

**Infrastructure for Large-Scale Tests in Marine
Autonomy**

by

Robert Andrew Hummel

S.B., Massachusetts Institute of Technology (2009)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2012

© Massachusetts Institute of Technology 2012. All rights reserved.

Author
Department of Mechanical Engineering
October 31, 2011

Certified by
Franz S. Hover
Finmeccanica Career Development Professor in Engineering
Thesis Supervisor

Accepted by
David E. Hardt
Graduate Officer, Department of Mechanical Engineering

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE FEB 2012	2. REPORT TYPE	3. DATES COVERED 00-00-2012 to 00-00-2012			
4. TITLE AND SUBTITLE Infrastructure for Large-Scale Tests in Marine Autonomy		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, MA, 02139		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)	147	

Infrastructure for Large-Scale Tests in Marine Autonomy

by

Robert Andrew Hummel

Submitted to the Department of Mechanical Engineering
on October 31, 2011, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

This thesis focuses on the development of infrastructure for research with large-scale autonomous marine vehicle fleets and the design of sampling trajectories for compressive sensing (CS). The newly developed infrastructure includes a bare-bones acoustic modem and two types of low-cost and scalable vehicles. One vehicle is a holonomic raft designed for station-keeping and precise maneuvering, and the other is a streamlined kayak for traveling longer distances. The acoustic modem, like the vehicles, is inexpensive and scalable, providing the capability of a large-scale, low-cost underwater acoustic network. With these vehicles and modems we utilize compressive sensing, a recently developed framework for sampling sparse signals that offers dramatic reductions in the number of samples required for high fidelity reconstruction of a field. Our novel CS sampling techniques introduce engineering constraints including movement and measurement costs to better apply CS to sampling with mobile agents. The vehicles and modems, along with compressive sensing, strengthen the movement towards large scale autonomy in the ocean environment.

Thesis Supervisor: Franz S. Hover

Title: Finmeccanica Career Development Professor in Engineering

Acknowledgments

This thesis would not have been possible without the support of many people. First, I would like to thank my advisor, Professor Franz Hover, whose guidance and inspiration were vital to my success. Also, I would like to thank my colleagues at the University of Southern California, Professors Urbashi Mitra and Gaurav Sukhatme and their research groups, for helping bring the kinematic compressive sensing project to life. And to those at Woods Hole Oceanographic Institution: Professors Hanumant Singh and Milica Stojanovic and your research groups, thank you for your invaluable assistance during our sea trials and for your insight into contemporary acoustic communication challenges.

The members of my own research group have all been immensely helpful. Eric, Brooks, Josh T., Josh L., Brendan, and Charlie: besides the specific contributions you have made to my project (of which there are many!), I am most grateful for our thought-provoking discussions—related to my work, marine autonomy and large-scale systems were common topics into which you all provided insight.

And to the undergraduates who have worked on our autonomous vehicles, your contributions are very much appreciated, although I'm sure my machining skills have atrophied as a result.

Of course I must also thank my friends and family, without whom I would certainly have had a much less enjoyable time these past few years.

This thesis is the result of research partially supported by the Office of Naval Research, Grant N00014-09-1-0700, and by the MIT Sea Grant College Programs Doherty Professorship in Ocean Utilization and the Finmeccanica Career Development Professorship.

Contents

1	Introduction	17
1.1	Motivations for Large Marine Fleets	17
1.1.1	Heterogeneous Fleets	18
1.1.2	Adaptive Sampling and Cooperation	18
1.2	Acoustics: the Big Challenge	19
1.3	Thesis Contributions	19
1.3.1	Autonomous Surface Craft	19
1.3.2	Low-cost Acoustic Modem	20
1.3.3	Compressive Sensing Mission Design	20
1.4	Thesis Structure	21
2	Design of Autonomous Surface Craft	23
2.1	Prior Work	24
2.2	Raft Mechanical Design	25
2.2.1	Hull	26
2.2.2	Frame	27
2.2.3	Thrusters	27
2.2.4	Electronics Enclosure	31
2.2.5	Summary of Raft Mechanical Design	31
2.3	Kayak Mechanical Design	32
2.3.1	Hull	32
2.3.2	Frame	34
2.3.3	Propulsion	34

2.3.4	Enclosure and Miscellaneous	40
2.3.5	Summary of Kayak Mechanical Design	41
2.4	Electrical System	41
2.4.1	Arduino Single-board Microcontroller	41
2.4.2	Communication	42
2.4.3	Navigation Sensors	43
2.4.4	Power System	45
2.4.5	Data Logging	49
2.4.6	Raft Thruster Control	50
2.4.7	Kayak Thruster Control	50
2.4.8	Summary of Electrical System	51
2.5	Control System	51
2.5.1	Raft Control System	51
2.5.2	Kayak Control System	58
2.6	Performance Testing	58
2.6.1	Raft Performance	58
2.6.2	Kayak Performance - Buzzards Bay Sea Trials	63
2.7	Conclusion	70
3	Development of a Low-cost Acoustic Modem	71
3.1	Underwater Acoustics	72
3.2	Commercial Acoustic Modems	73
3.3	Piezoelectricity and Underwater Acoustics	73
3.4	Digital Signal Modulation	75
3.5	Acoustic Signal Generation for Transmission	77
3.5.1	Frequency Generation	77
3.5.2	High-voltage Conversion	79
3.5.3	Signal Amplification	79
3.5.4	Summary of Transmission System	81
3.6	Acoustic Signal Reception and Decoding	83

3.6.1	Transmit/Receive Switch	83
3.6.2	Pre-amplification	85
3.6.3	Op-amp Selection	86
3.6.4	Filtering	86
3.6.5	Post-amplification and voltage limiting	89
3.6.6	FSK Demodulation	90
3.6.7	Summary of Acoustic Modem Receiver	91
3.7	Interface and Controller Board	92
3.7.1	Power	92
3.7.2	Universal Asynchronous Receiver / Transmitter	92
3.7.3	Control	94
3.8	Submersible Enclosure	95
3.9	Modem Parts List	95
3.10	Results	95
3.11	Conclusion	99
4	Compressive Sensing with Mobile Agents	101
4.1	Introduction	101
4.2	Compressive Sensing Theory	103
4.3	Related Work	104
4.4	Energy-constrained Compressive Sensing	106
4.4.1	Energy Budget and Costs	106
4.4.2	Mobility Models	107
4.4.3	Horizon Planning	109
4.4.4	Random Walk Optimization	110
4.4.5	Simulation Results	111
4.4.6	Discussion	115
4.5	One-step-ahead Kinematic Compressive Sensing	118
4.5.1	Motivation	118
4.5.2	Formulation	120

4.5.3	Examples	122
4.6	Experimental Results	127
4.7	Conclusion	132
5	Conclusion	133
A	Tables	135

List of Figures

2-1	Photograph of indoor test tank	25
2-2	Raft thruster schematic	26
2-3	CAD view and photograph of raft frame	28
2-4	Schematic of raft motor wiring	29
2-5	Motor-propeller assembly	30
2-6	Exploded view of thruster pod assembly	30
2-7	Kayak hull and thruster schematic	33
2-8	CAD views of kayak through-hull bearing assembly	35
2-9	CAD model of kayak thruster assembly	37
2-10	Kayak servo assembly	39
2-11	Photographs of kayak thruster and servo assembly	40
2-12	Photograph of kayak electrical system	42
2-13	Photograph of GPS unit installed on kayak	45
2-14	Schematic of power system	46
2-15	Flow chart of battery protection procedure	47
2-16	Power supply control circuit	49
2-17	Swistrack software vehicle tracking output	53
2-18	Motor calibration apparatus	56
2-19	Motor calibration data	57
2-20	Description of kayak position control system	59
2-21	Plot of raft station-keeping test	60
2-22	Step response of raft during station-keeping	61
2-23	Steady-state performance of raft during station-keeping	62

2-24	Heterogeneous raft fleet mission example	63
2-25	Multi-vehicle sea trials	64
2-26	Photograph of kayak and support ship during station-keeping experiment	65
2-27	Plots of kayak station-keeping performance	66
2-28	Environmental sampling tours	67
2-29	Photograph of kayak in calm conditions in Buzzards Bay	68
2-30	Photograph of kayak in moderate conditions in Buzzards Bay	69
2-31	Illustration of cross-track error	69
3-1	Illustration of the piezoelectric effect	74
3-2	Illustration of acoustic signal transmission and reception	74
3-3	Illustration of radially-expanding piezoelectric ceramic ring transducer	75
3-4	Photograph of radially-expanding piezoelectric ceramic ring transducer	76
3-5	Illustration of frequency shift keying (FSK) signal modulation	76
3-6	Modem transmitter board	78
3-7	Frequency generation circuit	79
3-8	Waveform of frequency generation output	80
3-9	Voltage regulator circuit	81
3-10	Waveform of amplified frequency signal	82
3-11	Modem receiver board	84
3-12	Schematic of transmit/receive switch	85
3-13	Schematic of a single stage of the bandpass filter	87
3-14	Maxim filter design software	87
3-15	Waveform of transducer input and bandpass filter output	88
3-16	Schematic of filter output amplifier with voltage clamp	89
3-17	Functional block diagram of FSK decoding with XR2211	90
3-18	Interface Board	93
3-19	Photograph of modem assembly	96
3-20	Modem testing results	97
4-1	Example of compressive sensing in one dimension	102

4-2	Typical robot paths	109
4-3	Illustration of random walk	109
4-4	Optimized PMF for random walk	112
4-5	Plot of compressive sensing results for three mission scenarios	114
4-6	Plot of measurements and energy budget	117
4-7	A representative 1D signal and measurement point selection	124
4-8	Summary of DCT results	124
4-9	Sea-surface temperature compression and reconstruction	126
4-10	Summary of 2D Haar results	127
4-11	Photograph of TSP sampling mission in Buzzards Bay	128
4-12	Experimental results from sea trials in Buzzards Bay, MA	130
4-13	CS reconstruction error	131

List of Tables

2.1	Raft hull parameters	27
2.2	Motor calibration results	57
3.1	Summary of commercially-available acoustic communication devices .	73
3.2	Modem parts list	98
A.1	Raft mechanical components	135
A.2	Raft electrical components	136
A.3	Kayak mechanical components, off-the-shelf	137
A.4	Kayak mechanical components, fabricated	138
A.5	Kayak electrical components	139

Chapter 1

Introduction

In recent years, the capabilities of unmanned marine vehicles have created new opportunities for scientific, commercial, and military endeavors. Increasingly, they are being called upon to perform complex, large-scale missions spanning great distances and incorporating multiple agents. Significant challenges to implementing such fleets include the severe limitations of underwater communications, and the highly dynamic and destructive nature of the ocean. These issues and others must be taken into account when designing large-scale marine networks, and often simulation is not adequate. With this motivation, the thesis charts the development of low-cost, scalable infrastructure that will assist research in large marine vehicle networks.

1.1 Motivations for Large Marine Fleets

There are many advantages to operating with multiple vehicles in a fleet, as opposed to using just one vehicle. Most simply, a fleet with more vehicles can generally cover a larger area in a given time. Multiple-vehicle fleets can also benefit from increased coverage per unit energy, because each vehicle can travel slower, and energy losses due to drag are decreased at lower speeds [15]. Additionally, a multi-agent system is robust to one or a few vehicles failing in their individual tasks. This is useful when operating in dangerous areas—vehicle redundancy allows for a more risk-averse behavior by individuals, because there are others available to compensate.

1.1.1 Heterogeneous Fleets

Multi-agent systems also allow for more complex missions by using heterogeneous fleets. In one scenario, a few vehicles with low-resolution sensors quickly survey a large area, guiding the deployment of vehicles with high-resolution sensors. Such a sampling network was tested for mapping of frontal dynamics in a coastal environment by Schmidt et al. [59] (the low-res sensors were stationary hydrophones and the high-res sensors were AUVs). Another example of a heterogeneous vehicle mission is described by Kunz et al. [35], who have developed a pair of AUVs, one designed to search wide regions near the ocean floor for hydrothermal vent plumes and the other designed for hovering in place and recording photographs and sonar maps of the newly-discovered vents. Heterogeneous fleets can achieve more than would be possible with a single general-purpose vehicle; the specialization and distribution of tasks improves the performance of the system as a whole, and provides a very compelling reason to move towards larger fleet sizes.

1.1.2 Adaptive Sampling and Cooperation

Multiple agents working together can map large-scale environmental features, such as temperature or salinity fields, with higher precision and in less time than individual agents. For example, Zhang et al. [83] develop a cooperative control law to map level curves in a scalar field, using measurements from a small fleet to maintain the center of the fleet on a level curve. Other adaptive sampling strategies also use real-time data to guide movement; these approaches work better when multiple vehicles cooperate to cover a larger region [25, 38–40]. Cooperative fleets can also use each other for navigation, thereby reducing the individual errors of each vehicle. One example is provided by Curcio et al. [14], where several GPS-equipped autonomous surface craft assist underwater vehicles with navigation. Also, Engel and Kalwa [21] propose a method in which several underwater vehicles navigate using dead reckoning, and mutual inter-vehicle range measurements are taken to reduce the rate of error covariance growth. Other large-scale, multi-vehicle missions include the NIMS system

of tethered robots to map atmospheric conditions in forests [55], the aquatic microbial observing system of robotic boats presented in [72], and underwater glider fleets traversing large regions of ocean for meteorological data and algal bloom detection in [25, 66].

1.2 Acoustics: the Big Challenge

In implementing large vehicle networks in a marine environment, several challenges exist. The most notable among those is the difficulty of underwater communication. Chitre et al. [8] describe recent progress and current challenges for underwater acoustic networking. Acoustic communications is difficult because of many factors, including the limited bandwidth, reflections and multi-path from the sea floor and surface, Doppler shifts, and time delays. In the early years of its technological development, deep water, point-to-point communication saw significant improvement; however, shallow-water communications remained difficult. In the past decade, advances in signal modulation and channel estimation have helped in these environments [36, 41, 45, 69, 70]. With these new capabilities, researchers are now looking at communication *networks* in the underwater environment. Such networks demand in-situ research development, and testing, which is expensive.

1.3 Thesis Contributions

With the motivations presented above, we develop in this thesis several capabilities that aid research in autonomous network design.

1.3.1 Autonomous Surface Craft

We develop here two scalable, low-cost surface craft capable of rapid deployment for acoustic networking research. The vehicles have the ability to quickly arrange themselves into a particular topology, hold station while tests are performed, and then seamlessly reposition for a new set of tests. Research with mobile nodes is also

possible. The low cost and ease of deployment of the vehicles enable large fleet sizes without a tremendous resource commitment. They can serve as stationary nodes, AUVs, surface vessels, and gateways.

1.3.2 Low-cost Acoustic Modem

Commercial acoustic modems have traditionally been impractical for a low-cost vehicle because their cost is often on par with that of the vehicle itself. To benefit the scalability of acoustic communication research fleets, we develop here a simple acoustic modem that can be built for a fraction of the cost of commercial modems. Our modem currently supports only the most basic modulation techniques; however, it is subject to all of the same challenges described above. Its applications include acoustic network research, as well as coastal environmental sampling stations, port and harbor operations, and other low cost-per-node situations where an expensive modem is undesirable or unnecessary.

1.3.3 Compressive Sensing Mission Design

As mobile robots are increasingly used to survey and map spatial phenomena over wide regions, a fundamental question underlying such missions is: “Given some prior knowledge of the field of interest, how should robots move to gather maximum information, subject to energy and other constraints?” As is well-known, these planning problems can scale poorly even for one vehicle, and groups of vehicles executing a planned mission will likely have very significant connectivity and navigation requirements [24, 56, 62–64, 80, 82]. At the same time, recent results in compressive sensing (CS) [6, 18] show that under certain conditions it is possible to reconstruct fields with far fewer measurements than traditionally required.

Using CS to reconstruct fields may provide a new opportunity for designing extremely simple but effective missions in autonomous vehicles. In this thesis, we investigate several approaches to mission design for CS with consideration for vehicle constraints, in particular energy budget and mission flexibility. We then present re-

sults from numerical simulations and experiments using the vehicles developed in this thesis.

1.4 Thesis Structure

In this thesis, we seek to contribute to the field of large-scale marine autonomy with the following work:

Chapter 2: Development of two low-cost, easily-deployable vehicles, the first a small raft (20 cm diameter) with a holonomic thruster configuration for precise movements and station keeping, and the second a small kayak (2 m length) with an azimuthing thruster providing high maneuverability and increased range and payload.

Chapter 3: Development of a low-cost, bare-bones acoustic modem providing underwater communications at over 100 bits per second (bps) for under 300 USD per unit.

Chapter 4: Implementation of the compressive sensing framework for sampling with mobile agents, including novel trajectory planning techniques incorporating engineering constraints such as movement and measurement costs, allowing for improved signal reconstruction from sub-Nyquist rate sampling.

