a background electric field in prior experiments, we now turn our attention to characterizing the effect of conductive objects on the field that we have established with the AFE17. Recall, we entitle this effect the *differential disturbance field*.

16.3.1 Experimental Configuration

A test interface circuit board was designed and manufactured to allow a single AFE17 card to easily connect with standard laboratory bench top instruments. The test configuration can be seen in figure 16.11.

Working from the top of the figure, two Agilent E3631A 80 Watt maximum continuous output power supplies are mounted side-by-side to provide the four voltages needed for operation (± 2.5 V, ± 5.0 V), an Agilent 33220A arbitrary waveform generator provides the transmit signal and output amplification, and a multi-core computer (equipped with sampling hardware) handles data acquisition, digital processing, and visualization functions.

The PC-hosted data acquisition hardware is from National Instruments. It consists of a model PCIe-6323 X-series DAQ unit attached by a shielded micro-coaxial cable to a Shielded Cable Breakout (SCB) interface circuit board (NI-SCB68). The SCB68 is visible in figure 16.11, just above the AFE17 mounted in the test fixture.

Protruding from the left edge of the figure is a cable attached to the AFE17's electrode interface port. At the other end of the cable is a 68-electrode array, designed by the authors, that we call the F68. The ABS plastic frame (in the shape of a fish!) provides positioning, structural support, and electrostatic focusing [FTS10]. The F68 appears in the upper portion of figure 16.12 as constructed. The exact dimensions and configuration of the F68 appear in the lower portion of the figure.

The F68 consists of four rows of 17 electrodes. Each row is connected to a single AFE17. The transmit electrodes are located on the midline between rows 2 and 3. For this study only a single AFE17 was involved, and only electrodes 1 and 2 (channel 2) of row 2 was used for detections.

The F68 is placed against a long wall of a 55 gallon (208 liter) capacity tank



Figure 16.11 – *A single AFE17 installed in a test fixture with a rack of support equipment. (from top) Two power supplies provide the four voltages needed for operation, an Arbitrary Waveform Generator (AWG) provides the transmit signal and output amplification, a multi-core computer (equipped with sampling hardware) handles data acquisition, digital processing, and visualization functions.*

filled with 50 gallons (189 liters) of drinking water (0 ppt salinity). Sodium Chloride was added in measured amounts to raise the salinity to 31.5 ppt. The tank is placed beneath our custom 2-axis gantry and a reference 2.54cm (1”) steel pipe is attached to the gantry carriage such that it may be freely repositioned by the moving gantry. A schematic of the experimental configuration appears in figure 16.13. The gantry used in laboratory testing is the same as the one involved in the ocean trials and makes an appearance in figure 7.2.




Figure 16.12 – *The F68 electrode array. 68 electrodes are arranged in four rows. The middle row is for transmit (excitation) electrodes.*

16.3.2 Model Validation

In [FHT11], we modeled the differential electrostatic field. With the advent of the AFE17, experimental validation of the model is possible. Data from trial #P20 of the Carderock dataset appears in figure 16.14 (*bottom*) and the corresponding simulation from [FHT11]’s model appears above it (figure 16.14 (*top*)). Note the strong agreement and conforming $1/d$ decay with distance (where d = distance from the sensor electrodes) indicated by the red trend line.




Figure 16.13 – *A reference metal pipe is mounted to a custom 2-axis gantry and raster-swept through the tank in front of an array of receive (sensing) electrodes with two interspersed transmit (excitation) electrodes.*

It should be noted that the model has not been scaled to real units and only represents the abstract shape of the surface expected from the differential measurement of a disturbing dipole. And, the scale and rotation of the two surfaces plotted (modeled and observed) are close, but far from identical, which explains any visual dissimilarities.

In this experiment (and in the rest of the experiments reported in this section), a reference metal pipe (diameter of 1”) was raster-swept through the tank. The resulting differential sensor voltage is plotted. In the figures, the X-Y plane corresponds to the position of the pipe (the sensor is stationary), while the Z axis

references the amplitude of the signal recorded by the stationary sensor when the pipe was positioned at the indicated location.

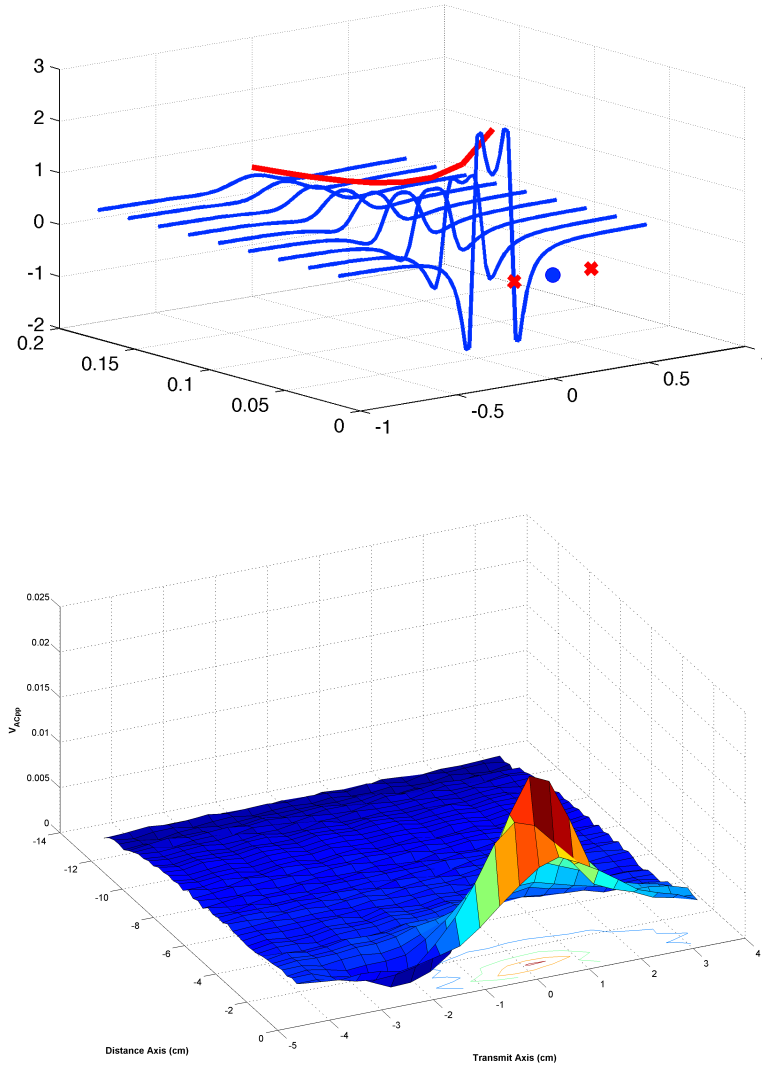


Figure 16.14 – A reference metal pipe is raster-swept through the tank. The resulting differential sensor voltage is plotted. The X-Y plane corresponds to the position of the pipe (the sensor is stationary). A model demonstrating expectation (top) and data from an actual laboratory trial (#P20) (bottom). Note the strong agreement.

For trial number P20, the excitation (transmit) frequency was 813Hz. This was the measured peak sensitivity of the selected receive channel (AFE17 differential receive channel number 2). The excitation amplitude was 1 V_{pp} delivered by the

Table 16.1 – *Trial Data Summary*

| Trial Number | Source V_{pp} | Waterline V_{pp} | Power (mW) |
|--------------|-----------------|--------------------|------------|
| P8 | 13.00 | 9.76 | 158 |
| P9 | 12.00 | 9.12 | 132 |
| P12 | 9.00 | 7.04 | 69 |
| P14 | 7.00 | 5.52 | 41 |
| P16 | 5.00 | 4.08 | 19 |
| P18 | 3.00 | 2.48 | 7 |
| P20 | 1.00 | 0.84 | 0.67 |
| P22 | 0.20 | 0.17 | 0.03 |

Agilent AWG. A waterline voltage of 0.84 V_{pp} was observed during excitation by a Rigol 100 Mhz digital sampling oscilloscope set to 2 V/div. Excitation was delivered from the middle pair of four transmit electrodes on the F68, resulting in a transmit axis length of just 2.80 inches. The other two excitation electrodes were left floating. 20 data frames were recorded at each position and averaged. Each data frame consisted of 100 samples captured at 15,625 samples-per-second. The center point of the receive axis is located 2.03 cm (0.8”) to the left (towards the head) and outside of the transmit axis and 1.52 cm (0.6”) above the transit axis.

16.3.3 The Effect of Output Power

Table 16.1 summarizes the results from this series of experiments, including the aforementioned trial #P20. The output power was decreased between successive trials until the minimum output power of the AWG was reached. Viewed raw, along the range axis (the 2D Y-Z view plane), the trial datasets appear as they do in figure 16.15. There are long-term-ionization and surface turbulence effects that disturb the location of the baseline over time.

Improved software and hardware was applied to conduct these studies. Namely, replacing the former Python scripts for automating instrument control with Mat-

lab™scripts and replacing the Agilent 34410A Digital Multi-Meters (DMM) with National Instruments PCIe-6323's, resulted in runtimes almost 5 times faster. These speed improvements had an unintended side-effect. Since the NI-DAQ card was so much faster than the DMM, the gantry would barely pause between spatial sampling points. This resulted in significantly more surface turbulence – and therefore more dislocation of the electrostatic differential baseline.




Figure 16.15 – A range-axis (Y - Z) view of the electrostatic disturbance field caused by a metal pipe for various output power levels. The y -axis is the received RMS amplitude versus its baseline.

To better visualize the problem, figure 16.15 shows the dataset of figure 16.15 viewed end-on (X - Z view plane) and sliced through at $Y = -0.5$ cm. The predicted pattern from the model is apparent, but is accompanied by a baseline shift in amplitude that is correlated with the amplitude of the excitation (output power).

The baseline shift may be compensated ex-post facto by taking the discrete-




Figure 16.16 – *The dataset of figure 16.15 viewed end-on (X-Z) and sliced through at $Y=-0.5$ cm. The y-axis is the received RMS amplitude versus its baseline.*

spatial derivative with a window size of 1 point. Running the trial datasets over this filter function results in figure 16.17, where it is still the dataset of figure 16.15 viewed end-on (X-Z) and sliced through at $Y = -0.5$ cm. After compensation via a discrete-spatial derivative filter, the effect of output power on the differential disturbance field is apparent, proportional, and monotonic. The shift is effectively removed. All of the following datasets have this filter applied.

16.3.4 Sensing Efficacy

To illustrate the efficacy of the AFE17 as a sensor, many reference scans (scans without targets present in the tank) were performed. The results from a typical scan appear in the top pane of figure 16.18 – this time in the Y-Z view plane. For comparison trial #T3 (performed under the same conditions as the previously referenced trial #P20, with the exception of an increased output power of 320 mW) with the metal pipe target present is shown in the bottom pane of figure 16.18.

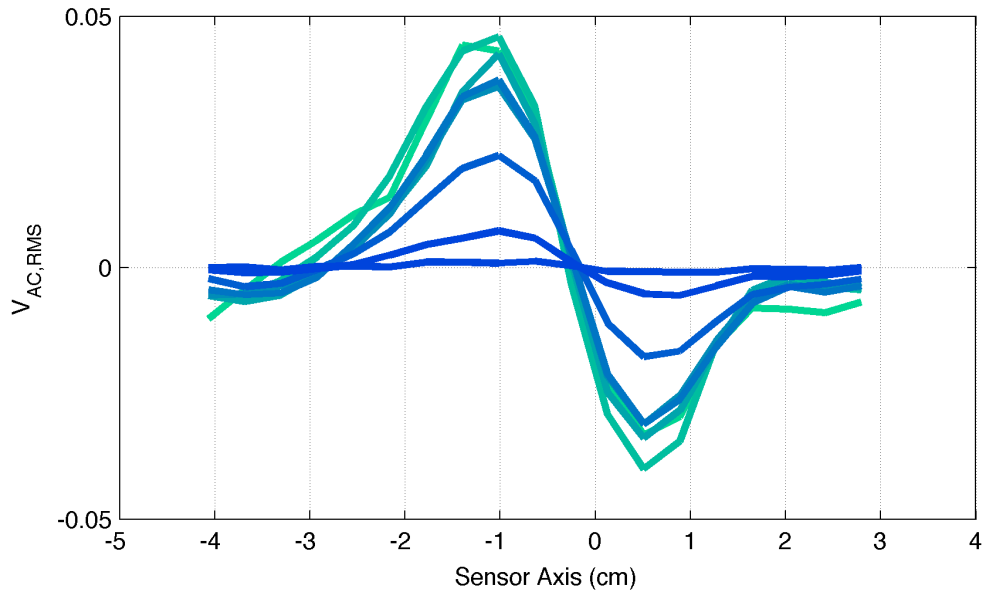


Figure 16.17 – *The dataset of figure 16.15 viewed end-on (X-Z) and sliced through at $Y=-0.5$ cm after compensation via a discrete-spatial derivative filter. The effect of output power on the differential disturbance field is apparent, proportional, and monotonic. The y-axis is the received RMS amplitude versus its baseline.*

The target presence condition is clearly indicated. There are theoretically no false-positives in BEI and we never observed any.