This thesis is organized as follows. Chapter 2 describes the autonomous platforms developed for multi-vehicle research. Chapter 3 details the development a low-cost acoustic modem for underwater communication. Chapter 4 introduces new mission design strategies for sampling with mobile agents using compressive sensing field reconstruction. Chapter 5 concludes the thesis.

Chapter 2

Design of Autonomous Surface Craft

The autonomous vehicles developed here are designed for large-scale, multi-vehicle research experiments and are intended as a platform that is adaptable to many mission types. Of major importance is the ease of manufacture. Additionally, the vehicles should be deployable with minimal infrastructure, personnel, and other resource requirements. They should be economical and scalable. We develop here a small, powered raft suitable for calm weather and indoor operations and a robust, maneuverable kayak for ocean experiments. The two platforms utilize similar electronics packages derived from a design that can be easily reconfigured for other vehicle types.

This chapter will provide insight into the design process and documentation of the vehicle subsystems. As stated above, the electronics and sensor package is nearly identical across the two vehicles, the only difference being the thruster drive circuitry. These similarities motivate the structure of this chapter, which is laid out as follows. After briefly summarizing some previously-developed autonomous surface vehicles, we discuss the mechanical design of our vehicles, first with the raft and second with the kayak. Then we discuss the electrical systems, first with the shared components such as power, computation, and sensing circuits, and then with the vehicle-specific components. Next, we describe the feedback-control systems used in the two vehicles. Finally, we provide results from field trials and then conclude the chapter.

2.1 Prior Work

Several autonomous surface craft (ASC) have been developed for research purposes. This section will describe some of them.

MIT Sea Grant built several versions of the ASCs, starting with ARTEMIS in 1996 [58], continuing to ACES in 1997 [47] and AutoCat in 2000 [48]. These vehicles were designed for precision bathymetry, and the final iteration, AutoCat, performed well in the mapping of Boston Harbor.

Also at MIT, in 2005 Joe Curcio developed an autonomous surface craft named SCOUT [13], which was built using a kayak hull three meters in length. The vehicles are currently used for research in cooperative autonomy and environmental sampling, among other uses.

The Robotic Embedded Systems Lab at the University of Southern California operates two autonomous surface craft built with an OceanScience Q-boat hull. They have been used for cooperative control research [2].

Commercial unmanned surface craft include the Wave Glider by Liquid Robotics and the Kingfisher M100 by Clearpath Robotics. The Wave Glider is well-suited to long-duration missions—it is propelled by wave motion and generates electricity from solar power. Its low speed makes it poorly suited to acoustic network research, however its almost unlimited range is perfect for environmental sampling. The Kingfisher is designed for data collection by environmental and civil engineers using remote control. An autonomy-ready version is also available. However, the 5.7 kilogram payload of the Kingfisher limits its usefulness for our application.

Additionally, some research has been conducted using larger unmanned boats, over twenty feet in length, particularly with military applications [57]. There is very interesting work in high-speed stability from those projects, however this work is not directly applicable to our vehicle.

After reviewing the prior work with ASCs, we chose to build our own versions in order to have complete control over the design of the systems. Our sea-going surface craft is five feet long, about half the length of the SCOUT, and while this reduces our




Figure 2-1: Photograph of indoor test tank. The tank size is approximately 10 m x 4 m x 1 m deep. The raft platforms are designed for use in this tank.

maximum payload somewhat, we benefit from the smaller size with increased ease of transport and deployment. Also, the educational benefit of designing and building a new system has been substantial to both graduate and undergraduate students.

2.2 Raft Mechanical Design

The autonomous raft is designed to be a multi-purpose test platform for operation in the testing tank located in our lab; the tank size is approximately 10 m x 4 m x 1 m deep (see Fig. 2-1). A major performance requirement of the vehicle is that it have very precise and accurate positioning capabilities; a correct and controllable ground truth is necessary for proper experimentation. High speed is not important, as the tank is relatively small, and, because the testing arena is indoors, the duration of experiments is limited only by the patience of the experimenter. With these requirements in mind, we developed a craft with a circular, flat-bottomed hull and a symmetric, holonomic thruster arrangement. The basic design can be seen in Fig. 2-2.

The remainder of this section will detail the specific decisions made for the design.

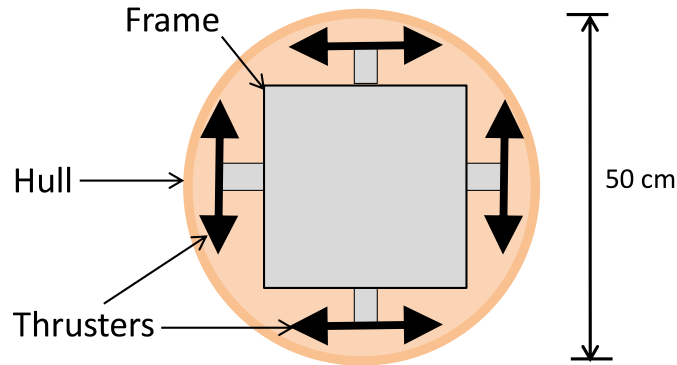


Figure 2-2: Schematic of raft hull and thrusters. The thrusters, depicted by arrows attached to the gray frame, produce thrust in the direction in which the arrows point.

A major consideration for every decision was minimizing the effort required for fabrication, assembly, parts procurement, repair, and operation. Our goal was a machine which could be built quickly and cheaply, allowing the number of vehicles to scale almost immediately to accommodate the desired experiment.

2.2.1 Hull

The circular hull is cut from urethane foam sheets (see Tables A.1 and A.2 for a complete parts list), available from McMaster-Carr and advertised as “ultra-machinable and high-strength.” See Table 2.1 for specific details about the hull. The size is a compromise between maximum payload and ease of deployment by a single person, and also to efficiently utilize standard foam sheets. The foam is cut to size with a jigsaw or wide-throat bandsaw, and the edges are chamfered slightly with a hand file. Holes for the frame and the motor wires (both discussed below) were drilled with a drill press. After the shaping and drilling of the hull, the entire piece is coated in several layers of urethane sealant. This prevents the foam from absorbing water and also strengthens it against abrasion damage. The construction time for the hull is about two hours.

Table 2.1: Raft hull parameters


Parameter	Value	Units
Diameter	50	cm
Height	10	cm
Density	0.16	g/cm ³
Volume	19.6	L
Mass	3.1	kg
Max. Payload	16.5	kg

2.2.2 Frame

To attach components to the hull we utilize an external framing system. We attach the frame securely to the hull and mount all components (motors, sensors, etc.) onto this frame. The frame consists of a top and bottom set of two identical L-shaped pieces of 1/4"-thick aluminum aligned to form a square. One of these squares is attached below the hull and one above. The frame is attached to the hull by passing threaded rods through the hull, placing spacers between the hull and the frame, and attaching nuts to both ends of the rods to secure the assembly (see Fig. 2-3a). The frame is easy and inexpensive to manufacture; the two-piece, L-shape, two-piece design is preferable to a single-piece design because several pieces can be nested into one sheet of stock and quickly machined on a water jet cutter. The spacers, threaded rods, and nuts are off-the-shelf products. The relatively large width of the spacers allows the frame to be tightened securely to the hull without damage to or indentations in the foam. The frame also serves as a handhold for lifting the raft. Mounting of external components is accomplished by drilling and tapping holes in the frame. These holes can be drilled during the initial waterjet machining or any time afterwards. Holes can even be drilled in situ, without disassembly of the frame. Frame construction requires less than thirty minutes of waterjet machining and a similar time for assembly.

2.2.3 Thrusters

The thrusters are built using off-the-shelf bilge pumps, which are designed for underwater use and are therefore waterproof. Alternatives to this design include (a) build-



(a) CAD view of raft frame



(b) Photograph of raft frame with attached motors

Figure 2-3: CAD view and photograph of raft frame

ing a custom waterproof enclosure for a motor, either depth-rated or filled with oil or (b) purchasing commercial ROV thrusters. We chose the bilge pump motors because custom sealed enclosures are complicated and prone to failure, and commercial thrusters cost several hundreds or thousands of dollars. The bilge pump motors are reliable, readily available, and inexpensive. They are available in a variety of speeds (rated in gallons per hour) and are sold as an assembly or just the “cartridge,” which consists of the sealed motor and impeller to allow replacement in situ when failure occurs. We selected pumps that were designed to survive operation in air, which allows benchtop testing without damage to the thrusters. (Some pumps require a constant source of water as lubrication for the shaft seal.)

Figure 2-3b shows the four sets of thruster pods mounted on the frame. A pair of motors are connected in an antagonistic arrangement together in a pod, and only one motor is run at a time. Two motors are used together because the thrust produced in the direction away from the thruster is substantially greater than in the opposing direction. This arrangement uses more thrusters than the minimum required for holonomic operation, however it simplifies the control of the vehicle substantially by

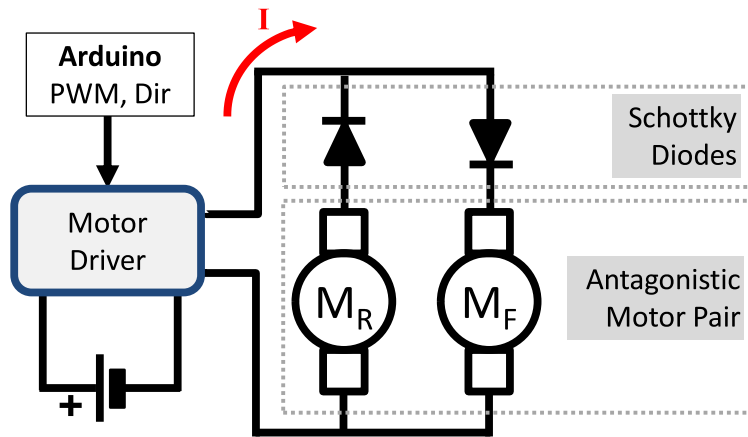


Figure 2-4: Schematic of raft motor wiring for a single thruster consisting of two opposite-facing motors. The Schottky diodes allow current to flow in only one direction. When the current, I , is positive, the forward thruster, M_F is active. When I is negative, the reverse motor is active. The motor driver controls the speed and direction of the thruster with a PWM signal generated by the Arduino microcontroller.

providing symmetric thruster performance from each pod. We use only one motor at a time, which conserves the energy that would be wasted by operating the motor in the direction that produces negligible thrust. To ensure the thrusters do not spin in reverse, a Schottkey diode is placed inline with each motor, which also allows a single motor driver to control both of the motors in a pod (see Fig. 2-4).

The motor assembly is prepared by replacing the pump's impeller with a propeller designed for a model RC boat. A propeller coupling secures the propeller to the motor; these couplings are fabricated with some amount of machining, but, because there are eight of them per raft, care was taken during the design to minimize the cost and time of fabrication. The coupling is shown in an exploded view of the motor-propeller assembly in Fig. 2-5; the caption details the steps necessary for its manufacture. We outsourced the fabrication and received the parts at a cost of three dollars each for 34 pieces (assuming \$30 per hour for machining time, only six minutes per part, neglecting material costs).

The thruster pod is assembled as follows. An exploded view is shown in Fig. 2-6.




Figure 2-5: Motor-propeller assembly. The propeller is attached to the motor with a custom coupling. The quantity of the coupling required per raft (8) motivated design for minimum machining steps. The part requires one cutoff operation, two drilling operations, and two tapping operations. Its production was outsourced for the price of three dollars per part for 34 pieces.




Figure 2-6: Exploded view of thruster pod assembly. The thruster pod consists of two bilge pump motors fit inside of a PVC pipe that has a slot removed. A clamping conduit hanger fits around the pipe. The hanger is then mounted to the raft frame.

To connect the motors together and provide for mounting to the frame, a PVC pipe four inches long is cut longitudinally at two places 45 degrees apart, creating a slot. Friction tape is applied to the ends of the motors and they are inserted into each end of the pipe. The wires coming out the back of the motors are routed through the slot in the pipe. The motor-pipe assembly is installed inside of a stainless steel, bottom-loading clamping hanger, which is normally used for hanging conduit and piping and matches the size of the pipe used here. The hanger is then attached to a tab which extends out from the frame. A total of four of these thruster pod assemblies is used in a raft. The approximate time for manufacture and assembly is two hours per raft.

2.2.4 Electronics Enclosure

An electronics enclosure is mounted to the top frame. The enclosure is waterproof-rated to NEMA 6P, that is, resistant to the ingress of water from splashes and temporary submersion at a limited depth. We have used the rafts outdoors in slightly choppy water with no water penetration problems.

2.2.5 Summary of Raft Mechanical Design


Tables A.1 and A.2 list the parts used in the raft, along with prices, vendors, and other useful information. The total cost is about fourteen-hundred dollars. Note that (a) many parts are off-the-shelf components with minimal modification, and (b) custom components require only a few fabrication steps. This vehicle platform is highly flexible—with the addition of batteries and motor control circuitry the vehicle can be configured to hold many types of sensors, such as a camera, Doppler velocity log, depth sounder, laser rangefinder, side scan sonar, GPS, inertial navigation system, among others. Thus the platform can be quickly reconfigured for a broad range of research topics in marine robotics.

2.3 Kayak Mechanical Design


The autonomous kayak is designed to have many of the same capabilities as the rafts, namely high maneuverability and precise positioning, with improvements in speed and seaworthiness as a result of its size and streamlined shape. The intended mission for the kayak is to be part of a multi-node, reconfigurable underwater acoustic network. Large-scale acoustic networks are becoming increasingly desirable, as discussed in Chapter 1, and the kayak provides a means to perform experiments in optimized acoustic networking. The kayak can perform as a static node utilizing its superb station-keeping abilities, as a mobile surface node, or as a proxy for an underwater mobile node. The kayak has several advantages to the alternatives for this type of research, which requires multiple nodes and quick reconfiguration of the network. As a static node, the kayak has an advantage compared to buoys or other stationary fixtures because it can be repositioned quickly and easily. As a mobile surface node, the alternative, a manned vessel, is more expensive and operationally challenging when the network size increases. Autonomous underwater vehicles are complicated and expensive; for shallow-water acoustic experiments, a kayak towing a submerged modem is a valid substitute that greatly simplifies the deployment. With this mission in mind, the performance requirements of the kayak are (a) high maneuverability and positioning accuracy while station-keeping and (b) reliable transit capability in terms of speed, stability, and maneuverability.

2.3.1 Hull

The kayak hull, shown in Fig. 2-7a, consists of a youth white-water kayak, 178 cm in length with a mass of 12 kg. The hull is designed to accommodate a paddler with mass up to 50 kg. It is made of high-density polyethylene (HDPE) and is very durable. The size of the kayak makes it easy to deploy; fully loaded, two people can be deployed off the side of a boat or dock and can likewise retrieve it even in moderate sea conditions. Its wide, flat bottom provides pitch and roll stability and large payload capacity.



(a) Kayak hull



(b) Kayak thruster schematic

Figure 2-7: Kayak hull (a) and thruster schematic (b). The hull consists of a youth-size whitewater kayak, six feet in length. The size is large enough to accommodate a large payload and small enough for easy deployment with two people. An azimuthing trolling motor is mounted through the hull of the kayak. The forward placement of the thruster gives passive stability, enhanced with two skegs mounted towards the rear. The kayak is steered by rotating the thruster. Image in (a) from Wave Sport.

2.3.2 Frame


The kayak hull is somewhat flexible, therefore we installed a frame to provide a rigid attachment points for vehicle components. The frame is made from one-inch T-slot framing extrusions from 80/20 Inc. Two rails are installed lengthwise on the inside of the hull, bolted to the floor with one-inch spacers lifting them off of the bottom. The two pieces are 38" long with 5-1/2" between them. In addition to rigidity, the frame gives us a level, flat surface for mounting other kayak components. Care must be taken when aligning the frame, as the other kayak components are referenced from it.

2.3.3 Propulsion


The propulsion system of the kayak is shown schematically in Fig. 2-7b. A thruster is mounted along the centerline, towards the bow of the vessel. It can rotate 135 degrees in each direction, providing steering to the vessel. The azimuthing thruster allows zero-radius turning, and the position of the thruster provides for yaw stability during forward motion. (Similar to a pendulum, the drag acts to the rear of the thruster providing a restoring moment on the vehicle. Locating the thruster to the rear of the center of drag results in an inverse pendulum-like instability.) To improve the stability further, we installed two symmetrical skegs to the underside of the craft, which help when heavy crosswinds are present.

Thruster

The thruster is made from a low-power trolling motor designed for single-occupancy inflatable rafts. The first step is to remove the original control box and mounting hardware, leaving only the motor, watertight shaft, and wiring. A hole is cut in the floor of the hull using a jig that fits on the frame for alignment. The shaft through-hull assembly constrains the shaft axially and laterally, while allowing for rotation without the ingress of water. An exploded view of the assembly is shown in Fig. 2-8a, and a section view is shown in Fig. 2-8b. Lateral shaft loads are opposed by a bronze



(a) Exploded view of kayak through-hull bearing assembly



(b) Section view of kayak through-hull bearing assembly

Figure 2-8: CAD views of kayak through-hull bearing assembly. The bearing block holds the sleeve bearing and shaft seal. The bearing block is attached to the hull with bolts connect it to the corresponding sandwich block inside the hull. Collars clamp onto the thruster shaft and constrain the shaft axially. Thrust bearings reduce the friction introduced by the collars.

sleeve bearing press-fit into a custom bearing block. Adjacent to the bearing is a PTFE spring-loaded shaft seal, which keeps out water. The bearing block is attached on the outside of the kayak by sandwiching the hull between the bearing block and a corresponding block (called the “sandwich block”) inside the hull. These blocks are bolted together very snugly to ensure loading is opposed by the friction between the hull and the blocks, rather than shear on the bolts or bolt holes. On each side of this sandwich assembly we attach a thrust bearing and shaft collar. The collars provide axial constraint to the shaft, and the thrust bearings limit the friction introduced by the collars.

Thruster Rotation

Rotation of the thruster is achieved with a position-controlled servomotor. The servo output is geared down at a 3-to-1 ratio to achieve 270 degree rotation of the shaft from the 90 degrees rotation available from the servo. A 60-tooth gear is attached to the servo output shaft using a clamping hub. The output shaft is constrained with a bronze sleeve bearing. The servo and sleeve bearing are attached to a 3d-printed ABS mounting plate designed for this purpose. This mounting plate also holds a sleeve bearing for the thruster gear shaft. The mounting piece is attached to the frame with three vertical aluminum plates, two bolted to the frame and a third attached to those for structural support. Figure 2-9 shows a CAD model of the entire propulsion assembly.

Attaching the 20-tooth gear to the thruster shaft was somewhat complicated; the following issues were considered in designing a solution. First, the gear must be securely attached to the shaft when in use, while also allowing for disassembly and removal of the thruster. To allow for disassembly, any components affixed to the shaft must be less than the shaft diameter or be removable. Additionally, the shaft material is plastic, so pins or set screws for torsional constraint would wear down the material in a very short time. Also, the motor wires must pass through the center of the shaft; an exit hole in the side would excessively weaken the shaft if it were to be big enough for practical assembly.




Figure 2-9: Kayak thruster assembly. The thruster goes through the Through-Hull Bearing Assembly and is rotated with a servo-motor attached to the Servo Mount Assembly. The Servo Mount Assembly is attached to the kayak with the Vertical Side Plates that mount to the rails that run the length of the kayak and are bolted to the hull. The Vertical Rear Plate provides rigidity to the assembly.

Our solution is as follows. An exploded-view CAD model is shown in Fig. 2-10. A cylindrical shaft coupling is made to press fit inside the thruster shaft. A longitudinal groove is milled out of one side to allow the motor wires to pass around the coupling. The center of the coupling is drilled out to accommodate a hardened steel shaft which fits the bore of the thruster gear. A flat is ground onto one end of the gear shaft, and a set screw goes through the coupling to secure the shaft in place. The gear is attached to the gear shaft with another set screw. The gear shaft is constrained by a sleeve bearing inserted into the servo mounting plate. This bearing provides the second constraint to the thruster shaft, and its proximity to the servo ensures the gears will mesh properly. An E-clip is attached to the gear shaft above the servo plate to hold it in position when the set screws are loose.

Figure 2-11 shows photographs of the propulsion system from below and inside the kayak hull.

Assembly and Fabrication

The assembly process is performed in the following order. First the shaft coupling is permanently press-fit into the thruster shaft. Then, the thruster gear shaft is installed in the servo mounting plate using the retaining e-clip. Next the thruster gear is put in place. The thruster is then inserted into the through-hull assembly, and the coupling is fit onto the gear shaft. The set screw of the gear is tightened so as to properly mesh with the servo gear. Then the two shaft collars of the through-hull assembly are tightened with the thruster at the desired height. Next the coupling set screw is tightened in place. This design allows for quick assembly and disassembly; the servo can be removed independently of the thruster and vice versa. The precision alignment is performed during the installation of the frame and vertical plates, steps which need to be performed only once and facilitate easy installation of the remainder of the components.

Fabrication steps for the propulsion system are minimal. In the through-hull assembly, the bearing block and corresponding sandwich block are drilled and bored. The intricate servo mount, which holds the servo, two bearings, and attachment points




Figure 2-10: Kayak servo assembly. The thruster is rotated with a servo-motor through a 1:3 gear ratio. The thruster shaft is connected to the 20-tooth gear with a shaft coupling and gear shaft. The Servo Mounting Plate is 3D-printed and contains mounting holes for the servo, the servo gear shaft bearing, and the thruster gear shaft bearing.




Figure 2-11: Photographs of kayak thruster and servo assembly.

for the vertical plates, is produced by 3D printer. The vertical plates are cut on a waterjet. The shaft coupling requires some mill and lathe work, and the gear shaft is ground flat with a diamond hand file. The majority of the work is by performed by computer-controlled machines, therefore the system is readily reproduced. All together, the system can be fabricated and assembled with a few hours of work.

2.3.4 Enclosure and Miscellaneous

A watertight enclosure houses the electronics for the kayak. This enclosure is attached to the frame on shock-absorbing rubber mounts. Several waterproof penetrators are installed on the hull above the water line, allowing for sensors and other cables to be routed into the hull. The top opening of the kayak is covered with a neoprene cover to repel water. There is also an automatic bilge pump installed on the floor of the kayak below the electrical enclosure, however, to date, the pump has not been necessary.

2.3.5 Summary of Kayak Mechanical Design

The autonomous kayak is designed for large-scale ocean experiments, especially for research into optimization of acoustic networks. To that end, we have designed a robust, seaworthy vehicle effective at precision station-keeping and also kilometer-scale transit. It is easily and simply constructed, and is deployed with minimal effort and personnel requirements. Its large available payload provides opportunities for a multitude of mission types. These features make the vehicle well suited for multi-vehicle research projects. Tables A.3, A.4, and A.5 provide a complete parts list for the kayak. The total cost is slightly under three thousand dollars.

2.4 Electrical System

The electrical system of the vehicles, as mentioned above, provides the basis capabilities necessary for many different types of autonomous platforms. Capabilities include communications, thruster control, navigation, power management, data logging, and computation. This section will discuss how these systems are implemented in our autonomous vehicles. A photograph of the main electrical enclosure is shown in Fig. 2-12.

2.4.1 Arduino Single-board Microcontroller

We utilize an Arduino Mega 2560 single-board microcontroller as the “brains” of the system. This board interfaces with and controls the other electronics. It contains memory, digital inputs and outputs, 10-bit precision analog inputs, digital communication capabilities, PWM generation for motor control, and servo control sequence generation. The core of the system is an ATmega2560 microprocessor operating at 16 megahertz.

The greatest benefit of the Arduino microcontroller in comparison to other similar single-board microcontrollers is the ease of learning the system. The large community of people working with the Arduino platform mean that (a) most coding or



Figure 2-12: Photograph of kayak electrical system. The enclosure contains the batteries, battery and power distribution boards, Arduino microcontroller, sensor board, and XBee radios.

system problems have already been encountered by another person and solved on the Arduino forum, (b) help for new problems is readily available on said forum, and (c) there is a vast amount of example code with which to get started. Additionally, the boards are inexpensive. They provide more than enough capability for fast, low-level hardware control and to perform simple feedback control techniques. The next step for these vehicles will be to permanently incorporate alongside the Arduino a small, GNU/Linux single-board computer to handle acoustic routing protocols and other high-level research objectives.

2.4.2 Communication

Communication from the user to the vehicles is achieved with XBee radios from Digi, which operate using electromagnetic radiation at frequencies of 2.4 GHz or 900 MHz. We use the XBees in two distinct modes. The first is transparent, point-to-point mode—basically a wireless serial connection. In this mode, any data input to the XBee is immediately transmitted, and received data is immediately output to the host device. This mode is useful for direct, human-piloted control of the vehicles,

or if only one vehicle is being used. The other mode we use is the mesh networking mode. In mesh mode, the address of the recipient is specified, and the other nodes in the system will route the data throughout the network until it reaches its destination. This mode is useful when controlling multiple vehicles. Because the address must be specified in mesh mode, each data packet contains some overhead, making human-piloted control in this mode more challenging to implement.

Digi offers a wide assortment of products varying in range, throughput, networking capabilities, and price. The modems we use have a one-square-inch footprint, excluding the antenna. For long range communication, we use the XBee Pro XSC 900 MHz, with maximum range of several miles, but with a throughput of only 9600 bits per second (bps) and no mesh networking capabilities. For short range, especially in our tank with multiple vehicles, we use the XBee ZB 2.4 GHz, with throughput of 250 kbps but range of only 120 meters. The XBee ZB is capable of mesh networking. Other models are also useful, depending on the circumstances.

The XBee is connected to a serial port on the Arduino. All data is generated and processed by the Arduino, including the control commands for mesh networking mode. We use external antennas routed into the electrical enclosure to maximize the communication range.

2.4.3 Navigation Sensors

The navigation sensors installed on the vehicles are a compass and a yaw-rate gyro, as well as a GPS when operating outdoors.

Compass

A digital compass is used to measure the heading of the vehicle. The device used is the HMC6343 compass from Honeywell, a MEMS sensor that incorporates a three-axis magnetometer, three-axis accelerometer, and microprocessor to calculate heading compensated for the pitch and roll of the vehicle. The compass updates at five times per second, and communicates with the Arduino using the I²C 2-wire serial bus.

Gyro

The rate gyro measures the rotational speed of the vehicle about the vertical axis, and is used for derivative control of heading, which helps to reduce overshoot. The device used is the ADXRS614 from Analog Sensors. The output of the sensor is an analog signal proportional to the yaw rate. The gyro has a range of ± 50 degrees per second. This range is often exceeded by the vehicle, however the smaller range gives higher resolution, which we decided was an acceptable tradeoff, because we are normally commanding the rotation speed to zero.

GPS

When operating outdoors the vehicles utilize a Garmin 17x HVS GPS, shown in Fig. 2-13. With the proper differential mode equipped, the GPS has a 95 percent accuracy radius of three meters, with resolution of roughly 10 centimeters. The GPS provides position, speed, course, and time. The data is transmitted in NMEA 0183 format using a serial port connected to the Arduino. The Arduino extracts the data and converts it to a more useable form; latitude and longitude are converted to distance in meters from a local reference point using the equirectangular projection of the Earth's surface. We calculate the distance of the vehicle from the reference location as follows

$$\Delta x = R \cdot (lon - lon_{Ref}) \cdot \cos(lat_{Ref})$$

$$\Delta y = R \cdot (lat - lat_{Ref})$$

where Δx and Δy are the distances north and east, respectively, from the reference point in meters, R is the radius of the earth in meters (6378.1 km), lon_{Ref} and lat_{Ref} are the longitude and latitude, respectively, of the reference point, and lon and lat are the respective longitude and latitude of the vehicle (all values in radians). This approximation is very accurate for distances less than several tens of kilometers, and it does not put undue load on the microprocessor; it uses only a single trigonometric function, which need be performed only once per reference point.



Figure 2-13: Garmin GPS installed on kayak

2.4.4 Power System

The main goals of the power system are to (a) protect the main battery pack from conditions that could cause permanent damage and (b) provide efficient distribution of power at the proper voltage levels to the other electrical components. The power system, discussed below, is shown schematically in Fig. 2-14.

Battery System

The main power source of the vehicles is 3- or 4-cell lithium ion-polymer (Li-poly) batteries (3-cell on raft, 4-cell on kayak, nominally 3.7 volts per cell). These batteries offer a power-to-weight ratio superior to that of lead-acid or other rechargeable battery types. One challenge, however, is that if the batteries are discharged below three volts per cell permanent damage will occur as a result of irreversible chemical reactions. This will cause the battery performance to degrade dramatically. To prevent this from happening, we control access to the Li-poly batteries and only use them if the voltage is at the proper level. This access scheme requires an additional source of power for times when the Li-poly batteries are not in use, therefore we use an auxiliary battery pack composed of six 1.2-volt low self-discharge AA-size nickel-metal hydride cells (NiMH).

As shown in the flow chart in Fig. 2-15, when the vehicle is initially powered on, the auxiliary pack is used. The Arduino checks the voltage of the main pack, and, if it

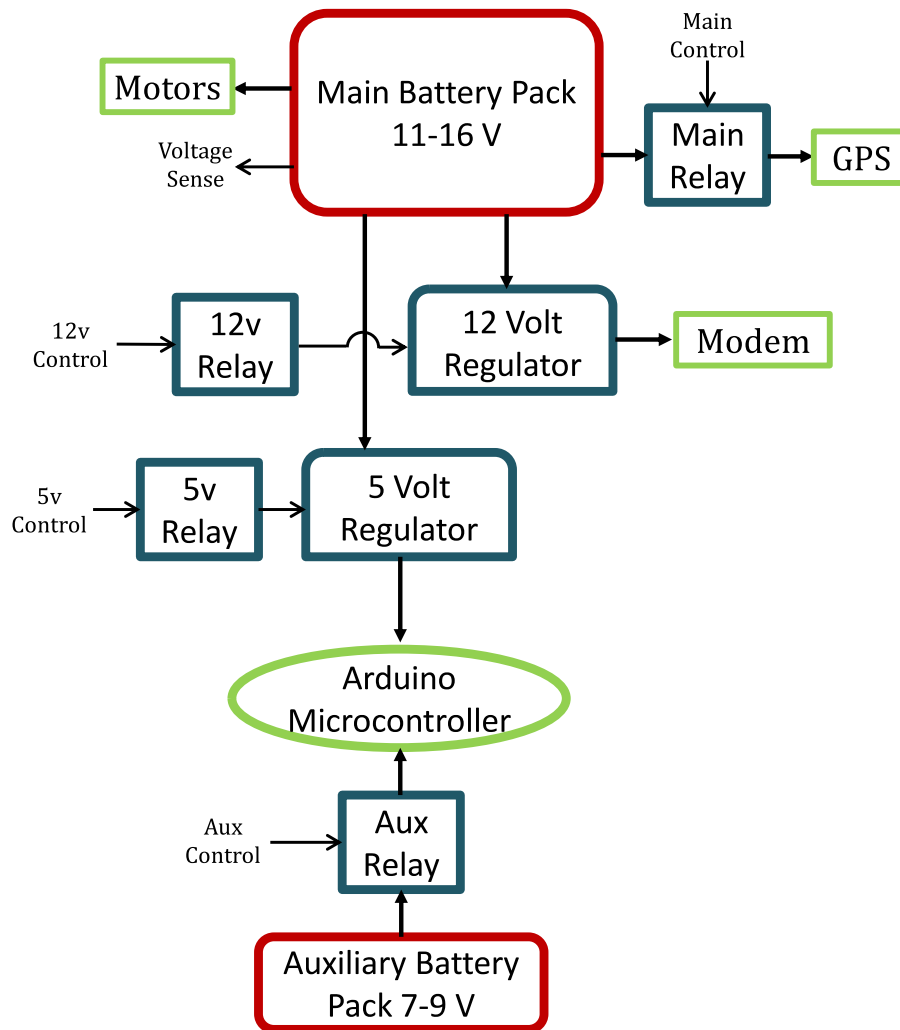


Figure 2-14: Schematic of power system. The Arduino controls distribution of power from the the main and auxiliary battery packs to the various subsystems. The relays allow the main battery pack to be disconnected if its voltage drops too low, which could permanently damage the pack.