As the AWG composites its output signal from a sequence of discrete voltage steps (the buffered and filtered output of a digital-to-analog converter), increasing the output gain of the AWG also increases its output noise. It was later determined that the noise visible in the long-tail of figure 16.19(*top*) was correlated with the AWG’s output power. Increasing the output power increased the amplitude of this noise band, reducing the output power reduces it, and disabling the output returns the amplitude to its near-zero baseline.

This presents an interesting situation whereby decreasing the output power can improve sensing range since the signal is not decreasing as rapidly as the noise. In the trials indicated, the SNR actually improves noticeably with decreasing output power up to a point. This phenomenon is explored quantitatively below. It




Figure 16.18 – *A comparison between a scan with no target present (top) and the reference metal pipe target present (bottom). The y-axis is the detected V_{RMS} versus its baseline value, which explains the negative values.*

is sufficient, for now, to qualitatively observe the difference in trial #P8 (figure 16.19(*top*)) and trial #P18 (figure 16.19(*bottom*)) in order to build an intuitive

understanding of the results that now follow.




Figure 16.19 – *The profile of the differential disturbance field as observed in the laboratory tank. The increase in transmit power results in an increase in output noise (top) as compared with a lower output power level (bottom). The x-axis is range in centimeters and the y-axis is V_{RMS} as in prior figures.*

16.3.5 Detection Range Composite Findings

We define detection range as the maximum extend of the differential disturbance field's primary lobe where a $1 \mu\text{V}$ band of amplitude (post filtering and amplification) defines a clean contour in space – a contiguous region without isolated




Figure 16.20 – *Output voltage (black) and output power (blue) versus the maximum detection distance observed by the F68. The point of optimality is indicated by the dashed red line.*

islands of included space⁴. Naturally, this is somewhat subjective and the cutoff (how much noise to consider visually unacceptable in the contour) was determined by hand introducing some uncertainty to the measurement. It is the opinion of the authors that the non-monotonicity in the associated figures is likely the result of our cutoff determinations, rather than being indicative of an underlying phenomenon. The detection range provides the final parameter for end-to-end system evaluation.

This data is compiled into figure 16.20. The blue (lighter) trace, which apexes left-most in the figure, is a plot of output power versus the maximum observed detection distance. The black (darker) trace is a plot of the output voltage versus the maximum observed detection distance. It is immediately obvious that increasing the output power beyond a certain point does not increase range as the transmitter output noise is increasing at a rate faster than the improvement in signal amplitude.

⁴Examples of such contours appear in figure 16.28




Figure 16.21 – The received voltage (blue) and output power (green) versus the output voltage into the water. The point of optimality is indicated by the dashed red line.

From figure 16.20, the optimal Q-point is an output voltage near 1 V_{pp}. This output level is indicated by the red vertical line in figure 16.21. The top (blue) trace of the figure is the received voltage versus the output voltage and the lower (green) trace is the output power under the same conditions. It is extremely encouraging to observe that the best performance is achieved at the lower end of the output power curve. The overriding conclusion is that you should apply “just enough” output power to create an observable signal – precisely, the ideal situation for energy-limited remote underwater sensors and vehicles.

16.3.6 Channel Sensitivities

To further improve performance, the channels of the AFE17 were calibrated a common reference level. The maximum amplitude of the disturbance field, across the sensing array, for different output powers is plotted in figure 16.22. The peak sensitivities occur near the excitation electrodes and, among these, the signal

entry electrode is more sensitive to targets than the excitation reference electrode. This explains the biological observation that the electric organ (transmitter) is typically oriented head-to-tail (strongest excitation near the head) and typically is not located past the operculum (the gill cover near the head). This positions the area of maximum sensitivity around the head and mouth where it would be most useful in hunting and collision avoidance.




Figure 16.22 – *The channels closest to the excitation (transmit) electrodes register the strongest changes to the differential disturbance field.*

It is clear from the figure that the choice of channel (location) matters. In the subsequent data, channel #12 was used.

16.3.7 Improved Electrodes

Having gone as far as possible with analog configuration and calibration of the AFE17, we turned to the electrodes themselves. A revised array (figure 16.23) was constructed making use of stainless-steel 1/4-inch round electrodes to replace the reacting steel of those used with the F68⁵. Reducing the output power and

⁵A revised F68 using non-reactive metal electrodes is currently in manufacturing



Figure 16.23 – *A revised array was constructed making use of stainless-steel 1/4-inch round electrodes to replace the reacting steel of those used with the F68.*

compensating the receiver-side gain can substantially improve the SNR by short-circuiting the transmitter noise problem (see figure 16.19). This has a number of interesting and notable effects on range, efficiency, and optimal output power.

Figure 16.24 recreates the results of figure 16.20 with the new electrodes and newly calibrated AFE17. The maximum detection range has almost doubled and the transmitter output noise is no longer the limiting factor due to the higher receiver gain and better dynamic range made possible by bias calibration. The upper x-axis of the figure is the output power in Watts.

As with figure 16.24, figure 16.25 is a recreation of figure 16.21 with the new test setup. The received voltage is now better correlated with the output power. The implications of these new findings will become apparent when a compared




Figure 16.24 – *Output voltage (black) and output power (blue) versus the maximum detection distance observed as in figure 16.20, but with new electrodes and a newly calibrated AFE17.*




Figure 16.25 – *The received voltage (blue) and output power (green) versus the output voltage into the water utilizing the revised electrodes.*

with the earlier F68 experiments.

16.4 Design Conclusions

Comparing the two trial groups, those with the F68 and those with the improved (“revised”) electrodes, results in a number of interesting conclusions. For illustration consider figure 16.26. The figure normalizes the two designs by expressing a measure of efficiency – how much sensor range is achieved for a given amount of output power. The operational energy of the receiver is largely invariant to receiver gain (it is dominated by quiescent power) allowing us to reduce overall energy efficiency to an expression exclusively determined by output power.

In the figure, the F68 trials begin with an initial advantage due to the lower receiver-side noise (given the lower receiver gain). However the deficit imposed by the transmitter output noise quickly overcomes this limited advantage. The combination of better electrode chemistry, better dynamic range tolerance (from calibration), and larger gain (from configuration) results in an efficiency curve that can reach almost twice the range, albeit at substantially greater output power.

The natural conclusion is that the output power should only be increased after the signal has reached the receiver’s noise floor and no further receive side gain is possible (either for lack of an absolute capability, or perhaps limited by dynamic range) – or more simply: receive gain is superior to transmit gain (an unusual phenomenon).

The range cannot simply be scaled indefinitely. Greater range requires greater output power, but greater output power, in turn, produces greater ionization in the intermediate salt-water environment. This results in greater conductivity and decreasing marginal utility for incremental output power.

Quantitatively, this is explored in figure 16.27 where the voltage-current characteristic is plotted. The dataset was created by sweeping output power with the Agilent AWG above and recording output current with an Agilent 33440A DMM as the transmitted signal transited our test tank at 31.5 ppt salinity. It shows




Figure 16.26 – A comparison of the two designs through a measure of efficiency – how much sensor range was achieved for a given amount of output power.




Figure 16.27 – The output voltage-current characteristic. The non-linearity is expected from the situational chemistry. This dataset is taken with the improved stainless-steel electrodes.

an increase in water conductivity for increasing output power. This is expected and covered in detail in our prior works [FTS10] [FHT12b]. It serves to explain the asymptotic behavior of the received voltage trace of figure 16.21. As the water conductivity increases, the required output power improves and a greater proportion of the supplied energy goes into ionizing the environment than being stored in the field. It also explains the beginning of the downward trend at the right of the revised trials of figure 16.26.

Since range cannot be improved further with output power, it is incumbent on us to examine other mechanisms. The next logical parameter is transmit electrode spacing. Figure 16.28 illustrates the effect of the transmit axis. The difference in detection range is significant. When the transmit axis is just 7.1 cm (2.8”), figure 16.28 (*top*) results. When the entire F68 body is employed at 18.3 cm (7.2”), the bottom contour results. Each of the contours are 1 μ V amplitude bands determined by the previously explained method.




Figure 16.28 – *The difference in detection range when the transmit axis is just 7.1 cm (top) and across the full length of the fish body at 18.3 cm (bottom).*

16.5 Conclusion

This chapter has detailed ocean and laboratory validation of BEI and the AFE17, revealed the efficacy of a number of design improvements made to BEI hardware and software, and conveyed a number of design axioms which will drive our future experimentation. The improved BEI system increased detection range by 97% and range-power efficiency by as much as 150% without a corresponding increase in the false positive rate. Both the differential and differential disturbance fields were shown to exist at greater than $7\times$ their respective electrode separation distances – the furthest reported in the literature to date.

CHAPTER 17

68-Electrode Biomimetic Electrostatic Imager

17.1 Overview

The previous chapter explored the performance of just one of the AFE17's many channels operating in an unstructured space against a reference target. In this chapter we'll explore the performance of multiple channels in unstructured space operating against the same reference target. This work will first explore multiple channels aboard a single AFE17, then expand to multiple AFE17's operating conjointly – a configuration entitled BEI68.

17.2 The BEI68 Platform

BEI requires tight sampling synchrony among channels and signal normalization with respect to the excitation signal and its source. For this reason, the BEI68 was developed to house all of the necessary analog hardware and maintain a known relationship among all of its component parts. This makes system building easier. The front panel of the BEI68 imaging platform appears in figure 17.1(a). The BEI68 chassis is two standard rack units in height and can support up to four AFE17 receive cards as well as one transmit card. The back panel of the BEI68 (figure 17.1(b)) contains the fan exhaust port, the USB interface to the host computer (for configuration commands and control feedback), the power input port, additional USB peripheral ports for attaching motion control platforms, and the four shielded mixed signal connectors that interface with the data acquisition



(a) The front of the BEI68 imaging platform. The chassis is two standard rack units in height and can support up to four AFE17 receive cards and one transmit card.



(b) The back panel of the BEI68. (*from left*) The fan exhaust port, the USB interface to the host computer for configuration commands and control feedback, the power input port, (*above*) additional USB peripheral ports for attaching motion control platforms, (*below*) the four shielded mixed signal connectors that interface with the data acquisition cards in the host.

Figure 17.1 – *The BEI68 imaging platform – as built.*

cards in the host.

The BEI68 interfaces with a host PC for control, acquisition, processing, and display functions. Given that the output signals from the AFE17's (and, therefore, the BEI68) are all analog, the BEI68 must interface directly with a data acquisition unit for signal conversion prior to additional digital processing. The acquisition units chosen were PCIe-6323's from National Instruments. Bifurcating




Figure 17.2 – *The BEI68 interfaces with this data acquisition sub-assembly.*

the hardware between analog and digital signal processing units allows either to be upgraded in the future without incurring the cost and delay of reproducing the other. Two PCIe-6323's were attached to a bracket system and combined with power and signal routing to form the acquisition sub-assembly. The sub-assembly, as deployed in this dissertation, appears in figure 17.2.

The BEI68 and the acquisition sub-assembly are connected by cable. The interface cable contains 32 micro-coaxial copper fibers in addition to several unisolated digital signaling lines. The outer shell of the wiring harness provides additional shielding.

On the other side of the BEI68, the electrodes of the F68 attach to the analog processing cards (AFE17's) via a commodity 36-pin mini-SAS connector. The intermediate cabling is tandem shielded. This connection is visible in figure 17.3.



Figure 17.3 – *The electrodes attach to the analog processing cards (AFE17's) via a commodity 36-pin mini-SAS connector.*

The BEI68 system itself consists of a few key architectural elements. The receive and control portions of the BEI68 architecture appear diagrammatically in figure 17.4. The four AFE17's are grouped into pairs and routed to one of the two data acquisition cards in the host PC. The darkest gray elements are located on the F68 electrode array. The next lightest elements are located on the BEI68 motherboard. The lightest gray elements are located on the AFE17. The white elements are located in the host computer.

The complete platform is shown in figure 17.5. In the figure, the F68 electrode array is shown connected to a single AFE17. When fully wired, the F68 requires five separate connections: four (one each) to the individual AFE17 receivers and one to the TX2 transmitter card (not shown in the figure).

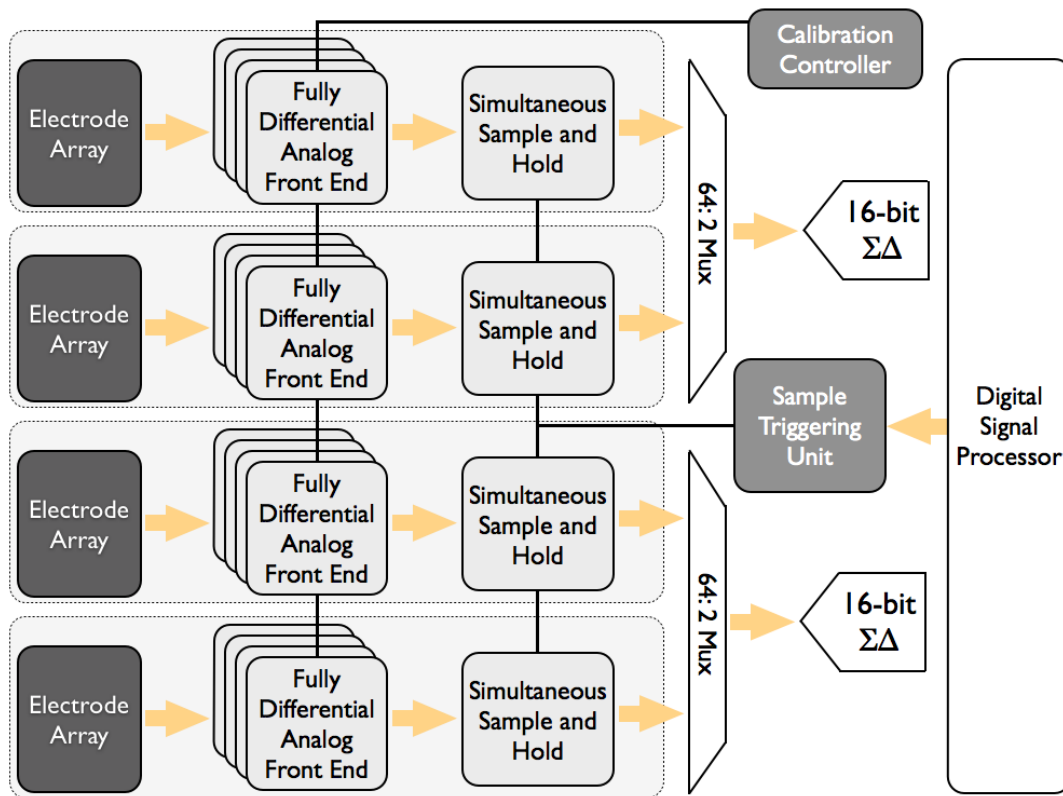


Figure 17.4 – *The receive and control portions of the BEI68 architecture. The four AFE17's are grouped into pairs and routed to one of the two data acquisition cards in the host PC. The darkest gray elements are located on the F68 electrode array. The next lightest elements are located on the BEI68 motherboard. The lightest gray elements are located on the AFE17. The white elements are located in the host computer.*




Figure 17.5 – *The F68 electrode array shown connected to a single AFE17. When fully wired, the F68 requires five separate connections: four (one each) to the individual AFE17 receivers and one to the TX2 transmitter card (not shown).*




Figure 17.6 – *The time-space variability of the target’s influence on a single channel of the BEI sensor array.*

17.3 BEI Detection and Tracking Performance

BEI works by establishing a known background field and then observing perturbations to that field caused by the presence of targets. Unfortunately, that background field contains noise. Elaborating on the experiments conducted in the previous chapter, where a reference metal pipe (the target) was raster scanned in a salt water tank by a motion-control platform, reveals the extent of the environmental variance. Figure 17.6 shows the recorded differential voltage, the output of one of the AFE17’s channels, as a function of distance from the sensor for various excitation output powers. Since the recording is performed sequentially, the spatial and temporal domains are correlated and the variation in baseline with distance also reflects a variation in baseline with time.

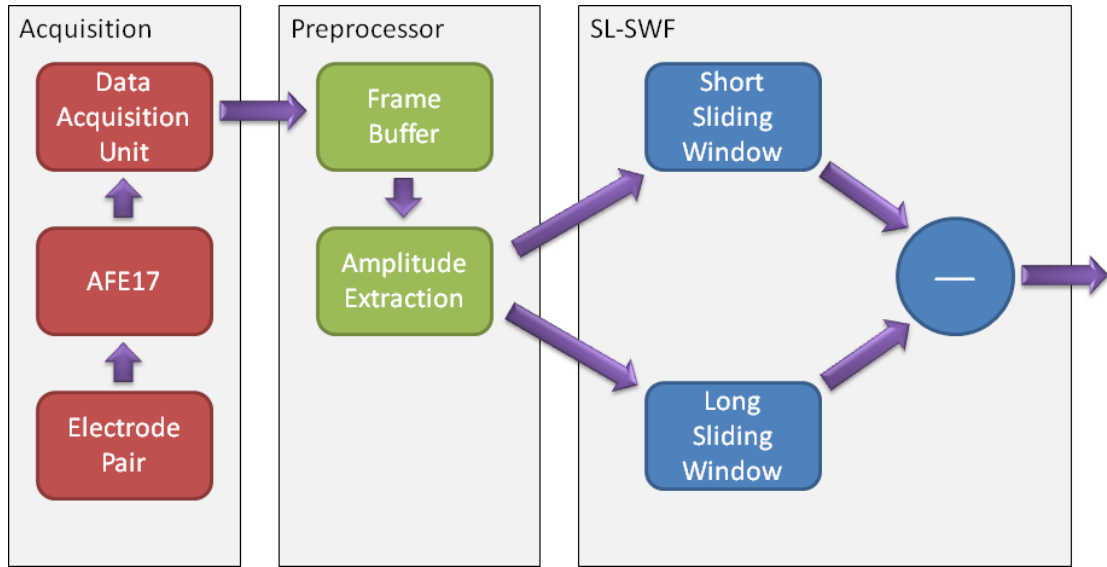


Figure 17.7 – *The console data visualization algorithm is a representation of the difference between a short and long running low-pass filter. This pathway is per channel and duplicated 16 times to cover the 17 electrodes of the AFE17.*

In the previous chapter, the discrete-spatial derivative with a window size of 1 point was used to remove these variances and normalize the background state. That worked exceptionally well *ex post facto*. Our objectives in this chapter extend to real-time and *in situ*. To that end, a slightly more elaborate filter scheme was devised (figure 17.7). Spatial filtering was reserved to a future work and a temporal approach applied in its place.

The Data Acquisition unit (DAQ) acquires readings from the AFE17 as a continuous stream. The stream is buffered and partitioned into 20 data frames. Each data frame consists of 100 samples captured at 15,625 samples-per-second. Each data frame is reduced to its peak-to-peak value by differencing its maximum and minimum¹. The resulting frame amplitudes (the output of the Preprocessor) are averaged into a single value representing the input to the discrete-time Short-Long Sliding Window Filter (SL-SWF).

¹In later experiments, not covered in this chapter, amplitude extraction was improved, first, by using a Fast Fourier Transform (FFT) to fit a curve to the samples and, later, by employing the sample-and-hold hardware of the BEI68.

The SL-SWF itself consists of two independent sliding window filters, one much shorter than the other. The output of the SL-SWF is merely the difference between the short window and the long window. This highlights recent changes in the dielectric condition of the environment from the background state, effectively creating a continuously self-calibrating sensor.

The obvious disadvantage is that the sensor becomes blind to stationary targets over time (much the same as the human eye). This is best handled by moving the electrodes (emulating the rapid movement of the human eye to counteract static blindness), tracking the state of detected targets, or both. It is important to note that this calibration process is not a viable substitute for bias calibration of the AFE17 as it cannot deal with analog dynamic range concerns.

The simplest SL-SWF configuration offers two tuning parameters, the lengths of the filters' sliding windows. In order to analytically optimize the tuning of the SL-SWF, let us model the input signal as a sequence of electrode observations made of a moving target. When the target is absent (outside the electrode's observation range) the electrode will record 1 – just the background. When the target is present, it will record 2. The output of either SWF becomes just:

$$\mu_{\text{swf}} = \frac{\min(\frac{20}{v}, 2W_s) + \max(W_s - \frac{10}{v}, 0)}{W_s} \quad (17.1)$$

where μ_{swf} is the output value of the SWF (the mean of the window's values), W_s is the size of the sliding window in number of samples, and v is the velocity of the target in sensors-per-decasample (number of electrodes it passes by per 10 samples).

We have expressed v in normalized units appropriate to the scenario. For clarity, consider the following dimensional analysis, in which knowledge of the sampling rate and layout of the electrode array readily convert v into the more recognizable unit of meters-per-second.

$$v = \frac{\text{electrodes}}{10 \text{ samples}} \times \frac{\text{samples}}{\text{second}} \times \frac{\text{spacing}}{\text{electrodes}} \quad (17.2)$$

The minimum and maximum functions in equation 17.1 are necessary to bound the SWF since a really slow target may exceed the window size. It is also of note that this model handles fractional samples to improve the visual quality of the resulting model output, an impossibility in our actual system.

The output of our tuning model, H_{filter} , is just:

$$H_{\text{filter}} = 20\log_{10}(\mu_{\text{short}} - \mu_{\text{long}}) \quad (17.3)$$

now parameterized by the length of each sliding window and the speed of the target, and expressed as a normalized attenuation in decibel form.




Figure 17.8 – *The Short-Long Sliding Window Filter evaluated over its design space in four dimensions.*

Visualizing this information requires a four dimensional space ($\langle\langle$ Short SWF length, Long SWF length, Target velocity, Filter output $\rangle\rangle$). In figure 17.8, the

solution space is plotted. The x and y axes indicate the length of the short and long filters respectively. The length of the filter is the size of its sliding window aperture. The z axis represents the speed of the target. The coloring is the resulting attenuation of the filter (in decibels). The transparent undefined spaces are impossibilities in which the short filter exceeds the length of the long filter (these conditions are identical to the reflected point of swapping which SWF is defined as “short” and “long”) or configurations that can produce negative filter output values. Operation at these design points should be avoided.

The space of figure 17.8 is sliced in a few representative locations revealing a number of trends. Filter attenuation is most severe where the filter lengths are equal given that the identical apertures result in zero difference between them. Increased length of the short filter may be compensated by increased length of the long filter for constant attenuation. Most interestingly, the long-window-v.-target-speed plane contains a parabolic region (w.r.t. velocity) that may be loosely described as an extruded C shape in three-dimensions. This is better visualized by the isosurfaces of figure 17.9.

That curvature results from the discontinuity in short-filter mean that results from its finite length. Consider the decelerating target. The progressive sample streams consist of $[0,1,0]$, $[0,1,1,0]$, $[0,1,1,1,0]$, et cetera, as the target enters the electrode pair’s observability and slows (more time is spent in front of the sensing electrode pair). With an aperture of 4 samples, the resulting means are: 0.25, 0.5, 0.75, 1, 1, 1, ... Taking a large difference in filter length as an example, at large target velocities attenuation is minimal as the short filter tracks the input signal near-perfectly. As the velocity decreases, the short filter is already saturated with signal and only the long filter mean improves (reducing the difference, ergo higher attenuation).