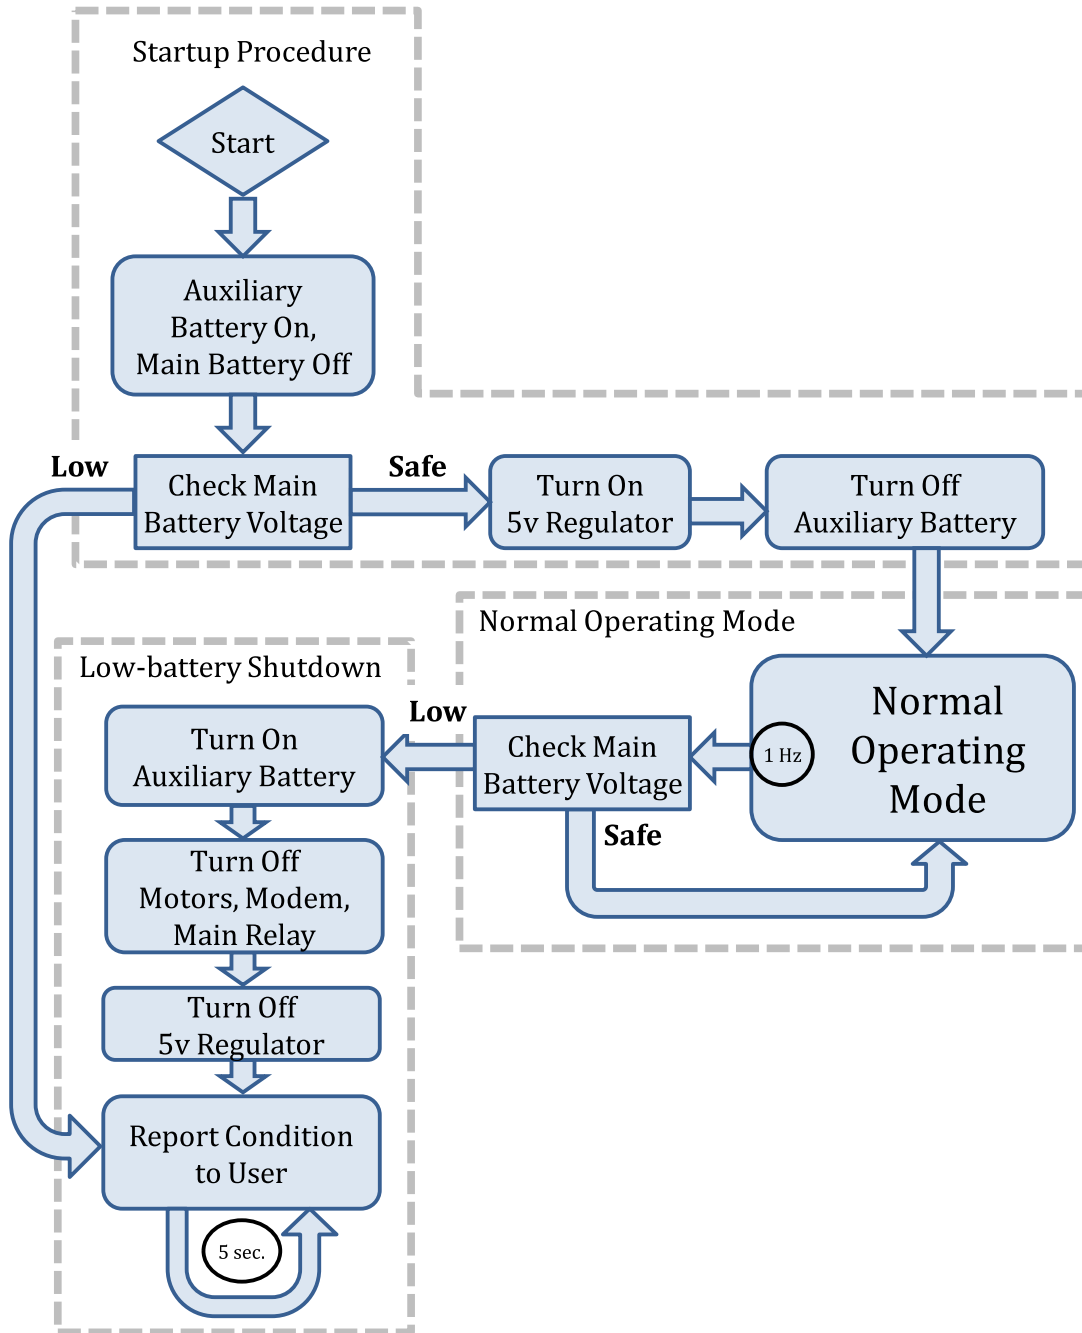


Figure 2-15: Flow chart of battery protection procedure. On startup, the auxiliary battery pack is used. The voltage of the main battery pack is measured, and, if it is properly charged, the main pack powers the system. Periodically, the main pack voltage is measured. If it falls too low, the Arduino enters the Low-battery Shutdown mode. The motors, modem, and GPS are turned off and the Arduino returns to auxiliary power. This prevents further depletion of the main battery pack which could cause permanent damage to the batteries.

is good, the main battery pack is powered on. Then the auxiliary pack is turned off. When the main battery pack voltage drops near a damaging level, an alert is sent to the user, and the main pack is turned off. In the event that there is no user nearby to turn off the system, the auxiliary pack will be depleted, however this does not damage the NiMH cells. (The NiMH cells are also less expensive than Li-poly.) The battery pack switching is controlled with several solid state relays. The auxiliary pack relay is a normally-closed (NC) type, making it the default power source. The main pack relay is a normally-open (NO) type, so the main pack is only used when explicitly engaged by the Arduino. Additionally, the power supplies (discussed below) that use the main pack are controlled with relays. The voltage of the main pack is monitored through a voltage divider with low-pass filtering capacitor, read with an analog input pin on the Arduino.

Power Distribution

The purpose of the power system is to control the flow of power to system components and regulate the voltages to the proper levels.

The Arduino requires five volts, and has an onboard step-down regulator. However, this is a linear regulator, which effectively burns off the excess voltage as heat. With a main battery voltage of 11.1 V in the raft and 14.8 V in the kayak, a linear regulator would waste the majority of our energy and get quite hot. Therefore, we use a switched-mode power supply (SMPS), which is nearly 90 percent efficient, to generate the 5-volt supply. The auxiliary pack is at 7.2 volts, so we feed this directly into the linear regulator, because the minimum input voltage of our SMPS is nine volts, and the energy loss is small.

Some sensors require 3.3 volts, and, conveniently, the Arduino contains a 3.3 volts regulator that is powered from the five volt line. We use this 3.3 volt source for the compass. Additionally, the acoustic modem that we use requires a regulated 12 volt source, which is provided with another SMPS. The GPS contains its own regulator, so we feed it power directly from the main battery.

The 5- and 12-volt power supplies each contain an on/off control pin. When the




Figure 2-16: Power supply control circuit. The voltage regulators contain an on/off control pin which can be used to enable or disable the device. When the pin is grounded, the device is off. When it is pulled high, the device is on. To drive the pin, a normally-closed relay is connected as shown with a pull-up resistor. When the relay is not activated, the regulator is off. When current is applied to the relay control input, the relay opens, turning on the regulator.

pin is grounded, the supply is off. When the pin is floating or pulled high, the supply is on. The circuit shown in Fig. 2-16 is used for each supply. The NC relay, when not activated, grounds the control pin. When the Arduino activates the relay with a high logic signal, the relay opens, pulling the control pin to the supply voltage and enabling the supply. These relays can be operated independently in order to disable the modem when not in use.

2.4.5 Data Logging

Data logging is performed using an OpenLog microSD flash memory card reader/writer. The card is very simple to use: on power-up, it increments a counter and creates a new file with the current count as the file name. Then, any serial data received by the device is written to that file. We connect an Arduino serial port to the device, as well as a digital output pin to the device reset line. When we want to create a new

file, we reset the device and send a short introductory message to identify the file with pertinent information such as the mission parameter values in use at the time.

2.4.6 Raft Thruster Control

Control of the raft motors is achieved with the use of the Arduino and four VNH3SP30 motor drivers from STMicroelectronics. The Arduino generates a pulse-width modulated (PWM) signal, the duty cycle of which corresponds to the proportion of full battery voltage to send the motors. This PWM signal can source only a few milliamps of current, so it is fed into the motor drivers which apply the full current available from the batteries at the duty cycle set by the PWM signal. As described in Section 2.2.3, we use one motor driver for each thruster pod pair of motors. The direction of the current flow, set by the Arduino and motor driver, determines which motor is active.

As a protection feature, inside the code that runs the PWM we set a maximum voltage which was empirically determined to correspond to the maximum safe operating current of the motors. If the positioning control algorithm computes a voltage higher than this safe maximum, the motor control code will threshold the output to the safe level.

2.4.7 Kayak Thruster Control

The control of the kayak thruster is very similar to that of the rafts, with the addition of the angular positioning control. The thruster speed is controlled by PWM from the Arduino, and that signal is sent into a discrete H-bridge motor driver developed by Pololu Robotics. This driver can handle the higher currents present in the kayak. To protect the motor we utilize the safe-operating-voltage thresholding procedure described above in the raft control section.

The angle of the thruster is determined by a servomotor. This servo is controlled by the Arduino, with a method of PWM different than that used to drive DC motors. For this type of servo PWM, a pulse is sent approximately every 20 milliseconds, and

the length of that pulse, rather than the specific duty cycle, sets the position of the servo.

We also utilize software control strategies to protect the servo; we have had problems specifically with the inertia of the thruster destroying the servo. Once, a bug in our position control algorithm caused the servo to rotate back-and-forth at high speed and frequency. While the thruster was spinning very quickly in one direction, the servo would be trying to spin it in the other direction. This violent motion caused a tooth on a titanium gear inside of the servo to shear off, and the tooth caused more damage to the other gears. To prevent these wild oscillations, we implemented inside of the servo control code a limit on the maximum rate of rotation. Thus, even if the high-level positioning algorithm requests a high rotation rates, the servo will instead move only as fast as the limit allows.

2.4.8 Summary of Electrical System

The electrical system described here provides the basic functionality any mobile robotic platform would require. As shown, the motor control system can be easily tailored to suit the needs of the particular application, while the remainder can be used with minimal modification. We have integrated motion control, sensors, power distribution, data logging, and communications, using an Arduino Mega 2560 as the computation and fusion center.

2.5 Control System

2.5.1 Raft Control System

Position control of the raft relies on keeping the thrusters aligned with the global cartesian coordinate frame. Using this alignment the X thrust maps directly to one pair of thrusters, and the Y thrust maps to the other pair. A high-speed (100 Hz) loop controls rotation, while a slower loop (1-15 Hz) controls translation. The translation control assumes the heading is accurately controlled.

An overhead camera tracking system provides position measurements when the rafts are operated in the lab. The tracking is performed using the open-source program Swistrack [44]. Images of the tracking procedure in Swistrack are shown in Fig. 2-17. Swistrack identifies the rafts using grayscale thresholding, pixel density, and size, and these parameters are tuned manually to match the markers placed on the rafts. For multi-vehicle operation, the markers contain codes which identify the vehicles. The tracking data contains the position of each vehicle with sub-centimeter accuracy. This data is sent to Matlab using an internal TCP/IP port. Matlab then computes the control action and sends the information to the rafts using XBee radios. The raft receives these messages at approximately 15 Hz and places a zero-order hold on the control actions between updates.

When using the raft outdoors we utilize GPS for position sensing. In this situation, all control actions are computed onboard the Arduino. GPS updates arrive at one second intervals.

The following is a description of the procedure used to calculate the control action and corresponding motor forces and voltages. First, the heading control torque, T , is computed using a proportional-derivative (PD) controller, with compass and gyro measurement inputs, θ and $\dot{\theta}$, respectively. The equation is

$$T = -k_p^\theta(\theta - \theta_0) - k_d^\theta\dot{\theta},$$


where k_p^θ and k_d^θ are the proportional and derivative gains for heading control, respectively, and θ_0 is the heading offset between the compass and the X-axis of the coordinate frame. We then calculate force, F_T needed by each motor to achieve this torque.

$$F_T = \frac{T}{4r},$$

where r is the radial distance of the motors to the center of the raft. Next, we threshold this force to leave at least some fraction of the available motor force free




(a) Camera Input



(b) Grayscale Thresholding



(c) Vehicle Detection



(d) Vehicle Identification

Figure 2-17: Swistrack software vehicle tracking output. The position of the rafts is measured with an overhead camera and Swistrack tracking software. Identifying markers are placed on the top of each raft. Figure 2-17a shows the camera image input to the software. The image is processed with conversion to grayscale and thresholding, shown in Fig. 2-17b. Vehicles are detected based on the size and density of the marker, shown in Fig. 2-17c. Identification of the vehicles is performed using the coded marker, shown in Fig. 2-17d. The position of the vehicles is output to Matlab over an internal TCP/IP port.

for X-Y position control.

$$F_t = \min(|F_T|, f \cdot F_{max}) \cdot \text{sgn}(F_T),$$

where F_t is the thresholded force, F_{max} is the maximum positive force produced by a single motor, and f is the maximum fraction of available force to use for heading control. Typically $f = 0.75$.

Next, the X-Y position control actions are combined with the heading control actions. When using camera-based position sensing, we use a PD controller, while when using GPS we use only a proportional controller. These body-referenced X and Y forces, F_x and F_y , respectively, are calculated assuming the vehicle heading is correct, and therefore the thrusters are aligned in the global X and Y directions, so that

$$\begin{aligned} F_x &= -k_p^x(x - x_0) - k_d^x\dot{x}, \\ F_y &= -k_p^y(y - y_0) - k_d^y\dot{y}, \end{aligned}$$

where x_0 and y_0 are the desired coordinates. If using GPS-only control, $k_d^x = k_d^y = 0$ because the velocities are not measured with high accuracy.

To combine the position and heading forces, we allocate all *available* force, that is, the force not being used for heading control, F_{avail} , to position control, as given by

$$F_{avail} = F_{max} - |F_t|.$$

We then determine the constant, c , that scales F_x and F_y to fit within the available force with

$$c = \min\left(1, \frac{|F_x|}{2F_{avail}}, \frac{|F_y|}{2F_{avail}}\right).$$

Then, we determine the force required by each of the four motors. Note the specific motor equations presented here depend on the motor orientations and naming

conventions.

$$\begin{aligned}
 F_1 &= F_t + c \frac{F_x}{2} \\
 F_2 &= F_t + c \frac{F_y}{2} \\
 F_3 &= F_t - c \frac{F_x}{2} \\
 F_4 &= F_t - c \frac{F_y}{2}
 \end{aligned}$$

Now we have the desired thrust of each motor and must calculate the voltage that achieves this thrust. We find the voltage of motor i , V_i , using a linear relationship between thrust and voltage, with experimentally-derived constants for each motor. Because each thruster contains two motors, the slope, m_i , and offset voltage, V_i^0 , depend on the specific motor being used in the pod, determined by the sign of the thrust. The voltage equations are therefore

$$\begin{aligned}
 V_1 &= m_1 F_1 + V_1^0 \\
 V_2 &= m_2 F_2 + V_2^0 \\
 V_3 &= m_3 F_3 + V_3^0 \\
 V_4 &= m_4 F_4 + V_4^0,
 \end{aligned}$$

where

$$m_i = \begin{cases} m_i^+ & \text{if } F_i > 0 \\ m_i^- & \text{otherwise} \end{cases}$$

and

$$V_i^0 = \begin{cases} V_i^{0+} & \text{if } F_i > 0, \\ V_i^{0-} & \text{otherwise,} \end{cases}$$

where m_i^+ and V_i^{0+} are the slope and offset voltage used when motor i is operating in the forward direction, and m_i^- and V_i^{0-} are the slope and offset voltage when the motor is operating in the reverse direction.




Figure 2-18: Motor calibration apparatus

Raft Motor Calibration

To determine the motor constants used in the calculations described above, the thrust of each motor is measured at several voltages, and a least-squares fit is performed on the data. The motor thrust is measured using the apparatus shown in Fig. 2-18. The motor is constrained in the rig and allowed only to rotate; a force sensor is attached to the raft frame parallel to the motor being tested. A voltage is applied to the motor, and the force is recorded. The raft is then rotated ninety degrees to test the next motor. The sensor is fit on one side of the rig to test the clockwise motors, and then moved to the other side to measure the counter-clockwise motors. The entire calibration procedure can be performed in about thirty minutes. Figure 2-19 shows data from the raft calibration procedure. Table 2.2 presents the results from the calibration procedure for a single raft. It is important to have equal and symmetric thrust between motors, because the position controller assumes the heading of the vehicle is properly aligned with the global reference frame. The irregularity between motors is an unfortunate side-effect of using OTS bilge pumps instead of commercial thrusters; however, we found the performance degradation caused by asymmetrical thrust to be minimal, and the calibration procedure erased the effects almost entirely. This work was performed by Alex McCarthy [52].




Figure 2-19: Motor calibration data. The calibration is performed using the apparatus shown in Fig. 2-18. During the calibration procedure, a voltage is applied to a motor. The motor output force is then measured. From this data, the slope and offset are determined for a linear voltage-force model. Figure from [52].

Table 2.2: Motor calibration results. The motor calibration procedure was performed to determine the voltage-to-thrust mapping for each of a rafts eight motors. The values for one raft are presented here. The motor voltage, V , is determined from the desired thrust, F , by the linear relationship $V = mF + V^0$. The unit of m is volts per newton, and that of V^0 is volts.