BEI deals with a quasi-static phenomenon in that it requires the electric field to be assumed static (only amplitude information). The AFE17 contains a very




Figure 17.9 – *Two isosurfaces from the design space of the Short-Long Sliding Window Filter.*

high-Q filter for this purpose [FHT12b]. In practice, changes in the environment and target location manifest as amplitude-modulation applied to the permitted narrow-band signal exiting the AFE17. This modulation is necessarily orders lower in frequency than the carrier as changes in ionization, environment, and target location occur much more slowly than the order-1 kHz field excitation.

Energy at the same modulation frequency leaves us without any filter options, save to appeal to historical context (tracking) or a change of vantage (apply more sensor channels). Accordingly, BEI performs best against targets that fit in between the low and high frequency bounds. High frequency noise originates in the sampling/acquisition process. Low frequency noise is predominantly calibration error from ionization [FTS10], environmental changes [FHT12b], changes in depth, or changes in output power [FHT12a].

Given that the SL-SWF output is the difference of two coherent low-pass filter functions, it is a bandpass filter. It has the added advantage of allowing independent control of the two cut-off frequencies through the choice of the short

filter length (high-frequency end of the pass band) and long filter length (low-frequency end of the pass band). Frequencies are not convenient when designing target tracking. It is more convenient to deal with target/noise velocities (rates of change). The filter length may be designed by asserting:

$$L_{\text{short}} = \frac{5}{v} \quad (17.4)$$

where L_{short} is the length of the SL-SWF's short sliding window aperture in samples and v is the velocity in sensors per decasample. The environmental filter may be designed from observing characterization data and setting T equal to the maximum number of contiguous samples of instability in your characterization dataset:

$$L_{\text{long}} = \frac{T}{2} \quad (17.5)$$

where L_{long} is the length of the SL-SWF's long sliding window aperture in samples.

Applying the filter length design equations to our datasets (of which figure 17.6 is representative), we observed setting $L_{\text{short}} = 2$ and $L_{\text{long}} = 50$ to result in excellent behavior in the majority of test cases.

17.4 Evaluation

17.4.1 Experimental Configuration

The experimental configuration (figure 17.10) follows the same basic layout as in the previous chapter. However, the F68 electrode array was used for this set of experiments. The F68 consists of four rows of 17 electrodes. Each row is connected to a single AFE17. The transmit electrodes are located on the midline between rows 2 and 3. For this study only a single AFE17 was involved. The transmit pair employed had its electrodes spaced 7.1 cm (2.8") apart about the vertical centerline of the array.




Figure 17.10 – *Annotated experiment configuration.*

The F68 is placed against a long wall of a 55 gallon (208 liter) capacity tank filled with 50 gallons (189 liters) of drinking water (0 ppt salinity). Sodium Chloride was added in measured amounts to raise the salinity to 31.5 ppt.

The reference target is a metal sponge coerced into a spherical shape roughly 3.8 cm (1.5”) in diameter. It is mounted to a wooden dowel that was soaked thoroughly in the tank’s salt water to minimize its dielectric signature. The reference target appears in figure 17.11 along with a United States Quarter-Dollar coin for size comparison.

At an output power of 94 mW, 100 data frames were recorded for each channel. Each data frame consisted of 100 samples captured at 15,625 samples-per-second. The amplitude of each frame was determined and averaged to create the calibration baseline. The AFE17 then begins operation by recording data frames simultaneously for each channel at a rate of between 15 and 25 frames-per-second varying based on the visual complexity of the resulting scene.



Figure 17.11 – *The reference target is a metal sponge coerced into a spherical shape roughly 3.8 cm (1.5”) in diameter. Shown with a U.S. Quarter Dollar.*

17.4.2 Tracking Moving Targets

The console displays 100 processed frames (the output of the SL-SWF) on its vertical axis and 16 channels on its horizontal axis. The console appears in figure 17.12. In the left pane no target is present and in the right a small metal sponge is approaching the array.

The coloring is relative and auto-ranging. When a target is absent all of the channels tend towards zero (the green color). Green is the middle value of the color scale. The resulting color variations are normally distributed about the zero value (that is why the mean is the middle color). They represent the subtle variations among the channels and reflect the residual acquisition system and environmental noise.

When a target is present, these minor variations are marginalized by the substantially stronger target signature. The result, evident in the right pane of figure 17.12, is the more uniform blue look of the background – blue being near




Figure 17.12 – *BEI68 console operator’s view without a target present (left) and with a target on direct approach (right). Color is auto-ranging.*

the minimum color.

The AFE17 detecting and tracking the target is shown in figure 17.13. The target is successfully detected and tracked as it moves back and forth across the array at a minimum distance of 2.5 cm (1 inch)². The slant of the tracks in the figure is position with respect to time. It, therefore, indicates the direction in which the target is moving. Positive slopes represent the target moving left-to-right. Negative slopes represent the target moving right-to-left.

Approaches orthogonal to the array and at intermediate angles are also easily distinguished – the limit of which being the orthogonal approach of figure 17.14. In the figure the target approaches and retreats along the same bearing and moves along a track orthogonal to the array and directly in front of electrode pair number 7 (electrode 8’s signal subtracted from that of electrode 7’s). The corresponding detection and track of the target by the BEI array is readily apparent.

The BEI68 can distinguish targets that come to rest in its field-of-view. An alternative description is that the target cannot evade detection by rapidly approaching the sensor and then coming to an immediate dead stop – intending to create a quick “blip” on the console that might go unnoticed by the operator.

²A video demonstration of this work is available: <http://vimeo.com/45524639>

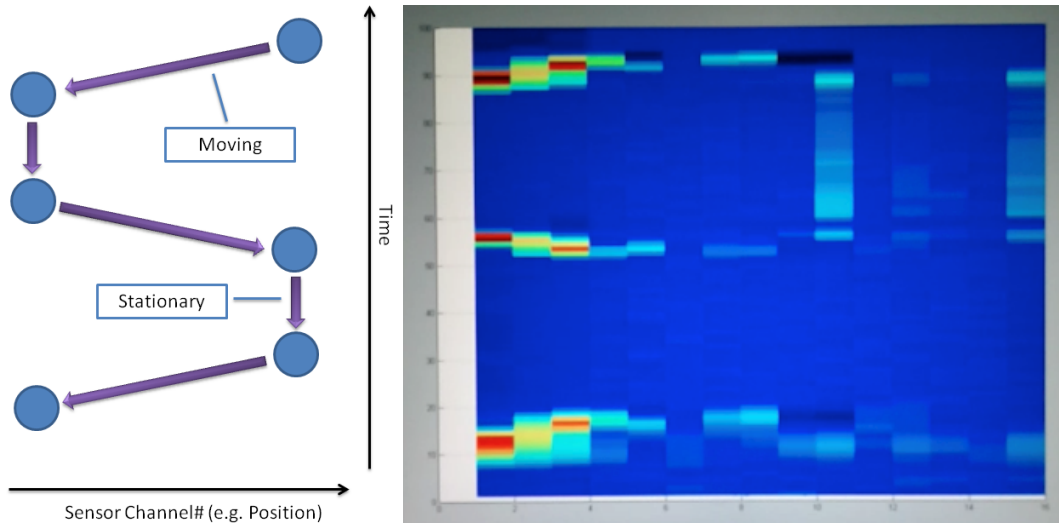


Figure 17.13 – BEI68 console operator's waterfall view of a conductive target oscillating back and forth in front of the array. The slant of the tracks indicates the direction.




Figure 17.14 – BEI68 console operator's view of a target making a direct-on slow approach and then retreating.

After the initial detection, the decay rate of the signal for a target *still present*, but immobile cannot be faster than a target that exits the sensor range due to the sliding window architecture. A target still present will continue to contribute a raw disturbance to the background field. This variation will necessarily be less of a difference from the previous (in time) recorded value (where the target was also

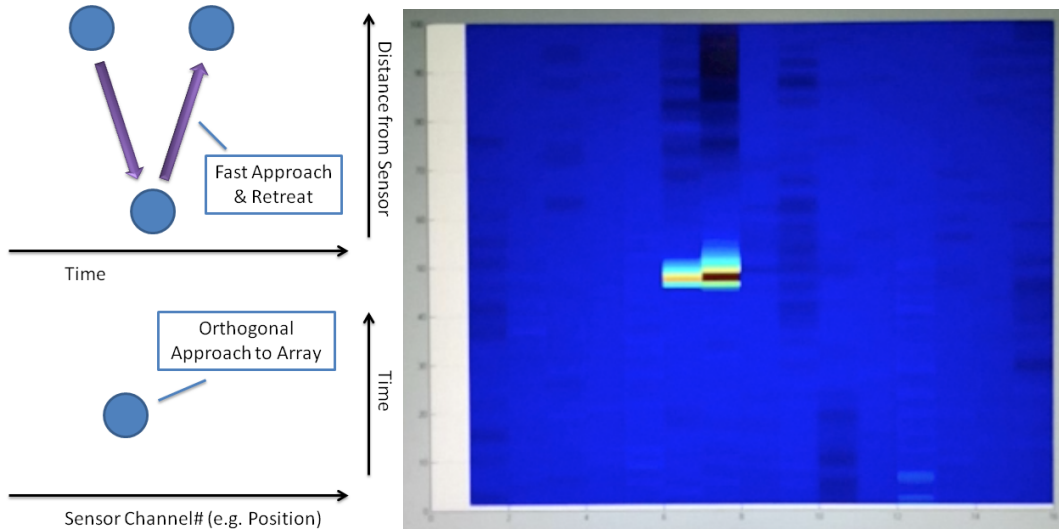


Figure 17.15 – *BEI68 console operator's view of a target making a direct-on approach, this time faster than in the previous figure, and then retreating.*

present) than the target's departure.

The SL-SWF, required to isolate background noise, will eventually render the sensor blind to the static target, but this will occur at a rate approaching the amplitude change from the target's initial detection (ΔA) divided by the length of the longer sliding window W_L when the longer window greatly exceeds the shorter one ($W_S \ll W_L$). A stopped target will produce a signature that decays linearly in time at a rate $\approx \Delta A/W_L$. In the case of figure 17.14, the signal disappears within 5 frames of the maximum value – rather than the near-50 we would expect if the target had merely stopped after its direct approach instead of retreating.

Figure 17.15 is a recreation of the scenario of figure 17.14 only, this time, the target approaches at an even greater velocity. A comparison between the two figures shows the qualitative effect that velocity plays on the visualization, while the time scale bar allows the target's velocity to be computed quantitatively. Analytically, the dynamic performance is limited by equation (17.6).

$$v = \frac{(n_e - 1)d_{sep}n_{fps}}{n_f} \quad (17.6)$$

In (17.6), n_e is the number of electrodes, d_{sep} is the separation distance between them, n_{fps} is the number of sample frames per second (the number of samples per second for each channel sampled in parallel – or number of sample stripes for sequential sampling), and n_f is the number of frames between a change in the detected position of the target.

For the AFE17, F68, and BEI console in their present state, equation (17.6) evaluates to ≈ 8 m/s or 18 mph – the biggest determinant being the frame rate. The frame rate is a pure engineering challenge, not a physics-limiting one. As a result, the performance may be increased going forward through the use of faster sampling hardware, more sampling hardware operating in parallel, and/or through the use of fewer simultaneously operating channels. Alternately the electrodes could be spaced further apart at the possible expense of creating holes in the sensor coverage.

The absolute limit is determined by the excitation frequency and is fixed to

$$v = \frac{(n_e - 1)d_{sep}}{f_{TX}} \quad (17.7)$$

where f_{TX} represents the excitation frequency.

17.5 Conclusion

The utility of BEI and the efficacy of the BEI68 platform should now be apparent. This chapter has covered necessary algorithms and their tuning to detect and track a small reference metal-mesh target in a salt-water tank. The construction and performance of the first practically capable BEI system has been reported.

Part VI

Conclusion

CHAPTER 18

Future Applications

18.1 Overview

BEI is well suited to redress two classes of problems: large-scale and micro-scale. Large-scale problems are those of maritime interest. Foremost among these are manned submarine collision avoidance and Unmanned Underwater Vehicle (UUV) navigation (collision avoidance, docking, and formation support). Other large-scale applications include harbor intruder detection, hull inspection of surface vessels, submerged pipeline inspection, ocean floor mapping, and anything else where being able to “see” short-distances underwater is useful. Operating at-depth in the ocean, there is no alternative sensor system which can match the temporal (latency) and spatial resolution of BEI.

BEI may also target micro-scale problems. BEI’s resolution is a ratiometric phenomenon determined by the dimensions and configuration of the electrode array that captures the signals. Dense high-resolution electrode arrays are already the norm in neuroscience laboratories. By replacing their simple passive backend electronics with BEI, it is my sincere hope to develop a machine capable of perfect brain tumor resection (surgical removal).

18.2 Real-Time In-Vivo Guidance of Surgical Tumor Resection

Brains consist of two types of cells: neurons, which perform the actual functions of the brain, and glial cells, which provide support functions (repair, defense, etc). Because neurons are so highly specialized, during the course of maturity they quickly lose the ability to reproduce. As a result they cannot become cancerous. The glial cells can and do. As a result, many brain tumors are considered inoperable as the cancerous glia become intertwined with healthy functional brain tissue.

At present, removing the tumor is done in bulk (the surgeon just tries to define an outer boundary and scoops out everything inside it). If the removal of the associated healthy tissue would remove brain functions (such as speech, vision, regulation of heart-rate, etc) where the quality of life would be much worse than living with the tumor, the tumor is left in place and the patient dies within a few-months to a year.

The surgical isolation and resection of tumor cells is confounded by the shared origin of abhorrent and normal tissue. The cells appear similar under visual inspection and tumor boundaries occur on a scale beneath the limits of human hand manipulation. As a result, modern surgical practice involves the use of resonant and functional imaging to identify the boundaries pre-operatively and then register this prior visualization against the patient intra-operatively (a few recently constructed facilities allow the patient to be moved back and forth via a co-located MRI system). In either case, surgical conservatism results in additional tissue loss than might otherwise be necessary.

What is needed is high-resolution spatio-temporal guidance for the removal of tumor tissue that can operate in-vivo, in real time, without precluding simultaneous surgical action.

The Golshani Laboratory for Neuroscience at the University of California,

Los Angeles, houses one of the world's few liquid-cooled two-photon microscopes capable of in-vivo 3D surgical resection. The instrument uses two laser beams that it directs to intersect at desired coordinates. The energy of a single beam is low enough to safely pass through intermediate tissue, but the point of intersection contains just enough energy to destroy the targeted cell. This provides on-demand cellular-scale attack capability.

The more difficult problem lies in the targeting of the cancerous cells, as distinct from the immediately surrounding healthy tissue. A study [SPL00] examined three different Dunning cell lines of various metastatic abilities, chosen because these cell lines are characteristically representative of a broad class of tumorous human cells. It concluded that tumor cells could be reliably detected with 95% accuracy and only a 12% false positive rate by their conductivity alone.

BEI can image conductivity in real-time. Therefore, a modified high-density neuroprobe (sheet of electrodes) placed non-destructively at the surface of the brain tissue (figure 18.1) and connected to my BEI platform could potentially provide the requisite targeting capability. The temporal resolution required has already been achieved in my prototype. The spatial resolution required for a cell-by-cell inspection requires that the BEI imager achieve sensitivities less than 1/10th the receive axis length – or about three times as sensitive as my current detector.

The goal of the next five years of my research is to create a combined pulsed multi-photon laser microscope and biomimetic electrostatic imager. The laser acts as the scalpel selectively destroying cells targeted by the BEI. The integration is self-registering as the laser and BEI electrode array can be mounted and calibrated co-axially with the laser beams entering the brain matter through the requisite spaces between adjacent electrodes. This would take the operative scale of brain tumor resection from the millimeters, at which human hands currently work, to the microns at which the laser and BEI can localize. Computer controlled,

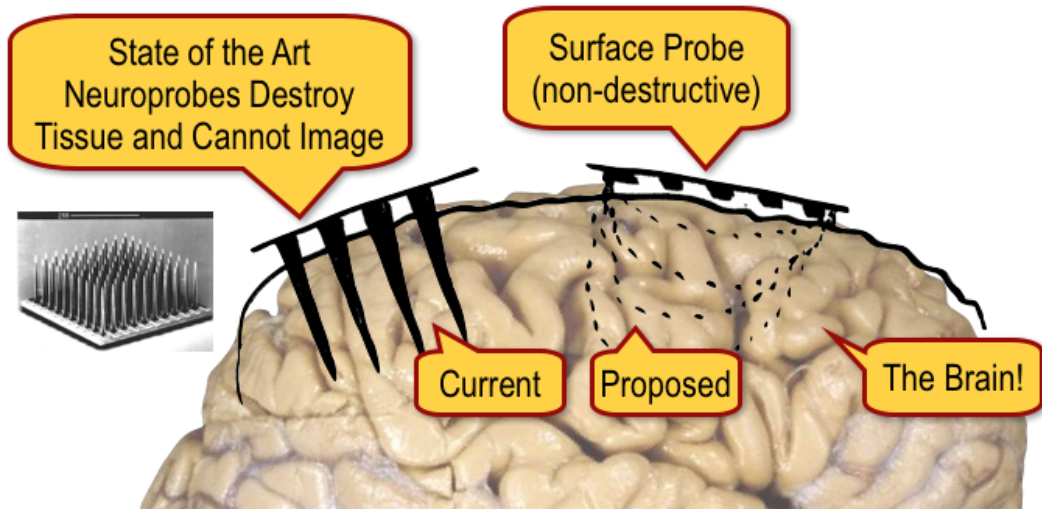


Figure 18.1 – *Traditional neuroprobes require electrodes to be placed in direct contact with cells. BEI proffers an alternative less-invasive approach while also providing anatomical imaging. Obviously not to scale.*

BEI-directed, laser surgery would end the notion of an inoperable brain tumor, reduce the invasiveness of the surgery, reduce the surgical skill required (reduces cost, improves availability), and substantially reduce epileptic post-operative complications.

18.3 Submarine Collision Avoidance and Navigation

At the other end of the scale-wise spectrum you find in-ocean applications which rely on the navigation and communication potential of BEI. The need for such technology is clear. As of 2001, the United States Navy (USN) submarine surfacing procedure calls for its submarines to maintain position underwater for at least 180 seconds while actively pinging its sonar in order to establish the presence of surface vessels posing a collision hazard as the submarine comes to the surface [Nav01]. In at least one instance, procedure was ignored to disastrous effect when the USN Greenville struck a Japanese civilian fisherman training ship with high school students aboard while surfacing [NT01] (figure 18.2). Had an electrostatic imaging



Figure 18.2 – *Damage to the Los Angeles class US Navy SSN Greeneville post collision with the Japanese civilian fishing boat Ehimemaru. Source: US Navy*

system been developed and deployed aboard, the crew would have had access to instantaneous real-time video imagery of the electrical disturbances caused by the presence of the surface vessel and the collision might have been avoided.

The problem is pervasive. In the past ten years (2001-2011) there has been an average of one major collision event per year involving military submarines at an average repair cost of \$78 million USD per incident¹[Tin09]. Biomimetic electrostatic imaging shows tremendous promise in submerged salt-water navigation and collision avoidance since it is built on an effectively instantaneous phenomenon and does not require an active target.

¹Repair estimates were not available for all incidents.

CHAPTER 19

Synopsis

19.1 Overview

Broadly, the contributions of my work include:

1. The application of theory to explain the observed performance of submerged oceanic electrostatics,
2. The development of a model to predict future design improvements,
3. Experimental proof of the existence and utility of the phenomenon,
4. An engineering validation of the rationale for the naturally observed weak-electric fish waveforms, and
5. The design and implementation of the first *Biomimetic Electrostatic Imaging* (BEI) systems – an engineered sensor which mirrors the biological phenomenon of electroreception.

This dissertation includes validations of the 3-electrode, 16-electrode, 17-electrode, and, ultimately, 68-electrode versions of the BEI hardware. It has detailed ocean and laboratory validation of BEI and the AFE17. And, it has revealed the efficacy of a number of design improvements made to BEI hardware and software. The improved BEI system increased detection range by 97% and range-power efficiency by as much as 150% without a corresponding increase in the false positive rate. Both the differential and differential disturbance fields were

shown to exist at greater than $7\times$ their respective electrode separation distances – the furthest reported in the literature to date.

As a result, this work has conveyed a number of design axioms which will drive future experimentation.

19.2 Contributions to Submerged Electromagnetics

Our early investigations into oceanic electrostatics met with limited success. A combination of experimental and conceptual shortcomings resulted in mixed-mode propagation wherein some signal energy was able to travel through the air above the water surface and couple into our instrumentation.

In our early laboratory testing we noted that no communication was apparent in free air. The SNR was around the receiver sensitivity of 6.5 dB, whereas in salt water, the average SNR was substantially higher (>12 dB). In fresh water it was 11.5 dB. Of note, the signal to noise ratio was decreasing with increasing distance in fresh water, whereas, in salt water it remained almost constant.

The standard model for electromagnetic RF propagation in salt water predicts a 10x reduction in transmission range for constant power with respect to fresh water [Rho07]. This was clearly not observed. In contrast, the electrostatic (transconduction) model predicts relatively little change in BER over distance given that the link remains open, which agreed with our results.

Our ocean expeditions continued the disturbing trend of better-than-predicted performance, yet non-compliance to either linear radiator electromagnetic or electrostatic theories. In the ocean, we successfully transmitted digital data packets over distances of up to 5 meters using GMSK at 8MHz and 250kbps data-rate on only 23mW of transmission power with the electrodes submerged. The packet length was 100 bytes per packet. Reducing the packet length helped re-open the link. Decreasing the packet length, given constant BER, increased the

number of successful packet transmissions on average. At only 10 bytes per packet, we demonstrated successful packet transmissions out to 9 meters.

While the electrodes were completely submerged in these experiments, feed lines were used to return the signals to on-shore acquisition and processing equipment. In order to completely isolate and submerge the transmitting and receiving systems, we designed and implemented a waterproof computer system which encapsulated the prior experimental setup hardware (USRP, TXLF, RXLF, Battery, Hard disk, et. al.) replacing the computer with a single-board variant based around the Intel Atom dual-core 1.6GHz 330 model CPU.

In these fully submerged and isolated experiments, the results were not reproducible. Given the similarity of my early laboratory and expeditionary findings to those reported by [ASS04] and others, I was convinced of the need for such an explanatory model.

19.2.1 The Folded Unipole Antenna

I found that model in the form of the Folded Unipole Antenna (FUA). Broadly, antennas fit into two groups: linear and aperture. The FUA is in many ways the superset of all linear antennas. The FUA has been well enough understood to be practically employed for over 50 years, yet, has not been studied or optimized sufficiently to enjoy widespread use [Rai07]. It is the integration of linear radiators with transmission lines and can take on near-limitless forms.

Electrically short monopole antennas have low resistance and high reactance. These antennas, therefore, require a capacitive matching network to tune out the natural reactance present in the short structure, but this matching network results in narrow bandwidth, relatively small aperture, and high losses. The cause of these losses is primarily the relatively high currents flowing at the feed-point.

The FUA reduces these losses because it reduces these currents. FUA's consist

of no less than six independent design variables which may be practically simplified to these five:

1. Frequency – easily determined as it is under experimental control.
2. Common-mode impedance of the folded wire pair (Z_C) – determined from the cross-sectional geometry or by measurement.
3. Differential-mode impedance of the folded wire pair (Z_D) – also determined from the cross-sectional geometry or by measurement.
4. Top terminating impedance (Z_{top}) – this may be calculated or measured for the device explicitly added to the structure to serve in the role of top termination. As it is provided by design, it usually has a target value. It has been demonstrated in this work that, in the case of our experiments, the ocean in concert with the artificial ampullae served in this role.
5. Base terminating impedance (Z_{base}) – similar to the top termination, the base termination is typically a specific designed-in value achieved through the addition of discrete elements such as a capacitors or by structures such as shunt transmission lines or environmental factors effectively serving in that capacity.