Motor	Direction	m	V^0
1	+	2.88	2.86
1	-	2.79	2.37
2	+	2.63	3.10
2	-	2.60	2.52
3	+	2.92	2.74
3	-	2.83	2.66
4	+	2.92	2.37
4	-	2.79	3.01

2.5.2 Kayak Control System

Position control of the kayak is performed with feedback from compass and GPS measurements. For simplicity, we utilize two distinct modes of operation, one for turning the vehicle (pivot) and one for forward motion (transit). The pivot mode is used when the heading error is beyond some threshold. Otherwise, the transit mode is used. If the vehicle position error very small (typically, one meter) then no thrust is applied, This section will discuss the two modes of operation and the transitions between the two, as shown in the chart in Fig. 2-20.

Pivot mode is used to rotate the vehicle in place until it points towards the target. When in pivot mode, the thruster is rotated perpendicular to the kayak major axis, ± 90 deg, the sign of which depending on the sign of the heading error. A constant thrust is applied, rotating the vehicle in place. Pivot mode continues until the heading error falls below a transition threshold (typically $\pm 45^\circ$).

Transit mode is used when the heading error of the vehicle is below the pivot transition threshold. In this mode, proportional control is used for thruster rotation and output thrust. We found that a proportional heading gain of unity provides reliable performance, in which case the applied force is aimed directly at the target. Typically, the proportional thrust gain is such that the output reaches the motor voltage threshold when more than ten meters away from target, so, in effect, there is an additional mode of operation with constant, full-power thrust and proportional heading control.

2.6 Performance Testing

2.6.1 Raft Performance

Precision Positioning

Figure 2-21 shows the trajectory of the raft during a station-keeping test. For this test, the camera position results in a pixel size of about 0.5 cm. As is shown in the step response plot in Fig. 2-22, the settling time is less than ten seconds. The

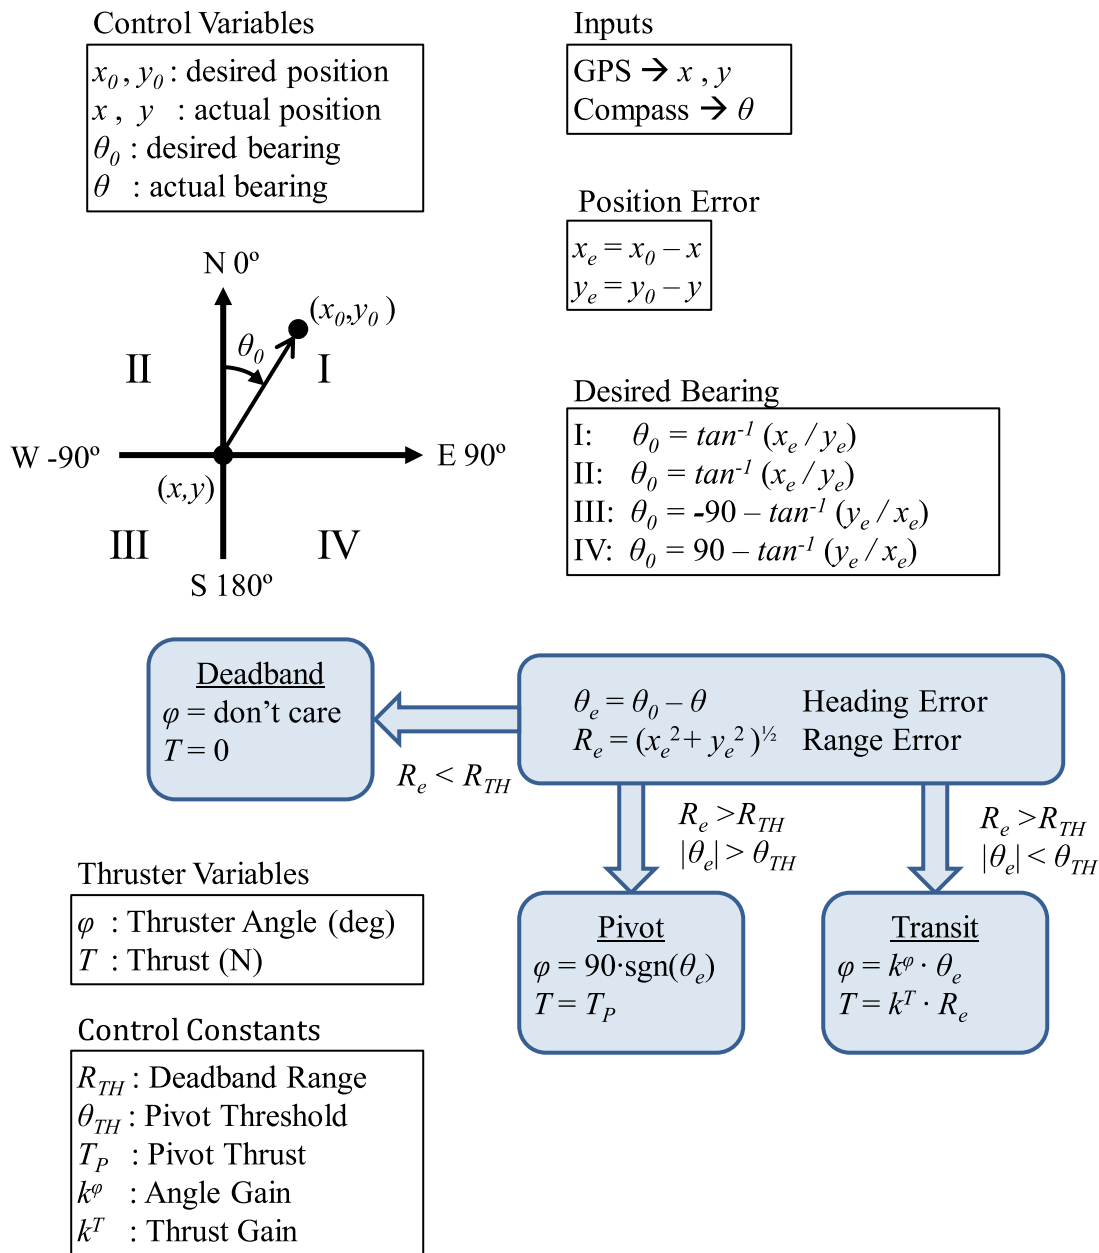


Figure 2-20: Description of kayak position control system. The desired position is set by the user with a set of one or several waypoints. The kayak measures its position with GPS and heading with a MEMS compass. Depending on the range and heading error, the controller enters one of three modes. If range error is minimal, the vehicle enters deadband mode and no control actions are performed. Otherwise the vehicle enters either pivot or transit mode. If heading error is large, pivot mode is used to rotate the vehicle until aimed at the target. When heading error is small, the vehicle uses transit mode with proportional control of heading and thrust. The thrust gain in transit mode results in full output power if range error is more than about ten meters, resulting in full-speed transit with proportional heading control.





Figure 2-21: Plot of raft station-keeping test. 1 pixel \approx 0.5 cm. The initial step response follows the trajectory entering the plot window from the top edge.

steady-state position, shown in Fig. 2-23, fluctuates within about two pixels of the target location.


Heterogeneous Sampling Scenario

A heterogeneous sampling mission was performed with two rafts. Both were outfitted with instruments to measure the floor of the tank. One raft held a downward-facing camera, the other an acoustic depth sounder. The camera was configured as a scalar sensor by reporting only the the mean value of the 10x10-pixel square in the center of the field of view, about 1 cm² on the floor of the tank. The depth sounder, with a 6-degree beam width, had a 10 cm diameter footprint, and reported the minimum depth within that footprint.

We introduced several different objects on the floor of the tank to demonstrate the usefulness of heterogeneity of fleet vehicles. One object, a cinder block colored similar to the tank floor, could only be detected using the depth sounder. Another object, a flat black piece of plastic, could only be detected with a camera. The third object, a black cinder block, could be detected by both sensors. A picture of the experimental setup is shown in Fig. 2-24a. The data from both rafts is shown in Fig. 2-24b, with the depth and camera data combined. The three objects were each




(a) Position (top) and thrust (bottom)




(b) Compass (top) and yaw rate gyro (bottom)

Figure 2-22: Step response of raft during station-keeping. In (a), the top plot shows the X and Y position of the raft (1 pixel \approx 0.5 cm), and the bottom plot shows the thruster control action. In plot (b), the compass and gyro measurements are shown. The X and Y control are turned on at time $t = 0$. Before this time, the raft is performing onboard heading control only.



(a) Position (top) and thrust (bottom)




(b) Compass (top) and yaw rate gyro (bottom)

Figure 2-23: Steady-state performance of raft during station-keeping. In (a), the top plot shows the X and Y position of the raft (1 pixel \approx 0.5 cm), and the bottom plot shows the thruster control action. In plot (b), the compass and gyro measurements are shown. The data in this plot is taken during the same trial as in 2-22 after the initial response has settled out.



(a) Experimental setup in tank



(b) Multi-vehicle data product

Figure 2-24: Heterogeneous raft fleet mission example. Two rafts were used to survey the tank to detect the three objects shown in (a). One raft, with a downward-facing camera, could detect objects with coloring dissimilar to the floor. The other raft, with an acoustic depth sounder, could detect objects with height differences from the floor. A survey was performed by both rafts of the area known to contain the objects. The data were combined and are presented in (b). All three objects were detected successfully. Images courtesy of [51].

successfully detected. This mission also serves to demonstrate the object avoidance and multi-vehicle communication capabilities of the system.

2.6.2 Kayak Performance - Buzzards Bay Sea Trials

In August 2011, we participated in sea trials in Buzzards Bay with collaborators at Woods Hole Oceanographic Institute. Buzzards Bay is directly west of Woods Hole, as shown in Fig. 2-25b. The motivating goal of the trials were to investigate newly developed adaptive routing procedures for multi-vehicle underwater fleets. Our fleet consisted of one SeaBED AUV developed and operated by Hanumant Singh and his lab, two Iver2 AUVs from OceanServer Technology, Inc. operated by Ryan Eustice,



(a)



(b) Source: Google Maps

Figure 2-25: Multi-vehicle sea trials were performed in August 2011 in Buzzards Bay. Participating vehicles were the kayak, one SeaBED AUV operated by Hanumant Singh, and two Iver2 AUVs operated by Ryan Eustice.

and our kayak. The kayak, SeaBED, and one Iver2 are shown in Fig. 2-25a.

Station-keeping Missions

On August 4 we performed several tests with the kayak holding station for an extended period of time. The kayak was equipped with a WHOI MicroModem acoustic modem. It was to hold station near the larger support ship while acoustic routing experiments took place. A picture of the kayak and ship during the experiment is shown in Fig. 2-26. Wind speed was about 10 knots with gusts up to 17-18 knots during the trials. GPS measurements collected during two station-keeping missions are shown in Fig. 2-27. The mean range error for the two missions was about 4.5 meters. Interestingly, the standard deviation of the error is small, less than one meter. The vehicle seems to be in a steady-state condition for the majority of the missions, where the proportional control force balances the wind and current forces.



Figure 2-26: Photograph of kayak and support ship during station-keeping experiments

Environmental Sampling Missions

During the Buzzards Bay trials we also used the kayak in an environmental sampling mission. The mission consisted of two tours of 100 randomly-selected points. At each point, the vehicle held station while measurements were taken. The scientific results of this mission are discussed in detail in Chapter 4. Figure 2-28 shows the path taken by the vehicle during two missions. In the figure, the cyan circles show the waypoints at which measurements were taken.

In these missions, the capture radius of the waypoints, that is, the range at which the vehicle is ‘at’ a waypoint and begins a measurement, is two meters and is equal to the deadband radius. This causes some slightly erratic behavior; in some cases, the vehicle drifts out of the deadband and attempts corrective action just as the measurement is completing. This can cause the vehicle to make small loops at some of the waypoints. A more careful coordination between sampling and stationkeeping could minimize these counterproductive actions, however they did not seem to substantially increase mission time.

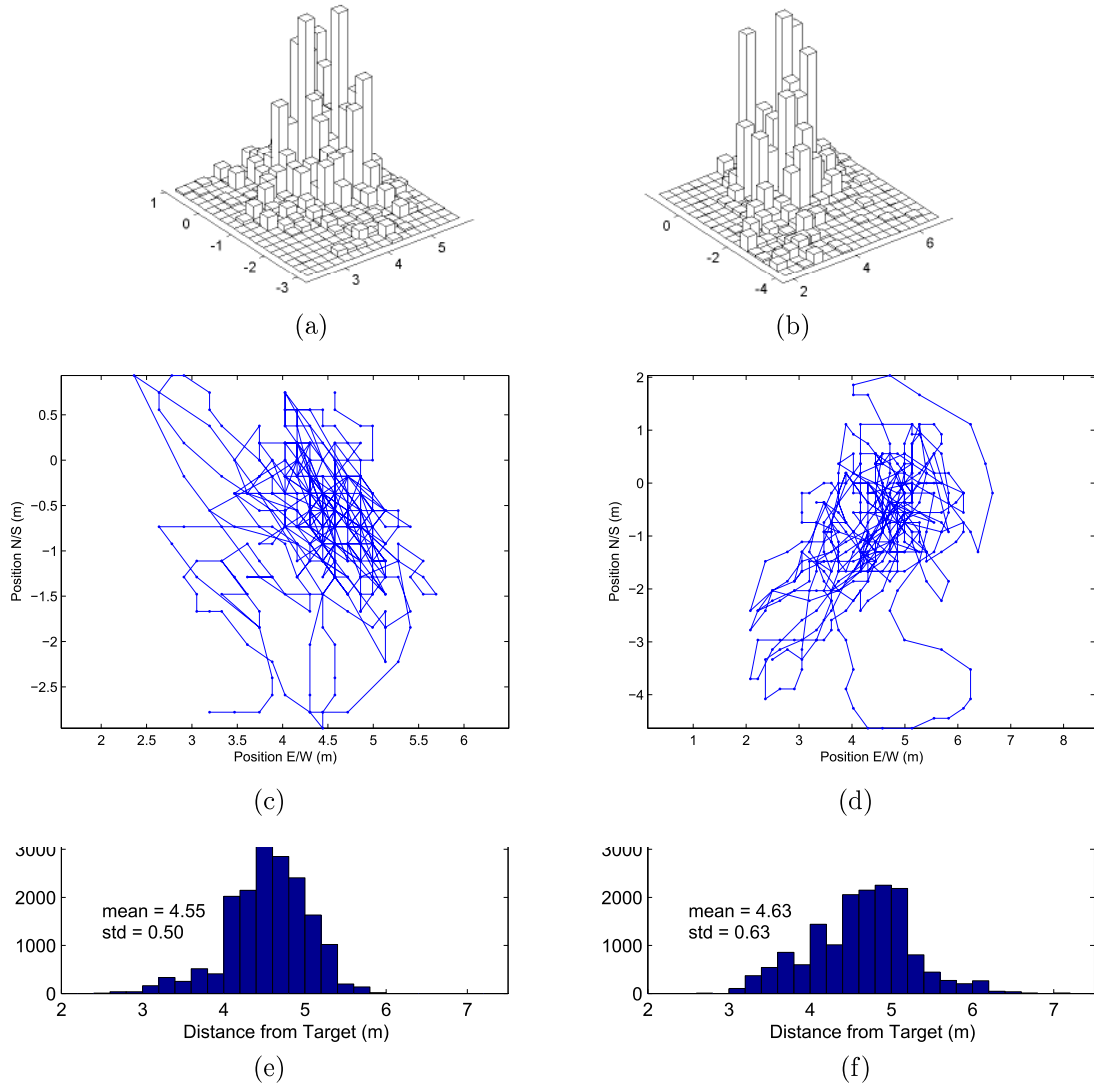


Figure 2-27: Plots of kayak station-keeping performance. Experiments were performed in Buzzards Bay, MA in support of a multi-vehicle acoustic routing experiment. Data consists of GPS measurements collected at one second intervals. Figures (a),(c), and (e) show Mission 1 which lasted for 39 minutes. Figures (b), (d), and (f) show Mission 3 which lasted 32 minutes. The target location is (0,0). Steady-state range error of 4.5 meters is a result of the proportional control balanced by the wind and current forces on the vehicle.

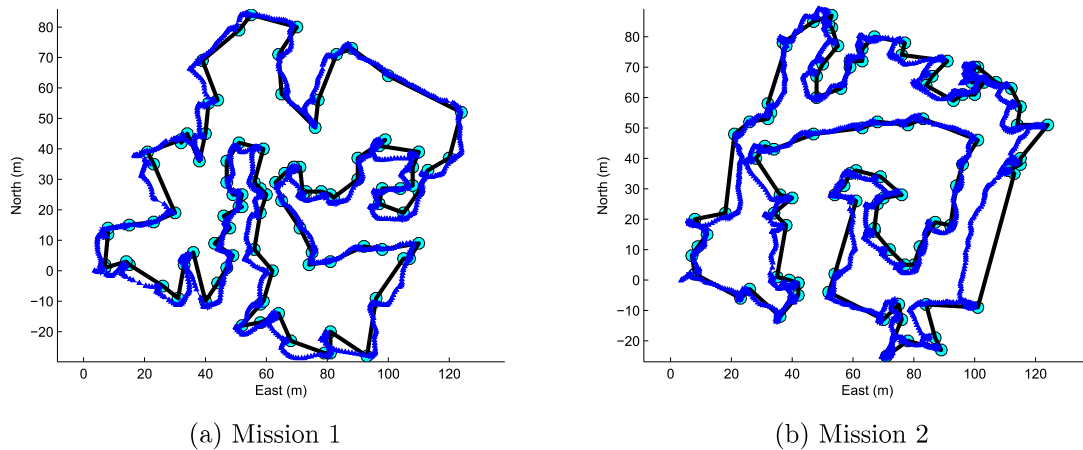


Figure 2-28: Environmental sampling tours. In these missions, the kayak was equipped with a depth sensor, fluorometer, and turbidity sensor. The missions consisted of visiting 100 random points within a 100 m x 100 m region, for the purposes of compressive sensing. *Legend: blue triangles - vehicle path (GPS meas.); cyan circles - sampling locations; black lines - ideal shortest path*

Evaluation of Vehicle Performance

Calm-weather Performance – Figure 2-29 shows a picture of the kayak in relatively calm water in Buzzards Bay. The vehicle performed well; it is very stable in roll and pitch. Using GPS measurements, we measured a top speed of 1.2-1.3 m/s. This speed will be improved with work currently in progress that will increase the thruster output force from 12 lb_F to 30 lb_F. In pivot mode, rotation rates of over 90 deg/s were observed. (Our sensor is accurate only to 50 deg/s, however we observed 360-degree rotation in under four seconds.) When holding station in calm weather, visual observation from a fixed position shows the vehicle position changing by up to several meters, however the GPS measurements shows an error fluctuating about zero and never more than about one to two meters. Our conclusions are that the vehicle can track the GPS signal to a high precision, and the fluctuations of the GPS signal are greater in magnitude than the tracking performance of the vehicle. Improvements to vehicle positioning beyond GPS precision could include Kalman filtering of inertial measurements, however the potential gains are limited.

Seaworthiness – We have operated the vehicle in many conditions, from calm seas



Figure 2-29: Photograph of kayak in calm conditions in Buzzards Bay

and light wind to moderate seas and wind. Figure 2-30 shows the kayak in moderate conditions. In the worst conditions, the average wave height was more than one meter; in fact, from our vantage point on the support ship, the kayak would disappear from view inside the trough of each wave. Wind speed at the time was 13-15 knots, with gusts nearing 20 knots. At no time did the vehicle appear in danger of capsizing. One issue was encountered when the vehicle trajectory was diagonal to the waves. The vehicle would pivot towards its target, as described in Sec. 2.5.2, however the next wave would perturb the vehicle outside the threshold angle, causing it to enter pivot mode again. This significantly slowed the forward progress of the vehicle. Our solution was to increase the threshold angle so that the vehicle would enter transit mode more frequently. This caused a longer path to be taken than the direct route, yet we observed that the kayak worked itself into a trajectory perpendicular to the waves, allowing it to make forward progress.

A bigger issue in strong weather was the drag introduced from the wind. At the time, we had a large computer box protruding about ten inches out of the top of the kayak (seen in Fig. 2-30). Making way upwind proved rather slow, about 0.25-0.5 m/s, and moving perpendicular to the wind pushed the kayak far off course. The new thruster will improve the top speed. To minimize cross-track error, we envision using a technique such as the space-based integral control implemented by Greytak [27] shown in Fig. 2-31. Also, refinement of the electrical enclosure is under way, to fit a computer within the hull and minimize the surface area exposed to the wind.



Figure 2-30: Photograph of kayak in moderate conditions in Buzzards Bay




Figure 2-31: Illustration of cross-track error. With proportional control, environmental forces will push the vehicle off of the path. Proportional-integral control can compensate for the environmental forces while staying on the path. Image from [27].

2.7 Conclusion

Future work with the autonomous surface vehicles will focus on improving the ease of deployment of the vehicles. The kayak electrical enclosure is currently undergoing work to separate the batteries from the other electronics to enable the quick exchange the battery packs to increase the duration of deployments. Also, a single-board computer is being added as a master to the Arduino microcontroller; this computer will handle external radio communications, high-level mission planning and execution, and acoustic communication protocols. Also, a higher-power thruster is being retrofitted to improve the top-speed of the vehicle.

The indoor rafts could be improved by increasing the coverage of the tracking system to encompass the entire tank. This would likely require multiple cameras, however it should be easily manageable, as Swistrack advertises this capability. Additions to the raft, such as laser range finders and other sensors, can be readily added as the desired mission dictates.

We have presented here the design and development of two autonomous surface craft intended for research in large-scale autonomy. The vehicles are low cost and scalable. Emphasis has been placed on ease of fabrication and deployment. We have further shown that the vehicles are capable of high-performance operation within their domains, and therefore will contribute to the field of distributed and multi-vehicle ocean systems.

Chapter 3

Development of a Low-cost Acoustic Modem

This chapter discusses the development of a low-cost acoustic modem. Our motivation for such a modem is the desire for large vehicle fleets with acoustic communications capabilities. As discussed in Chapter 1, commercial modems can be very expensive. A large reason for this high cost is the typical deployment scenario for a modem. When deployed on a single AUV, which can cost hundreds of thousands or millions of dollars, a ten thousand dollar modem is rather inexpensive. However, when deployed on thousand-dollar vehicles, such as those discussed in Chapter 2, modem cost can become the limiting factor to fleet size. Therefore, we endeavor to create a low-cost modem for research in large-scale autonomy.

The organization of this chapter is as follows. First, we describe underwater communications and the use of acoustics. We then briefly summarize the most common commercially-available acoustic modems. Then we look at piezoelectricity and how we use it for our modem transducers. Next, we discuss our digital-analog modulation technique. Then we detail the design of the modem circuits, first with the transmitter, then the receiver, and finally with the interface between these components and the user. We then briefly describe the underwater enclosure built for the modems. Last, we present modem performance results and conclude the chapter.

3.1 Underwater Acoustics

Underwater communication has been extensively studied using three different technologies: acoustics, optics, and radio. Optics and radio have the potential for higher data rates than acoustics, however they are not currently possible at ranges more than about 100 m. For that reason, acoustics is by far the most prevalent technology used for underwater communications. Although there are severe limitations with acoustic communications, acoustics will remain the dominant technology for some time.

Acoustic communication relies on the transmission of pressure waves through water. Although the data rates can be increased by using higher-frequency waves, the attenuation of sound in water also increases with frequency, and only frequencies up to about 100-300 kHz are useful. Lower frequencies, around 10 kHz are used for kilometer-length communication, while the higher frequencies can be used for shorter ranges. Several thousands of bits per second can be achieved over short distances, however at full ocean depth a rate of a few hundred bits per second is common. The contrast between these rates and million-bit-per-second communication above the surface highlights the severity of the data-rate constraint imposed on underwater fleet coordination and control. Additionally, signals can take several seconds to travel to their target. (The speed of sound in water is about 1500 m/s.) Reflections, absorption, noise, and unknown and fluctuating water properties compound these challenges. As stated in Section 1.2, protocols for acoustic networking must be built from scratch with these issues in mind.

The difficulty of acoustic communications motivates research into optimal use of the limited bandwidth that *is* available. Our modems operate using many of the basic principles used in commercial, state-of-the art modems, and are subject to the same physical limitations. Therefore, they should provide a useful piece of equipment for research in large-scale ocean systems.

Table 3.1: Summary of commercially-available acoustic communication devices

Vendor: Product	Reference	Range	Rate (bps)
Linkquest: (multiple models)	[43]	10 km	2500
Benthos: (multiple models)	[73]	2-6 km	15360
Desert Star: SAM-1 Miniature Modem	[16]	1 km	150
Tritech: Micron Data Modem	[75]	1 km	40
EvoLogics: (multiple models)	[22]	10 km	28k
Sonardyne: USBL (multiple models)	[67]	6 km	N/A
WHOI: Micro-Modem	[65, 77]	11 km	5400

3.2 Commercial Acoustic Modems

Table 3.1 gives an overview of several commercially available acoustic modems and modem manufacturers. The values presented are those advertised by the manufacturer, with the following exceptions. The WHOI Micro-Modem does not advertise a maximum range, however the paper referenced in the table provides full-ocean-depth experimental results. The Sonardyne USBL systems do not advertise a data rate, because the products are primarily used for tracking purposes, however they can use the measurement pings to also transport data. Note that normally with acoustic modems the ranges and data rates advertised are the maximum of those specifications possible under ideal conditions, and those rates and ranges may not be simultaneously achievable (i.e. only low data rates possible at maximum range, and maximum data rates possible only at short range).

3.3 Piezoelectricity and Underwater Acoustics

Piezoelectricity is defined as the linear electromechanical relationship between electrical potential and strain which occurs in some materials. In the direct piezoelectric effect, charge accumulates within a material as a result of mechanical strain. In the reverse piezoelectric effect, strain develops as a result of an electrical potential across the material. Figure 3-1 [34] shows an illustration of the piezoelectric effect. The direct and reverse effects were first demonstrated in 1890 and 1891, respectively, in the naturally-occurring crystals quartz and topaz, among others. Since that time,




Figure 3-1: Illustration of the piezoelectric effect. A positive electric potential across the material results in a positive displacement (tension). A negative potential results in a negative displacement (compression). The effect is also reversible – a deformation of the material results in an electric potential across it. Image from [34].




Figure 3-2: Illustration of acoustic signal transmission and reception. Transmission (left) is performed by applying an electric signal to a piezoelectric material, inducing a mechanical deformation which generates a pressure wave. Reception (right) occurs when the pressure wave causes a mechanical deformation of the receiving piezoelectric material which transforms the deformation into an electric signal. Image from [46].

man-made materials, particularly ceramics, have been developed with improved piezoelectric sensitivity. Today, the most commonly used is lead zirconate titanate, also known as PZT, which is the material chosen for our transducers.

Piezoelectric devices are used for underwater acoustics as both a transducer and receiver. As shown in Fig. 3-2 [46], a signal is transmitted through the water by applying a voltage to a piezoelectric device. The strain that develops creates pressure (acoustic) waves that travel through the water to (hopefully) the piezoelectric receiver. Here, the pressure waves cause a mechanical strain, which develops a voltage across the receiver. This voltage is processed to extract the communications signal.

The shape of the transducer affects the direction in which the acoustic waves will





Figure 3-3: Illustration of radially-expanding piezoelectric ceramic ring transducer. An electric signal is across the electrodes on the inner and outer surface of the cylinder, causing a radial deformation that produces acoustic waves. The waves are omnidirectional in the horizontal plane.

travel. We use a hollow cylindrical transducer with electrodes on the inner and outer faces, which produces waves that are omnidirectional in the horizontal plane, shown in Fig. 3-3. The specific transducers we use are shown in Fig. 3-4a. They are 20 millimeters in diameter, 5 mm tall, and 1.5 mm thick. Wires are soldered to the electrode and the device is potted in silicone sealant to protect it from water. The potted transducer is shown in Fig. 3-4b.


3.4 Digital Signal Modulation

To encode data in the acoustic signal we use frequency shift keying (FSK) [78]. As shown in the illustration in Fig. 3-5, a digital signal is represented as the frequency of the acoustic signal. In binary FSK, performed here, two frequencies are used. The frequency corresponding to a logic high signal is called the *mark* frequency, f_m , while that corresponding to a logic low is the *space* frequency, f_s .

FSK modulation was chosen because of its suitability to acoustic communications. The incoming signal can be filtered heavily to remove noise, leaving only the frequency band containing the expected signal. Additionally, FSK is less sensitive to the amplitude of the received signal, which can vary over several orders of magnitude. Also, as described below, integrated circuits are available to perform most of the functions



(a) Piezoelectric ceramic transducer. Image from Steiner and Martins, Inc.



(b) Transducer potted in silicone sealant

Figure 3-4: Photograph of radially-expanding piezoelectric ceramic ring transducer. The transducer is 20 millimeters in diameter, 5 mm tall, and 1.5 mm thick. Wires are soldered to the electrode and it is potted in silicone sealant to protect it from water.




Figure 3-5: Illustration of frequency shift keying (FSK) signal modulation. Frequency shift keying (FSK) encodes digital data in the frequency of the transmitted signal. In binary FSK, shown here, one frequency corresponds to a digital high signal, and one to a digital low signal.

required for FSK modulation and demodulation.

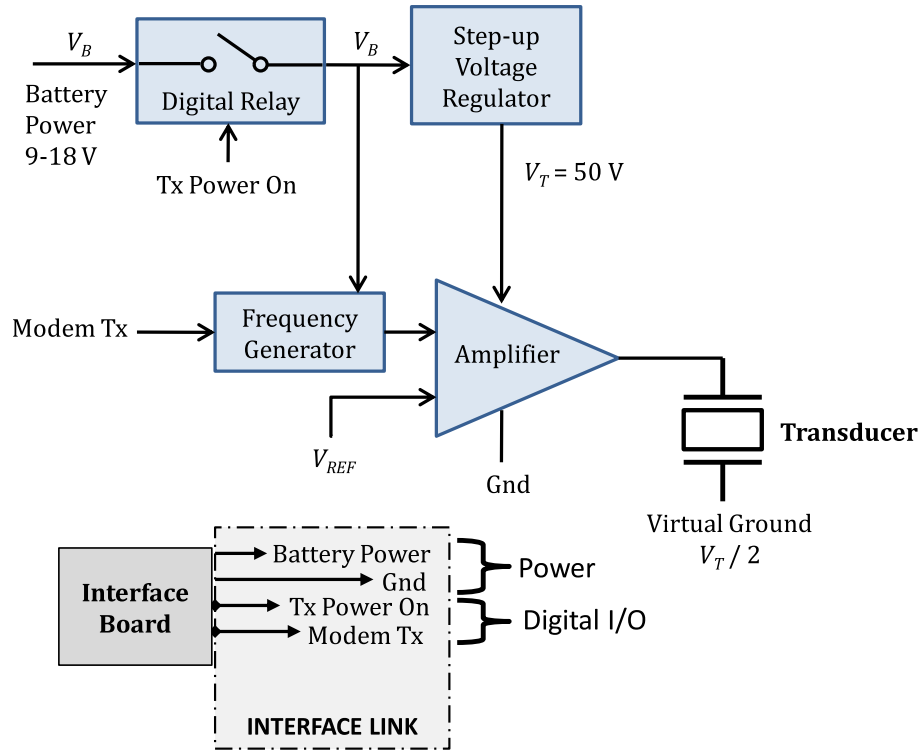
3.5 Acoustic Signal Generation for Transmission

The purpose of the transmitter is to generate acoustic waves at either the mark or space frequency, as described in Section 3.4. Fig. 3-6a shows a schematic of the modem transmitter. A frequency generator IC creates the signal at the proper frequency. A boost converter creates the high-voltage power necessary to drive the transducer. A high-voltage op-amp amplifies the signal and sends it to one electrode of the transducer. The other electrode is held constant at half of the high-voltage level.

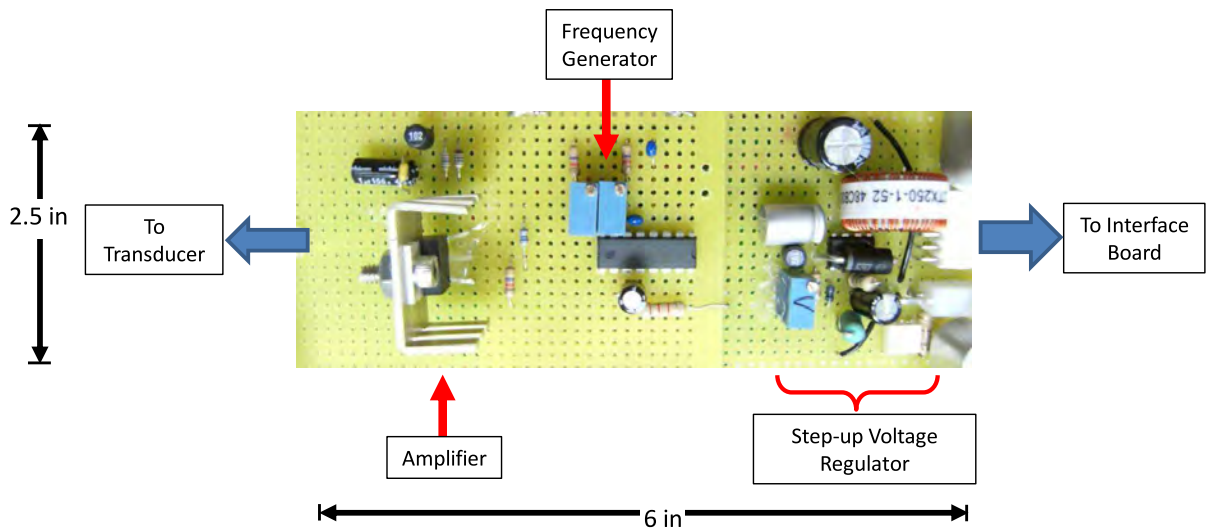
This section will discuss the subsystems of the transmitter: frequency generation, high-voltage conversion, and amplification. These functions are integrated into a single circuit board independent of the rest of the modem. Data, power, and control come from a link to a separate interface board. A schematic of the transmitter and the transmitter circuit board are shown in Fig. 3-6.