19.2.2 Ocean-Terminated Folded Unipoles

This dissertation is the first work to demonstrate the applicability of the Raines' FUA model to in-ocean submerged electrostatics. Working clockwise around the schematic of figure 5.1, we described how the FUA model mapped to our experimental configuration. It should be noted that the experimental configuration is rotated 180° from the figure, its feed point being generally considered to enter from above, given that the other end of the cable reaches the electrodes which may be submerged in the ocean. It may be simpler to perform the rotation the

other way and imagine the ocean as being at the top of figure 5.1. In either case, the signal generator, being connected to the driven arm of the electrodes' feedline, forms the first leg of the FUA. The impedance of the electrodes themselves provides the top termination and the undriven arm of the feedline forms the return leg of the unipole. No base termination is applied ($Z_{base} = 0 + 0j$, it is a perfect short circuit), consistent with the feedline's direct ground connection to the frame of the signal generator.

Given that the geometry of the feedline readily yields values for Z_c and Z_d , Z_{base} is shorted, and a typical current transformation ratio value of $1/2$ is reasonable, the only missing value necessary for a Raines' analysis of the FUA model is the top termination. Consequently, the impedance of the electrodes must be determined.

In this dissertation I have shown that the electrodes, of my earliest designs, were well approximated by a large flat structure with little inductance to speak of. In short, they may be well modeled as a parallel plate capacitor whose capacitance, C_{AA} , may be estimated by:

$$C_{AA} = \frac{\epsilon_r \epsilon_0 A}{d} \quad (19.1)$$

where $\epsilon_0 \approx 8.854 \times 10^{-12} \text{ F/m}$, ϵ_r is the relative permittivity of the medium, A is the effective overlap among the plates, and d is the separation between them.

When the electrodes are not submerged, they offer little capacitance and almost infinite resistance (a near-perfect open-circuit). When the electrodes are submerged in the ocean, the entire situation changes – *even if they are barely submerged*. The coaxial feedline connecting the electrodes to the transceiver is now top-terminated by the ocean and electrodes, acting collectively, as a capacitive reactance.

The results of applying the Raines approach for folded unipole antenna analysis *with the electrodes submerged* in the ocean, appeared in figure 5.3. In that analysis, the zero points in the reactance shifted to much smaller fractions of the wavelength, repeated with greater frequency, and occurred simultaneously with usable values

of resistance. The resonances no longer occurred at near-zero or infinite resistance. Not only will the retuned structure radiate very efficiently at several points, but the probability of operating at, or near, one of those points increased almost 66% over the air-terminated monopole.

Practically, this technology shows promise for buoys, periodically surfacing crafts, and other devices operating at or near the ocean-surface. An experiment (table 5.2) was conducted to demonstrate the relative advantage of the ocean-terminated FUA under various scenarios. All manner of combinations were explored in which an ocean-terminated FUA could communicate with a second ocean-terminated FUA. Each was either submerged or surfaced and either co-located or isolated from the other.

The resulting propagation times fell into three categories: (1) The fastest group, which propagated through the air (Points 1-3), (2) a middle group, which propagated through both the air and the water, and (3) a slowest group, which propagated almost exclusively through the water.

The amplitude and propagation time were correlated, which was encouraging. There was a clear performance benefit for submerged electrodes. The strongest recorded signal occurred when both electrodes were submerged in the same tank (point 12). If the distance between the electrodes was held constant, but the electrodes were moved into two separate tanks, the signal received decreases dramatically – $\approx 59\%$ between points 11 and 12.

Three general trends with respect to signal strength were clear: (1) submerged performance always exceeds dry performance, (2) more distance results in more attenuation, and (3) electrostatic propagation results in less attenuation than electromagnetic propagation over short tank distances.

19.2.3 Summary

In this dissertation I have shown that surface-wave propagation and ocean-terminated folded-unipole antennas have utility in ocean network applications where surface-breaching antennas may be applied. I have also demonstrated that certain submerged communication performance claims based on experimental setups similar to those of chapter 4 are inappropriately applied to the fully submerged case. I have offered the first application of the folded unipole antenna, and the Raines' approximation in particular, to explaining the observed performance of my experiments and those of others ([ASS04]). The model is of great utility as it may be used to not just explain observed performance, but to design and predict the performance of various alternatives.

19.3 Contributions to Far-field Analysis

It has been my aim to contribute, not to just the specific application of biomimetic electrostatics, but to the larger practice of sensor design. To this end, I have proposed the Twisted Quadratic Distribution (TQD) as a means to speed the three-dimensional approximation of radiation power density problems.

19.3.1 The Twisted Quadratic Distribution

The TQD layout is visualized in figure 10.4 where it is obvious that the TQD produces a similar concentric circle layout, however the TQD is twisted completely once around when tracing from top to bottom on the sphere. Comparison of the two plots (which show the azimuthal distribution) reveals that despite the TQD containing 33 fewer points in total it looks denser than the reference distribution. Numerically, its distribution is 320% more uniform by our D metric definition.

Array composition analysis is performed by computing the impact of each

virtual interferer on each element of the array. Therefore, a distribution which produces greater uniformity can provide the same analytical capability with fewer points in the cloud. This relationship is roughly linear allowing our TQD based analysis a better than $3\times$ runtime improvement without penalty to the fidelity of the results. In concert with the Far-field Arbitrary Array Analysis Simulation Toolkit (FAAAST), developed as part of this work, which employs it, I was able to demonstrate this expected simulation performance improvement ($1/3$ the original runtime).

19.3.2 Array Steering and Optimization

Beamforming (or beam steering) is the process of constructing interference patterns among array elements to yield a composite far-field pattern with highly directional sensitivity. It is a requisite task in array signal processing applications such as radar, sonar, acoustics, astronomy, seismology, communications, and medical-imaging. To achieve imaging of a scene, the beam must be steered, but the creation of far-field focus via interference is limited to just that – the far-field. Near-field situations, such as biomimetic electrostatics, are generally considered to be unable to benefit from array synthesis.

Radiative far-field synthesis works because the electromagnetic energy travels as a wave. The energy is ejected from the antenna and travels through space. As such, the point of origin of any forces exerted on charges in the environment is the current location of the energy quantum radiating through space. The origin moves because the wavefront propagates. If two coherent waves are emitted in parallel with opposite phase and traveling in the same direction, an object equidistant between them and collinear with the wavefront will experience zero net forces. However, if the forces were considered originating from some distance away (as opposed to collinear with the object), the net force would have a component in the direction between the wave origin and the object's location. This is the case

for electrostatic charges.

This dissertation has revealed that, even given this reality, all is not lost. While unable to truly cancel field energy, the number, location, polarity, and intensity of charges has an impact on the final electric force field. When evaluated differentially, as in the case of biomimetic electrostatic imaging, a surprising degree of verisimilitude with electromagnetic synthesis capabilities is achievable.

In figure 10.7, the addition of four additional charges, from the configuration of the figure which precedes it results in an enlarged primary lobe and reduced (relative to the primary) secondary lobes. The illustrated configuration is an example of the use of an array transmitter to increase effective range and minimize side lobes. The array transmitter provides a means to overcome the absolute amplitude limits imposed by the chemistry of the marine system as discussed in section 6.2.1 as the enlarged pattern results without increased output amplitude at any one electrode.

Array focusing of electric fields in the near-field via phasor interference is not possible. However, a more limited capability to focus the field may be achieved through field complication (adding more charges to the field). In figure 10.8, the six charges from its preceding figure were reweighted to a new relative polarity/amplitude.

The results of this work are have been quite encouraging. I have shown that in at least one specific instance, the side lobes have been reduced by approximately 50%, while the primary lobe has been reduced in amplitude by only 14.4%. As expected, unfortunately, little far-field focusing has been achieved with the reference beamwidth of 61.2 degrees declining little to 61.0 degrees when array synthesis was performed.

The effect of this particular array weighting is covered in figure 10.9 the data being the difference between the reference and focused beam patterns. The side

lobes have been strongly attenuated while the close end of the hyperbolic null region has been filled in by additional field strength. Despite limited ability to reduce beamwidth, the array still provided benefit. Comparing the three figures (figures 9.8(a), 10.8, and 10.9), it was clear that the array weighting vector played an extremely important role. Discovery of the optimal weighting vector is best handled by application of the FAAAST tool mentioned as a previous contribution.

19.4 Contributions to RADAR

Inspired by the Brown Ghost Knifefish's ability to foveate, interfere with, detect, and track quasi-static electric signals, in this work, I have proposed a new signal encoding technique that makes use of Angle-of-Arrival (AoA) information (direction/foveation) in conjunction with local interferometry (jamming response) to improve target location estimation. I called this process Angle-of-arrival-assisted Relative Interferometry (ARI). In the context of RADAR systems design, it has a number of desirable attributes including the ability to reduce the synchronization, network, and hardware requirements when operating as the sole RADAR modality and its ability to augment existing pulsed and CW RADAR algorithms. In oceanic electrostatics it proffers a path to higher resolution imagery and target decluttering.

Over the various chapters of this work, the Bio-inspired waveform encoding technique of ARI was explored in detail. Chapter 12 covered the simplest form of ARI based on a continuous-wave transmission. Having demonstrated the approach and its utility it was incumbent upon us to demonstrate its application. To this end, I demonstrated a Software Defined Radio (SDR) approach to an Angle-of-arrival-assisted Relative Interferometric (ARI) RADAR transceiver. A software-defined model in Matlab was developed and implemented. The model was used to explore the implementation parameters and demonstrate the stability of the scheme in the face of clock jitter. At the end of the chapter a SDR-ARI transceiver was designed,

simulated, and implemented utilizing USRP hardware and gnuRadio software.

In chapter 13, I proposed a Pulse-Wave (PW) extension to the ARI technique to dramatically reduce the scan-time and the number of vantage points necessary to obtain high-fidelity target position estimation. Accordingly, I called this enhanced process PW-ARI. PW-ARI is the fusion of data from three domains: time (time-of-flight), phase (relative phase-of-arrival), and angle (direction-of-arrival). It has a number of desirable attributes. Foremost among these are its ability to rapidly image extremely large volumes with high accuracy and dense target clusters, support an infinite number of friendly submarine or surface vessels while providing stealth operational support, and present a spectrally small footprint to both hostile (evade detection and counter-measures) and friendly (maximize electromagnetic compatibility) systems.

19.5 Contributions to Submerged Electrostatic Fundamentals

This dissertation presents the highest resolution study ever reported involving the generation of a quasi-static electric field, recorded differentially and at-depth, in the ocean. As reported, an expedition was undertaken to the Pacific Ocean to a site in the coastal city of Marina Del Rey, California. A six and half hour study was conducted by a cantilevered purpose-built robotic gantry system, which collected over 10,000 data points at 635 micron-square resolution. The differential field was shown to exist at greater than $7\times$ the electrode separation distance – the furthest reported in the literature at this time. A mathematical model was derived from first principals, presented, and used to corroborate the findings.

A recurring point of confusion with this work, among readers at first encounter, is its differentiation from radio – the propagation of electromagnetic (EM) waves. Figure 19.1 illustrates this difference. EM is a dynamic phenomenon. Radio waves




Figure 19.1 – *Electromagnetics is a dynamic phenomenon in which quanta of energy travel to the target. Electrostatics is a static phenomenon in which the quanta interact from their origins.*

result from attempting to balance the conservation of energy equation [HRW97]. When mobile charges reach the end of an antenna, they must stop (decelerate). The momentum they’ve acquired must transfer somewhere. That energy is radiated outward as an electromagnetic wave. The wavefront contains quanta of energy know as photons.

In the left pane of figure 19.1, two photons are travelling past a charged “sensor” particle. Because the wavefront propagates at the speed of light, the electromagnetic forces cannot effect the sensor until the photons arrive (the forces being orthogonal to the direction of travel). As shown in the figure, at this time the two equivalently charged photons exert equal force on the sensor particle from opposite directions resulting in zero net force.

In the right pane, the energy quanta (the former photons) are just stationary (static) charged particles. The force field easily exceeds the velocity of the quanta and results in a force vector in the direction of the “sensor” particle. Given the symmetry, the lateral components cancel, but the vertical components superpose

to result in a net force away from the electrostatic charges' axis. These forces degrade with distance because the force originates from the static charges.

In the EM case, the forces originate from the propagating photon. So, *ceteris paribus*, the force it exerts on the “sensor” particle does not degrade with the distance between the EM source and particle as the origin of those forces is always the wavefront when it reaches the particle.