3.5.1 Frequency Generation

The frequency generation is performed with XR-2206 Monolithic Function Generator IC from Exar Corporation. The device contains a voltage-controlled oscillator that produces a signal at the frequency determined by external discrete components. Figure 3-7 shows a schematic of the system. For FSK generation, the user selects a timing capacitor, C , and two timing resistors, R_1 and R_2 . Internal current switches determine which resistor is active from the FSK input pin. This input pin is connected to the Modem Tx line coming from the interface link. If the FSK input pin is set high, then the VCO frequency is equal to $1/R_1C$, while if the input pin is low the frequency is $1/R_2C$. Thus a digital signal is converted to the mark and space frequencies used in FSK modulation. The square wave output is used, with a pull-up resistor connected to the battery supply voltage, V_B . The output of the frequency generation circuit is a square wave at the desired frequency with low voltage of zero volts and high voltage of V_B . Waveforms showing the FSK input and frequency output are



(a)



(b)

Figure 3-6: Acoustic modem transmitter schematic (a) and circuit board (b). The frequency generator creates a waveform at the mark or space frequency set by the Modem Tx input from the interface board. The signal is amplified and sent to one electrode of the transducer. The other electrode is connected to a virtual ground equal to the mean of the amplified signal. The high-voltage amplifier supply is created with a step-up DC/DC converter. Power to the board is turned on and off with a relay controlled from the interface board.

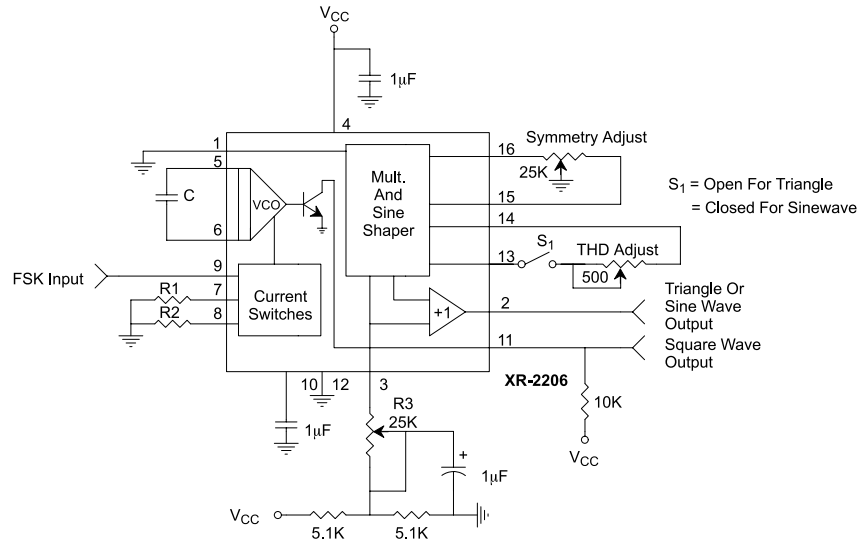


Figure 3-7: Frequency generation circuit using the XR-2206 monolithic function generator. Timing capacitor, C , and resistors, R_1 and R_2 set the frequency of the output. The FSK input pin determines which of the two frequencies is active. Schematic from XR-2206 datasheet.


presented in Fig. 3-8.

3.5.2 High-voltage Conversion


As described in Section 3.3, a relatively high voltage is required to achieve significant displacement within piezoelectric materials. The high-voltage conversion section boosts the 9-18 volt DC supply of the battery to $V_H = 50$ V DC to power the amplifier circuit. The conversion is performed with an MC34063 Buck Boost Inverting Switching Regulator IC from ON Semiconductor. As the name suggests, the IC contains the basic functions needed in a switching voltage regulator and can be used in a step-down (buck), step-up (boost), or inverting configuration. A functional block diagram of the IC is shown in Fig. 3-9. The specific circuit used for a step-up regulator is shown in Fig. 3-9b.

3.5.3 Signal Amplification

Signal amplification is performed using an LM675 Power Operational Amplifier from National Semiconductor configured as a comparator. The signal, created as described



(a)



(b)

Figure 3-8: Waveform of frequency generation output. The frequency generation circuit creates a square wave with the selection of space or mark frequencies (f_s or f_m) depending on the state of the FSK input pin. The output is connected to the battery supply, V_B , with a pull-up resistor, therefore the output swings between ground and battery voltage.




Figure 3-10: Waveform of amplified frequency signal. The input signal (Trace 2) is amplified with a comparator circuit. When the input is low the comparator output goes to zero, and when the input is high the output goes to to the high-voltage supply level. The trapezoidal shape is a result of the limited slew rate of the device. One electrode of the transducer is connected to the amplified signal; the other is connected to a virtual ground equal to half of the high-voltage value.

mode power supply. The amplified frequency signal is then sent to the piezoelectric transducer which converts it to an acoustic signal.

3.6 Acoustic Signal Reception and Decoding

As described in Section 3.3, when the transducer is deformed by an acoustic signal a voltage is developed across the two surfaces of the transducer. The purpose of the modem receiver is to amplify and decode this voltage into the appropriate digital signal. For FSK encoding, this process requires determining the frequency of the acoustic wave and whether it matches the mark or space frequency. As in the transmitter, the receiver circuitry is integrated into a stand-alone circuit board, and data, power, and control come from a link to the interface board. A schematic and corresponding printed circuit board (PCB) are shown in Fig. 3-11.

3.6.1 Transmit/Receive Switch

The first stage in the receiver circuit consists of an important feature called a Transmit/Receive (T/R) switch that protects the circuit from the high voltages used during transmission. There is one T/R switch on each of the two electrodes of the transducer. The T/R switch consists of a diode bridge as seen in Fig. 3-12. The T/R switch is AC coupled at the input and output with capacitors $C = 0.1\mu\text{F}$. The diodes are biased by the current which flows from V_+ to V_- , the magnitude of which is set with the two equal resistors R . A larger bias current causes less noise to be introduced to the received signal, however higher current consumption will also occur. A resistor value of $20\text{ k}\Omega$ was observed to introduce minimal noise with only 2.5 mW power consumption.

The T/R switch operates as follows [32]: when the transducer is transmitting, the voltage at the input is much greater than the bias voltage, causing the diodes to be reverse-biased, therefore no current flows. When the transducer is receiving, the voltage at the input is typically no greater than several tens of millivolts, much less than the diode reverse-bias threshold voltage of roughly 700 mV , therefore the

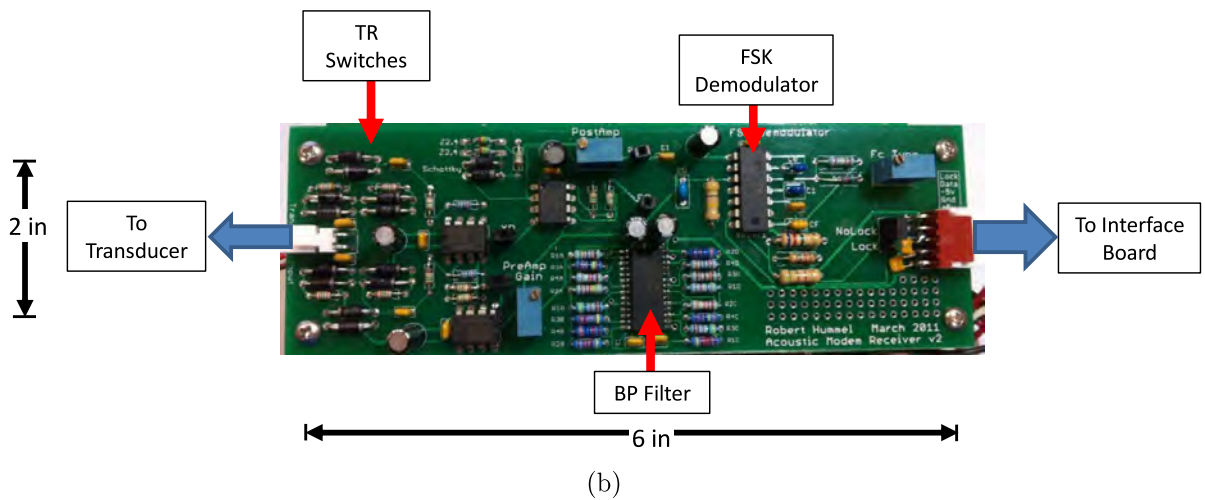
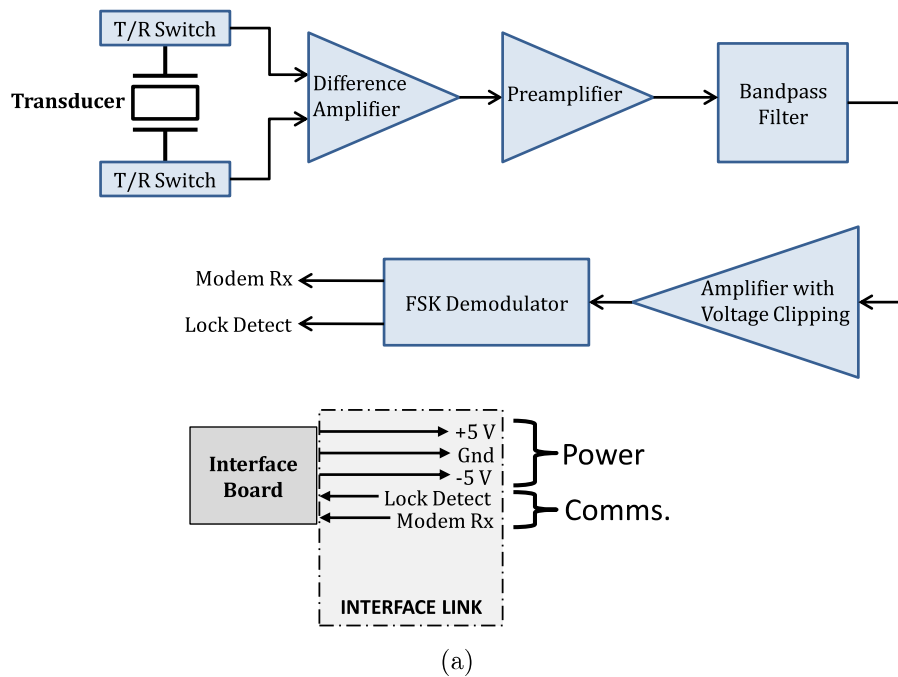


Figure 3-11: Acoustic modem receiver schematic (a) and printed circuit board (b). The acoustic signal is received with the transducer and fed through transmit/receive switches that protect the circuit from the high voltages produced during transmission. The signal is amplified, filtered, and amplified again. The signal then enters the FSK demodulator which produces a digital signal depending on the frequency of the input. This digital data is sent to the interface board for processing.




Figure 3-12: Schematic of transmit/receive switch. The AC-coupled diode bridge is connected to each of the transducer inputs on the receiver. The switch protects the receive circuitry from the transmission-level voltages would cause damage. The diodes do not block low voltages such as those present when receiving an acoustic signal.

signal passes through unimpeded. In practice, T/R switches are also implemented with MOSFETs in place of diodes, the advantage being that no bias current is required, so steady-state power consumption is lower, however the MOSFETs require external and precise control, and the automatic operation of the diode protection circuit was seen as preferable in our application.

3.6.2 Pre-amplification

After the T/R switch, the next stage of the receiver consists of an operational amplifier (op-amp) circuit which amplifies the voltage difference between the outputs of the two T/R switches. This amplifier is arranged in the common difference amplifier configuration [54], with a gain of 10. The purpose of amplifying the difference between the two signals is to combine them into a single sinusoidal signal with zero mean. This signal then goes to a pre-amplifier with adjustable gain, implemented using another op-amp in the standard inverting amplifier configuration [54]. The gain is adjusted between 10 and 660 with a potentiometer on the feedback resistor. This stage must provide large amplifications without introducing substantial noise, therefore component selection, while important in every part of the receiver, is of utmost importance in the pre-amplifier.

3.6.3 Op-amp Selection

Our objectives in amplifying the signal are to (a) provide sufficient amplification to a signal which may be only a few millivolts in magnitude (b) operate at 50 kilohertz (c) introduce minimal noise to the signal. To meet these requirements, we need an op-amp that can provide amplification of at least 500 V/V at 50 kHz; the gain-bandwidth product (GBW) of the op-amp must be at least 20 megahertz. Also, the output must swing through its full range at 50 kHz, about 10 volts peak-to-peak in 10 microseconds, requiring a slew rate of at least 1 volt per microsecond (V/ μ s). Within these constraints, we must also look for an op-amp with the proper supply voltages, low offset voltage, and low cost. To meet these requirements, we selected the OPA228 manufactured by Texas Instruments [74], which features a slew rate of 10 V/ μ s, and GBW of 33 MHz, and noise of 3 nV/ $\sqrt{\text{Hz}}$, available from Digikey for about two dollars each. We utilize three of these op-amps in each receiver.

3.6.4 Filtering

After the pre-amplifier stage of the receiver, the signal passes through a band-pass filter. We perform the filtering using a MAX274 8th-order, Continuous-time, Active Filter integrated circuit from Maxim Integrated Products Inc. [49]. The chip consists of four cascaded 2nd-order op-amp filters with the performance of each filter set by external resistors (see Fig. 3-13). The values of the resistors can be determined analytically from the design equations in the datasheet or, more conveniently, with the provided design software [50]. Figure 3-14 shows the design process. The filter is designed as a Butterworth filter with a center frequency of 50 kHz, passband width of 10 kHz, stopband width of 20 kHz, maximum attenuation in the passband of 3 dB, and as much attenuation in the stopband as feasible. A Butterworth filter was chosen because of its flat frequency response within the passband [76]. Figure 3-15 shows the transducer input and bandpass filter output.

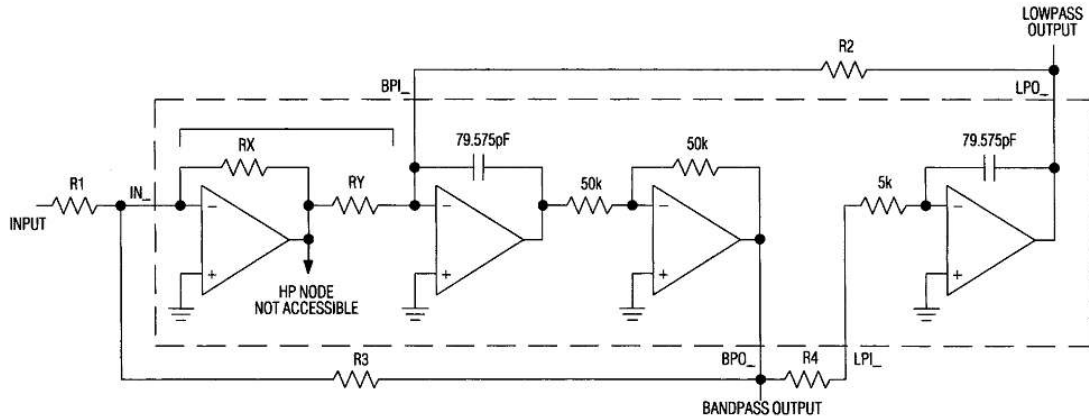


Figure 3-13: Schematic of a single stage of the bandpass filter. The filter is implemented using a MAX274 Filter IC [49]. The resistors R_1 , R_2 , R_3 , and R_4 determine the frequency response of the filter. In the MAX274, four of these filter stages are cascaded to achieve an 8th-order filter. Schematic from MAX274 datasheet.






Figure 3-14: Maxim filter design software [50]. The filter design software guides the selection of resistor values based on a desired frequency response.



(a)



(b)

Figure 3-15: Waveform of transducer input and bandpass filter output. The filter properly extracts the space (a) and mark (b) frequencies from the noisy input.




Figure 3-16: Schematic of filter output amplifier with voltage clamp. The amplifier is positioned after the filter output to amplify the signal before it enters the FSK demodulator. The voltage clamp consisting of two zener diodes and two signal diodes limits the voltage to protect the FSK demodulator IC. Image from [26].