For intuition building purposes, one might consider electrostatic coupling as the low frequency end of radio communication since the majority of energy coupling to any receiver within a wavelength of a radio transmitter is overwhelmingly electrostatic or magnetostatic (the near-field or inductive field region) [Wes91] [Mil93] [RWD94]. In air, where this discussion is usually applied, the main advantages of radio wave energy propagation, as compared to electrostatic coupling, are higher bandwidth and longer range.

Another facet of this longer range is energy efficiency. Given that range can typically be improved with increased output power, of prime interest is the ratio between the two rather than absolute numbers. In air and at range, EM transmission is a clear winner. At shorter ranges (within a wavelength or two), the energy advantage dissipates given the high carrier frequencies necessary to keep short range receivers more than several wavelengths away.

At 2 meters, the operating distance of the carrier frequency of a radio transmitter would need to be in excess of 150 MHz to operate at far-field in air – *regardless of the amount of data carried across the channel*. This comes with a large energy penalty as energy and frequency are linearly correlated to the first order. Compared with a base-band electrostatic transmission, which can operate at the frequency of the data payload, a huge inefficiency occurs. For a 10 kbit/second data stream, electrostatics achieves an advantage of more than 15,000× assuming the receiver is within range.

19.6 Contributions to Electrostatic Imaging

Biomimetic Electrostatic Imaging (BEI) is a technique modelled from nocturnal ocean animals unable to employ light-dependent forms of perception to navigate their surroundings, hunt, or avoid predators. For these animals the primary means of perception involves the passive detection of electrical fields (electroreception) necessarily created by muscle activity in their prey (a process similar to the human electrocardiogram). In some species, the process is an active one in which the generation of a Voltage gradient in the intermediate environment enables the discernment of non-emissive navigational hazards (such as rocks) by the manner in which they disturb the self-generated field.

It has been the core work of this dissertation to journal the efforts of myself, and those with whom I have collaborated, in creating and validating an engineered sensor which mirrors this biological system's ability to detect targets with different conductivities from the background ocean environment.

In summary, this work may be broadly classified as an electrostatic system with the following characteristics:

1. Submerged in-ocean operation – the goal of the system is in-ocean and submerged operation, mimetic of the lifeforms that inspired it
2. A surface array of electrodes – the targets are forward of, and not surrounded by, the sensor as is necessary for tomographic applications of electrostatics.
3. Conjoined transmit and receive (in the case of imaging) – transmitter and receiver may share a ground and timing reference of extreme quality.
4. High-impedance receive electrodes – high sensitivity measurements from electrodes in direct contact with the seawater. Input and output current (leakage) is minimized.

5. Low-impedance transmit electrodes – direct coupling of a voltage source via the electrodes to the seawater.
6. Difference-linear detection – determination and detection of the *differential disturbance* field (in the case of imaging) by voltage differencing spatially adjacent electrodes in the array. Direct differencing allows for high-sensitivity measurement by applying the Common-Mode Rejection Ratio (CMRR) against environmental noise.
7. Single-excitation – detection is simultaneous and does not require a sequence of *excitations* to produce a measurement as is the case in Magnetic Resonance Imaging (MRI), X-ray Computed Tomography (X-CT), or the various forms of Electrical Impedance Tomography (EIT).
8. Dielectric backing – mirroring the fish’s body, BEI makes use of a dielectric (e.g. nonconductive) backing to the electrode array to focus the field-strength forward of the array.

BEI stands apart from other common electrostatic systems such as Electrical Impedance Tomography (EIT). In EIT, the data is determined by exciting pairs of electrodes and observing either the total impedance of the path, the resulting voltages around the target (by surrounding it with electrodes), or both. In BEI, no such sequencing of excitation is required. The target determinations are made directly from a single excitation of the field. This makes BEI substantially faster and better at imaging targets that are outside the transmit axis. Naturally, BEI is inferior in volume resolution for targets intermediate the transmit electrodes (e.g. near the transmit axis) and potentially shorter in range if the EIT electrode ring is sufficiently large in diameter.

BEI’s electrostatic approach offers a number of additional features of interest. The sensor is theoretically impossible to shield against since coating the target in

a material equal in dielectric coefficient to the ocean (in an attempt to make it invisible) would simply cause the excitation field to penetrate that material as well. As a static phenomenon there is no wave propagation to phase cancel and any active jamming emission would be detectable.

In terms of Uncrewed Underwater Vehicles (UUV's), BEI is extremely attractive as a non-contact "bumper" given that, unlike acoustics, its transducers are extremely cheap (effectively, just the exposed end of a wire or screw). As such, they are also tiny, can be made conformal, and offer an extremely wide aperture (e.g. they are natively unfocused and cover the full hemi-spherical field of view).

Returning to the biological origin, from which this work emerged, comes a clearer statement of the class of problems I intend to target. Sensing and communication at long ranges is best handled through acoustics and chemotaxis. At shorter ranges, in fresh and brackish waters, needs are best dealt with via optical means. In the ocean, at the close ranges where an electric field may be extended, electroreception is the clear choice. It has been the work of this dissertation to define the optimal design and operating conditions for such an engineered biomimetic system and to demonstrate when and how it could be used to greatest effect.

REFERENCES

- [Aqu09] AquaComm. “Underwater Wireless Modem.” 2009.
- [ASS04] A.I. Al-Shamma’a, A. Shaw, and S. Saman. “Propagation of electromagnetic waves at MHz frequencies through seawater.” *Antennas and Propagation, IEEE Transactions on*, **52**(11):2843–2849, Nov. 2004.
- [Bal05] C. Balanis. “Antenna Theory, 3rd. Edition.” *ISBN: 0-471-66782-X*, 2005.
- [Bax96] Larry K. Baxter. *Capacitive Sensors*. John Wiley and Sons, 1996.
- [BHG69] E. G. Barham, W. B. Huckabay, R. Gowdy, and B. Burns. “Microvolt electric signals from fishes and the environment.” *Science*, **164**:965–968, May 1969.
- [BJB06] A. Brannon, M. Jankovic, J. Breitbarth, Z. Popovic, V. Gerginov, V. Shah, S. Knappe, L. Hollberg, and J. Kitching. “A Local Oscillator for Chip-Scale Atomic Clocks at NIST.” pp. 443–447, June 2006.
- [Blo09] Eric Blossom. “Exploring GNU Radio.”, April 2009.
- [Boa] BoatHarbors.com. “Pier 44 Marina, Mindinao Road, Marina Del Rey, CA.” ”http://www.boatharborslocator.com/ca/los_angeles_county/ca_los_angeles_pier_44.htm”.
- [BOA91] A. Black, E. Ozbay, B.A. Auld, and D.M.Bloom. “Colliding Pulse Phase Detector For PicoSecond Resolution Timing Measurement.” *Electronics Letters*, **27**(18):1620–1621, August 1991.
- [Bro07] David W. Browne. *Exploiting the Spatial Structure of Radio Communication Channels*. PhD thesis, Univ. of California, Los Angeles, June 2007.
- [Bul99] Theodore H. Bullock. “The Future of Research on Electroreception and Electrocommunication.” *J Exp Biol*, **202**:1455–1458, 1999.
- [Cab09] Dr. Danijela Cabric. “Software Defined Radio; Lecture Notes.” *EE 209AS, University of California, Los Angeles*, Winter 2009.
- [Che] V.S. Chernyak. “Fundamentals of Multisite Radar Systems, Multistatic Radar and Multiradar Systems.”.
- [Che07] M. Cherniakov. *Bistatic Radar: Principles and Practice*. Wiley: West Sussex, 2007.

- [Chi71] Der Tau Chin. “Anodic Mechanism of Electrochemical Machining: Study of Current Transient on a Rotating Electrode.” *Journal of The Electrochemical Society*, **118**(1):174–179, 1971.
- [Cor01] Loren Data Corp. “Forward Looking Ground Penetrating RADAR Market Survey.” CECOM Acq. Center-Washington, AMSEL-AC-WB-B, May 2001. 20010510/ASOL012.HTM (W-128 SN50L5K3).
- [Dar70] S. Darlington. “On digital single-sideband modulators.” *Circuits Theory, IEEE Transactions*, **17**(3):409–414, August 1970.
- [DR00] W. Daily and A. Ramirez. “Electrical Resistance Tomography, Final Report.” 2000.
- [EGH99] Deborah Estrin, Ramesh Govindan, John Heidemann, and Satish Kumar. “Next Century Challenges: Scalable Coordination in Sensor Networks.” *MOBICOM*, Aug. 1999.
- [Ett09] LLC Ettus Research. “The Universal Software Radio Peripheral (USRP).” April 2009.
- [Eva97] David Hudson Evans. *The physiology of fishes*. ISBN 0849384273, 1997.
- [FB99] L. Fishelson and A. Baranes. “Distribution, morphology, and cytology of ampullae of Lorenzini in the Oman shark, *Iago omanensis* (Triakidae), from the Gulf of Aqaba, Red Sea.” *The Anatomical Record*, **251**:417–430, Jan. 6th 1999.
- [FCS08] J. Friedman, Z. Charbiwala, T. Schmid, Y. Cho, and M. Srivastava. “Angle-of-Arrival Assisted Radio Interferometry (ARI) Target Localization.” *Proceedings of MILCOM 2008*, November 2008.
- [FG94] D.D. Feldman and L.J. Griffiths. “A projection approach for robust adaptive beamforming.” *Signal Processing, IEEE Transactions on*, **42**(4):867–876, Apr 1994.
- [FHT11] Jonathan Friedman, Henry E. Herman, Newton Truong, and Mani B. Srivastava. “A 16-Electrode Biomimetic Electrostatic Imaging System for Ocean Use.” *International Conference on Sensors and Sensor Systems (IEEE SENSORS)*, October 2011.
- [FHT12a] Jonathan Friedman, Henry Herman, Newton Truong, Young H Cho, and Mani B Srivastava. “Applied At-Depth In-Ocean Differential Electrostatics – An Experimental Study.” *IEEE MTS/OES OCEANS Conference*, Oct 2012.

- [FHT12b] Jonathan Friedman, Henry E. Herman, Newton Truong, Juyao Dong, Dustin Torres, and Mani B. Srivastava. “Considerations for the Design of an Epipelagic Biomimetic Electrostatic Imaging Element.” *The Sixth ACM International Workshop on UnderWater Networks (WUWNet)*, December 2012.
- [Fri07] Jonathan Katzenstein Friedman. *IRIS: The Design of a Ultra-low Power Infrared Communications and Obstacle Avoidance Array for Fleets of Mobile Sensor Nodes Operating in Dynamic Environments*. Univ. of California, Los Angeles, 2007.
- [FTC10] Jonathan Friedman, Dustin Torres, Young H Cho, and Mani B Srivastava. “Submerged Biomimetic Electrostatic Imaging in Salt Water.” *The 9th Annual IEEE Conference on Sensors*, Nov 2010.
- [FTS10] Jonathan Friedman, Dustin Torres, Thomas Schmid, Juyao Dong, and Mani B. Srivastava. “A Biomimetic Quasi-static Electric Field Physical Channel for Underwater Ocean Networks.” *ACM Workshop on Underwater Networks (WUWNET)*, 2010.
- [GA07] Jayne M. Gardiner and Jelle Atema. “Sharks need the lateral line to locate odor sources: rheotaxis and eddy chemotaxis.” *Journal of Experimental Biology*, **210**:1925–1934, June 2007.
- [Ger87] A. B. Gershman. “Robust Adaptive Beamforming in Sensor Arrays.” *Int. J. Electronics and Communications*, **53**:1365–1376, October 1987.
- [Haw11] J. Hawkins. “Atlantic City Radio Technical Page.” *Online. Available.*, 2011.
- [HLP08] T.E. Humphreys, B.M. Ledvina, M.L. Psiaki, and Jr P.M. Kintner. “Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer.” 2008.
- [HPX88] S.M. Huang, A.B. Plaskowski, C.G. Xie, and M.S. Beck. *Electronics Letters*, **24**:418–419, March 1988.
- [HRW97] David Halliday, Robert Resnick, and Jearl Walker. *Fundamentals of Physics, 5th ed.* John Wiley and Sons, Inc., 1997.
- [HWB00] J. Hightower, R. Want, and G. Borriello. “SpotON: An indoor 3D location sensing technology based on RF signal strength.” *UW CSE 00-02-02, University of Washington, Department of Computer Science and Engineering, Seattle, WA*, 2000.
- [iCa08] iCasualties.org. “Iraq Coalition Casualty Count.” *Compiled from news reports. Online.*, 2008. <http://icasualties.org/oif/Stats.aspx>.

- [Ins06] Woods Hole Oceanographic Institute. “Micro-Modem Overview.” 2006.
- [Ins09] Texas Instruments. “SN54HC595 / SN74HC595: 8-Bit Shift Registers With 3-State Output Registers.” *TI Datasheet Catalog*, pp. 1–27, Nov. 2009.
- [Int] Trittech Intl. “Trittech Micron Data Modem - Small ROV Communications.”
- [JK61] C.W. Harrison Jr. and R.W.P. King. “Folded Dipoles and Loops.” *IRE Transactions on Antennas and Propagation*, pp. 171–187, March 1961.
- [K 98] S.R. Judah. K. L. Chan. “A Beam Scanning Frequency Modulated Continuous Wave Radar.” *Instrumentation and Measurement, IEEE Trans.*, **47**:1223–1227, October 1998.
- [KBS07] B. Kusy, G. Balogh, J. Sallai, A. Ledeczki, and M. Maroti. “inTrack: High precision tracking of mobile sensor nodes.” *Lecture Notes in Computer Science*, **4373**:51, 2007.
- [KCA02] J. Kositsky, R. Cosgrove, C. Amazeen, and P. Milanfar. “Results from a forward-looking GPR mine detection system.” *Proc. Of SPIE Conf. on Detection and Remediation Technologies for Mines and Minelike Targets*, **4742**:206–217, 2002.
- [Kra94] Bernd Kramer. “Communication behavior and sensory mechanisms in weakly electric fishes.” *Advances in the Study of Behavior*, **23**:233–270, 1994.
- [Kra99] Bernd Kramer. “Waveform Discrimination, Phase Sensitivity, and Jamming Avoidance in a Wave-Type Electric Fish.” *Journal of Experimental Biology*, **202**:1387–1398, April 1999.
- [Kre] I. Kremer. “Prostranstvenno-vremennaya obrabotka signalov.”
- [Lab05] Ocean Climate Laboratory. *World Ocean Atlas*. U.S. National Oceanographic Data Center, 2005.
- [Lee06a] David Chun-Keung Lee. *Ragobot: A Mobile Sensor Platform for Controlled Exploration*. Univ. of California, Los Angeles, 2006.
- [Lee06b] Mark Lee. “The art of capacitive touch sensing.” *EE Times*, March 2006.
- [LF92] M. F. Land and R. D. Fernald. “The evolution of eyes.” *Annual Review of Neuroscience*, **15**:1–29, October 1992.
- [Lin] Inc. LinkQuest. “19,200 Baud, Power Efficient, Highly Robust Underwater Acoustic Modem Family.”

- [LMP55] J. Leonhard, R.D. Mattuck, and A.J. Pote. “Folded Unipole Antennas.” *IRE Transactions on Antennas and Propagation*, pp. 111–116, July 1955.
- [LNN01] Michael A. Lobardi, Lisa M. Nelson, Andrew N. Novick, and Victor S. Zhang. “Time and Frequency Measurements Using the Global Positioning System.” *Time and Frequency Measurements Using the Global Positioning System*, pp. 26–33, July-September 2001.
- [LNP90] Weihong Li, Ken Nobe, and Ame J. Pearlstein. “Potential/current oscillations and anodic film characteristics of iron in concentrated chloride solutions.” *Corrosion Science*, **31**:615 – 620, 1990. Passivation of metals and semiconductors.
- [Lor05] Stefano Lorenzini. *The curious and accurate observations of Mr. Stephen Lorenzini of Florence*. Printed for Jeffery Wale, 1705.
- [LS05] J. Li and P. Stoica. *Robust Adaptive Beamforming*. Wiley-Interscience, October 2005. ISBN: 978-0-471-67850-2.
- [LVS08] A. Ledeczi, P. Volgyesi, J. Sallai, B. Kusy, X. Koutsoukos, and M. Maroti. “Towards precise indoor RF localization.” *Proceedings of the 5th Workshop on Embedded Networked Sensors (HotEmNets '08)*, June 2008.
- [LW93] K. Li and J. Wang. “Experimental Investigation on the relationship between resistivity and water saturation in low permeability rock with low salinity.” *Internal Report*, 1993.
- [LWL03] G. Liu, Y. Wang, J. Li, and M.R. Bradley. “SAR imaging for a forward-looking GPR system.” **5089**:322, 2003.
- [Mar07] Inc. Marineland Labs. “Instant Ocean Hydrometer.” *Product Catalog*, 2007.
- [Max55] James Clerk Maxwell. *On Faraday’s Lines of Force*. 1855.
- [Max65] James Clerk Maxwell. “A dynamical theory of the electromagnetic field.” *Philosophical Transactions of the Royal Society of London*, **155**:459–513, 1865.
- [Mil93] J. P. Mills. *Electromagnetic Interference*. Prentice-Hall, Inc., 1993.
- [MMD98] G. Mavko, T. Mukerji, and J. Dvorkin. *The Rock Physics Handbook*. Cambridge University Press, 1998.
- [Mol03] M. Molinari. “High Fidelity Imaging in Electrical Impedance Tomography.” *Ph.D. dissertation*, 2003.

- [Moo67] Richard R. Moore. “Radio Communications in the Sea.” *IEEE Spectrum*, **4**:42–51, Nov 1967.
- [MSM01] Malcolm A. Maciver, Noura M. Sharabash, Mark, and Mark E. Nelson. “Prey-capture behaviour in Gymnotid electric fish: motion analysis and effects of water conductivity.” *J. Exp. Biol*, **204**:534–557, 2001.
- [MT76] H. Momma and T. Tsuchiya. “Underwater communication by electric current.” *Oceans, IEEE Journal of*, pp. 631–636, Sept 1976.
- [MVD05] M. Maroti, P. Völgyesi, S. Dora, B. Kusý, A. Nadas, Á. Lédeczi, G. Balogh, and K. Molnar. “Radio interferometric geolocation.” November 2005.
- [Nav01] United States Navy. “The Right Thing (Commander Scott Waddle, USN Greeneville).” *Report of Proceedings*, April 2001.
- [Nel] Mark E. Nelson. “Electrolocation.” *Electrosensory Signal Processing (Nelson) Lab, Beckman Institute*.
- [Nel11a] Mark E. Nelson. “Biological smart sensing strategies in weakly electric fish.” *J. Smart Structures and Systems*, **8**, July 2011.
- [Nel11b] Mark E. Nelson. “Nelson Laboratory Website.” October 2011.
- [NLL04] L.M. Ni, Y. Liu, Y.C. Lau, and A.P. Patil. “LANDMARC: indoor location sensing using active RFID.” *Wireless Networks*, **10**(6):701–710, 2004.
- [NT01] National Transportation Safety Board (NTSB). “Ehime Maru and USS Greeneville Collision.” *Marine Accient Brief*, 2001.
- [Pol02] N. Polydorides. “Image Reconstruction Algorithms for Soft-Field Tomography.” *Ph.D. dissertation*, 2002.
- [PR96] Stuart B. Palmer and Mircea S. Rogalski. *Advanced University Physics*. Taylor and Francis, 1996.
- [Pre07] James Preisig. “Acoustic propagation considerations for underwater acoustic communications network development.” *SIGMOBILE Mob. Comput. Commun. Rev.*, **11**(4):2–10, 2007.
- [Pro10] Ltd. Process Tomography. “Electrical Capacitance Tomography.” *Technology Report*, March 2010.
- [Rai07] J. Raines. “Folded Unipole Antennas.” *ISBN-10: 0071474854*, p. 400, April 2007.

- [Rho07] Mark Rhodes. “Electromagnetic Propagation in Sea Water and its value in Military Systems.” *SEAS DTC Technical Conference*, 2007.
- [Rob47] William Van Roberts. “Input Impedance of a Folded Dipole.” *RCA Review*, **8**:289–300, June 1947.
- [Rob05] R.D. Roberts. “Method of locating object using phase differences among multiple frequency beacons transmitted from spaced apart transmitter sites.” U.S. Patent 7132981B1, June 2005. Harris Corporation.
- [RWD94] S. Ramo, J. R. Whinnery, and T. Van Duzer. *Fields and Waves in Communication Electronics, Third Edition*. John Wiley & Sons, Inc., 1994.
- [Sem] Cypress Semiconductor. “Cypress TrueTouch Touchscreen Solutions.” *Product Brochure*.
- [SFC08] Thomas Schmid, Jonathan Friedman, Zainul Charbiwala, Young H. Cho, and Mani B. Srivastava. “Low-power high-accuracy timing systems for efficient duty cycling.” In *ISLPED '08: Proceeding of the thirteenth international symposium on Low power electronics and design*, pp. 75–80, New York, NY, USA, 2008. ACM.
- [Shi93] O. Shivers. *BodyTalk and the BodyNet: A Personal Information Infrastructure, Personal Information Architecture Note 1*. MIT Laboratory for Computer Science, December 1993.
- [SK73] M. Siegel and R. King. “Electromagnetic propagation between antennas submerged in the ocean.” *Antennas and Propagation, IEEE Transactions on*, **21**(4):507–513, Jul 1973.
- [SL05] Y. Sun and J. Li. “Adaptive learning approach to landmine detection.” *IEEE Transactions on Aerospace and Electronic Systems*, **41**(3):973–985, 2005.
- [SPL00] D.G. Smith, S.R. Potter, B.R. Lee, H.W. Ko, W.R. Drummond, J.K. Telford, and A.W. Partin. “In Vivo Measurement of Tumor Conductiveness with the Magnetic Bioimpedance Method.” *IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING*, **47**(10):1403–1405, Oct. 2000.
- [SSP00] E.M. Sozer, M. Stojanovic, and J.G. Proakis. “Underwater acoustic networks.” *Oceanic Engineering, IEEE Journal of*, **25**(1):72–83, Jan 2000.
- [Sta06] Robert W. Stacey. “Electrical Impedance Tomography.” *SGP-TR-182*, June 2006.

- [Sys11] Industrial Tomography Systems. “Fast Electrical Impedance Tomography.” *Company Website*, 2011.
- [Tin09] C. Tingle. “Submarine Accidents: A 60-Year Statistical Assessment.” *American Society of Safety Engineers’ Journal of Professional Safety*, Sept 2009.
- [Tri09] Intl. Tritech. “Tritech AM-300 Acoustic Modem.” 2009.
- [Tur68] Karl K Turekian. *Oceans*. Prentice-Hall, 1968.
- [TZ05] S.K. Tallarovic and H.H. Zakon. “Electric Organ Discharge Frequency Jamming During Social Interactions in Brown Ghost Knifefish, *Apteronotus Leptorhynchus*.” *Animal Behavior, Journal of*, **70**:1355–1365, 2005.
- [Wes91] D. A. Weston. *Electromagnetic Compatibility*. Marcel Dekker Publishers, 1991.
- [Wol11] Andy Wolski. “Electromagnetic Waves in Conducting Media.” *PHYS370 – Advanced Electromagnetism; Part 3*, October 2011.
- [WYH06] Jack Wills, Wei Ye, and John Heidemann. “Low-power acoustic modem for dense underwater sensor networks.” In *WUWNet ’06: Proceedings of the 1st ACM international workshop on Underwater networks*, pp. 79–85, New York, NY, USA, 2006. ACM.
- [YP03] W Q Yang and Lihui Peng. “Image reconstruction algorithms for electrical capacitance tomography.” *Measurement Science and Technology*, **14**(1):R1, 2003.
- [Zel11] Randy Zelick. “Bi 463/563 Sensory Physiology Lecture Notes.” March 2011.
- [Zim95] Thomas Guthrie Zimmerman. “Personal Area Networks (PAN): Near-Field Intra-Body Communication.” *M. S. Thesis*, Sept. 1995.
- [Zim96] Thomas Guthrie Zimmerman. “Personal Area Networks: Near-field intrabody communication.” *IBM Systems Journal*, **35**(4):609–617, 1996.