3.6.5 Post-amplification and voltage limiting

The bandpass filter stage, while greatly reducing the signal magnitude at undesirable frequencies, also attenuates the signal within the passband. If we were to feed this signal directly into the FSK demodulator, the signal would be too small for consistent decoding. Therefore, we utilize a post-amplification circuit after the filter. The circuit is similar to that used in the pre-amplifier, with one addition. The FSK demodulator is sensitive to large signals; the datasheet suggests limiting the voltage entering the chip to $3 V_{\text{RMS}}$ to prevent overheating. The circuit shown in Fig. 3-16 amplifies the circuit with a user-configurable gain, while also limiting the output voltage to an acceptable level. The circuit utilizes zener diodes (D_{Z1} , D_{Z2}) and signal diodes (D_1 , D_2) to clamp the voltage at $|V_{\text{max}}| = V_{Zf} + V_{Zr} + V_{Df}$, where V_{Zf} is the zener forward voltage drop, V_{Zr} is the zener reverse breakdown voltage, and V_{Df} is the signal diode forward voltage drop [26]. Our selection of zener and signal diodes was limited by our desire for through-hole components, however we managed to obtain parts resulting in $|V_{\text{max}}| = 3.75 \text{ V}$, giving an RMS voltage of 2.6 V assuming a sinusoidal wave. In practice, the clamping resulted in a trapezoidal wave, which has a higher RMS value than a sinusoidal wave, however we observed experimentally that the output remained within the acceptable region in all test cases.

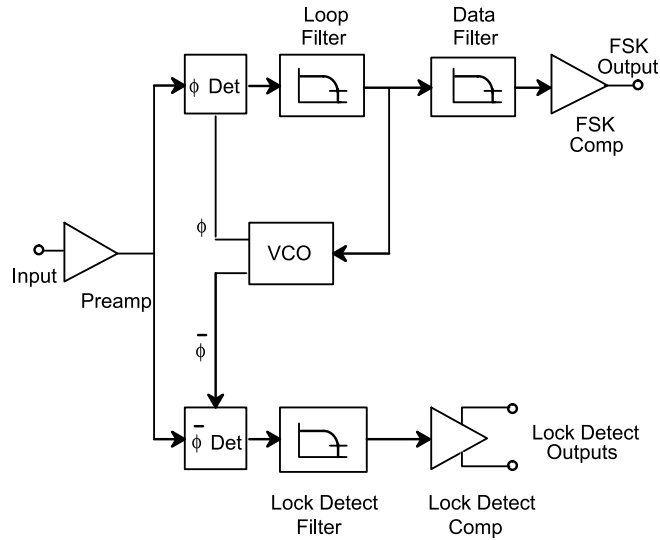


Figure 3-17: Functional block diagram of FSK decoding with XR2211 [23]. The voltage-controlled oscillator (VCO) free-running frequency is configured to match the acoustic signal. The Lock Detect Filter determines if the input is ‘near’ that frequency. The Loop Filter and Data Filter determine if the input is higher or lower than the VCO frequency, i.e. whether the input is closer to the mark or space frequency. This binary decision is output as digital data at the FSK output. Diagram from XR2211 datasheet.

3.6.6 FSK Demodulation

The FSK demodulation stage involves taking a sine wave of either the mark or space frequency, encoded as in Section 3.4, and converting it back to a high or low signal compatible with a microcontroller. Many commercial and research acoustic modems use a sampling-based approach to determine the frequency of the signal. However, this method requires sampling well above the Nyquist-rate, in our case approximately 100 kHz, and performing a frequency extraction technique such as the FFT. This method is highly taxing on the onboard processor; usually an expensive and complicated digital signal processor (DSP) is used. As an alternative approach, we implemented our demodulator in hardware, based around the XR2211 FSK Demodulator/Tone Detector from Exar Corporation [23], a monolithic integrated circuit containing all of the components necessary for FSK decoding.

The operation of the chip is shown in the functional block diagram in Fig. 3-17. First, a pre-amplifier amplifies the signal to a constant high level. The signal then

branches into two phase detectors. The loop phase detector compares the phase of the signal to that of a voltage-controlled oscillator (VCO) whose free-running frequency is set by the user with external discrete components as determined by provided design equations. The output of the loop phase detector is fed back into the VCO to allow the VCO to follow the input frequency. Additionally, the output is filtered into a DC level which depends on the input frequency. A frequency higher than the VCO free-running frequency will result in a logical low output, while a lower frequency will result in a logical high. The quadrature phase detector determines whether the VCO is properly tracking the input signal, and outputs a logical high level when the input is within the proper range of the free-running frequency.

Utilizing these functions, the system is able to detect whether a proper acoustic signal is being received, in addition to decoding the signal from the mark and space frequencies into the proper logical high and low levels, thus completing the acoustic modem receiver circuit. The FSK output shown in Fig. 3-17 is connected to the Modem Rx line going to the interface board shown in the receiver schematic shown above in Fig. 3-11a. The lock detect output is also sent to the interface board.

3.6.7 Summary of Acoustic Modem Receiver

The purpose of the acoustic modem receiver is to decode an acoustic signal into the proper logic levels used for data transmission. The sensing element is a piezoelectric transducer, which converts mechanical strain from incident pressure waves into a voltage that has a linear relationship to the input. The receive circuitry is protected from the high voltage transmission signals with a transmit/receive switch. The voltage which develops at the transducer is first amplified with a low-noise, high-gain amplifier. Then it is passed through a bandpass filter, leaving only the desired frequency components. The bandpass filter output is amplified again, with care taken to limit the voltage for protection of the next stage of the circuit, the FSK demodulator. The FSK demodulator receives the amplified bandpass filter output and determines a) whether the received signal falls within the desired frequency range, i.e. whether the signal is part of an acoustic transmission or is only noise, and b) whether the

received signal is closer to the mark or the space frequency. The outputs of the FSK demodulator provide logic-level signals which can be read by a digital device such as a microcontroller.

3.7 Interface and Controller Board

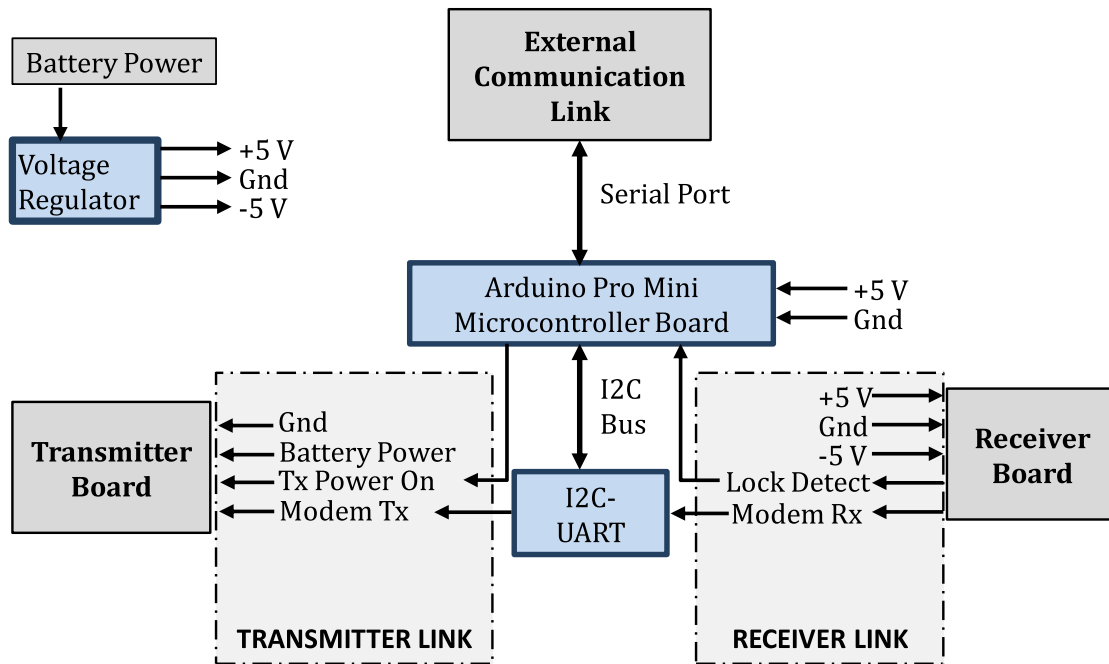
The interface board bridges the link between the user and the acoustic transmission and reception boards discussed above. Figure 3-18 shows a schematic and photograph of the board.

3.7.1 Power


The modem is powered with a DC supply of 9-18 volts, accommodating many of the most common battery types and sizes. The supply is regulated to ± 5 volts with an NDTD1205C DC/DC converter produced by Murata Power Solutions. These regulated voltages are used on the interface board and the receiver board. The transmitter board uses the unregulated battery supply which is routed from the interface board.

3.7.2 Universal Asynchronous Receiver / Transmitter (UART)

The digital signals that are converted to/from acoustic signals are controlled with an integrated circuit, the SC16IS750 single UART with I²C-bus/SPI interface and 64 bytes of transmit and receive FIFOs. We will refer to this device as the acoustic UART. The acoustic UART generates the digital signals that are sent to the transmitter to indicate which frequency should be transmitted. The output of this device is the FSK Input signal described in Section 3.5.1. The acoustic UART also reads in the digital signals generated by the FSK Demodulator on the receiver board, so the input is the FSK Output signal described in Section 3.6.6. The acoustic UART can be configured for various packet configurations and speeds (i.e. baud rate, bits per second, bps), and contains 64-byte input and output buffers to store data until it can be processed. It communicates with the Arduino using the I²C 2-wire serial bus.



(a) Interface Board Schematic



(b) Interface Board PCB

Figure 3-18: (a) Interface Board Schematic (b) Interface Board PCB
 The Interface Board controls the low-level functions of the modem and provides a communication interface for user interaction with the modem. The I²C-UART handles communication of digital signals to and from the transmitter and receiver boards. The Arduino is programmed with a user interface and control code for operating the I²C-UART to transmit and receive acoustic signals. The board also controls power regulation and distribution.

It should be noted that the microcontroller on the interface board also contains a UART port. The reasons to use a dedicated UART instead of that on the microcontroller are two-fold:

- Microcontrollers and computers are generally meant for much higher-speed operation than acoustic modems. Most microcontroller and computer UARTs cannot be easily programmed in the 10-100 bps range.
- A dedicated UART to communicate with the transmitter and receiver boards leaves the microcontroller UART available for external communication, as the serial port is the simplest way to connect to a host computer or microcontroller.

3.7.3 Control

The modem is controlled with an Arduino Pro Mini 5V/16MHz. This component contains an ATmega328 microcontroller providing digital input and output pins to communicate with and control the modem, as well as a serial port to communicate with the host device or user. The Arduino is programmed to send the appropriate commands to the acoustic UART, which generates and interprets the digital communications between the acoustic transmitter and receiver boards. The Arduino can also directly control the FSK Input and Output pins of the acoustic boards, which is used for calibration and debugging. The Arduino also controls the power state of the transmitter board; the power is turned off when no transmission is being made both to save energy and because the received signal will be overpowered when the transmitter is active.

In the main operating mode of the modems, the acoustic link behaves as a transparent serial link. When data is received on the Arduino serial port, the transmitter is turned on and the Arduino sends the data to the acoustic UART, which then sends it to the transmitter board. When acoustic data is received by the receiver board, it goes to the acoustic UART, which then sends it to the Arduino, where it is transmitted to the external serial port. Configuration of the acoustic UART baud rate and packet configuration (stop bits, parity bit, etc.) is performed by sending a

special control sequence to the Arduino to exit transparent mode. Additionally, for transmission, a buffer and timer are used to allow several bytes to accumulate before the transmitter is powered on and the data is sent.

3.8 Submersible Enclosure

The submersible enclosure for the modem is shown in Fig. 3-19. It is 16 inches long and 4 inches in diameter. Fabrication of the end caps and PCB mounting blocks was outsourced for about forty dollars per modem. The transducer cage, O-rings, PCB rails, eye bolts, and other parts are purchased for under thirty dollars. Total cost of the enclosure, including assembly, is estimated at one hundred dollars.

3.9 Modem Parts List

Table 3.2 shows a list of the components used in the modem and the cost of each. The total cost is less than three hundred dollars, including the submersible enclosure. Further cost reductions could be achieved by using surface mount components, which would (a) reduce the cost of the discrete components, (b) reduce the area and therefore cost of the printed circuit boards, and (c) reduce the volume of the device thereby lowering the cost of the enclosure.

3.10 Results

Success rates for testing the modem are shown in Fig. 3-20. For each of the tests, we sent randomized sequences of 8-bit characters using the serial protocol of one start, followed by the eight data bits, and one or two stop bits. Our results come from tests of 500-character sequences. As shown by our results, using two stop bits tends to reduce errors, however it also reduces the *effective* transmission rate. With one stop bit, 80 percent of the transmitted bits contain data, while with two stop bits, only 72 percent of bits contain data. The rates presented here are the actual bit rates

(corresponding directly to the duration of a single bit), not the effective rate which depends on the protocol.

In the first test, we submerged the two modems together in a small bucket. With one stop bit, rates up to 200 bits per second (bps) are achieved with 100 percent accuracy (effective data rate of 160 bps). Using a second stop bit increases the







Figure 3-19: Photograph of modem assembly. The modem is enclosed in a polycarbonate tube sealed at both ends with O-rings fitted in the end caps. The modem circuit boards stack on top of one another with nylon spacers in between them. The stack includes blocks which fit on two threaded rods attached to the bottom end cap. The transducer wiring is fed through a hole in the bottom end cap and sealed with silicone. A waterproof connector passes through the top end cap to provide power and connect to the external serial port. Eye bolts are screwed partway into the top end cap to provide an attachment point from which to hang the modem.



(a) Bucket Test - 0.1 meters



(b) Tank Test - 1 meter



(c) Tank Test - 5 meters

Figure 3-20: Modem testing results. Rates up to 300 bps are achieved with zero errors when both modems are submerged together in a bucket (a) or in a tank separated by one meter (b). With the modems in the tank separated by five meters (c), rates up to 100 bps are achieved with perfect accuracy. For the short-range tests, using two stop bits gives improved performance. In the longer-range test, the number of stop bits has little effect.

Table 3.2: Modem parts list showing items and associated cost

Hardware	USD	Interface Board	USD
Enclosure	100	Voltage regulator	30
Transducer	10	Microcontroller	18
Potting	5	I2C-UART	15
		PCB	10
Transmitter		Receiver	
Hi-power op amp	5	Diodes	6
Frequency generator	3	Op amps	10
Boost regulator IC	2	Bandpass filter	8
Capacitors	10	FSK demodulator	3
Inductors	6	Resistors	3
Resistors	1	PCB	10
Relay	2		
PCB	10		
		Total	267

maximum rate to 300 bps (effective data rate of 218 bps). Faster speeds resulted in less than ten percent success rates.

The second test was performed in the tank shown in Fig. 2-1. The modems were hung from the ceiling and suspended with their transducers halfway between the water surface and the floor. The modems were an equal distance from the wall, about 50 centimeters, and separated from each other by one meter. At this distance, we achieved nearly identical results to the bucket test. With two stop bits, the maximum rate is 300 bps, and with one stop bit the maximum rate is 200 bps.

In the third test we increased the separation to five meters, with the modems submerged in the tank as in the second test. Here, the maximum reliable rate is 100 bps, at which errorless transmission is achieved. At 150 bps, greater than 20 percent errors occur, and using an additional stop bit does not offer substantial improvement. At 200 bps, more than 90 percent errors are recorded.

Several attempts were made to operate the modems off of a dock in the Charles River. These attempts were unsuccessful, and the solution to the problem remains unclear.

3.11 Conclusion

Improvements to the modem would first deal with the transmission power level—a higher power level would increase the maximum effective range of the system. Other improvements could include adding digital potentiometers to the resistors that control transmission and reception frequencies. The ability to change these frequencies in real time could improve the performance in difficult acoustic environments, such as those with high noise levels or multi-path.

We have developed here a low-cost acoustic modem capable of transmission rates up to 100 bits per second at several meters distance, and 300 bps at shorter range. Future work to improve the transmission system could increase the maximum effective range of the modem. The modem can be fabricated for under 300 USD, making it particularly well-suited for acoustic network applications requiring a low cost per node.

Chapter 4

Compressive Sensing with Mobile Agents

4.1 Introduction

This chapter considers mission design strategies for mobile robots whose task is to perform spatial sampling of a static environmental field, in the framework of compressive sensing (CS). The application of this theory has the potential to dramatically reduce the number of measurements required for high-fidelity field reconstruction. Here we investigate sampling strategies for CS with numerical simulations and field experiments with the vehicles described in Chapter 2.

According to compressive sensing theory, we can reconstruct fields with much fewer measurements than the convention based on the Shannon Sampling Theorem [61], which says that a signal can be fully reconstructed if sampled at more than twice the frequency of the highest-frequency component of the signal, the well-known Nyquist Rate. Specifically, CS methods allow reconstruction with the number of measurements scaling with the logarithm of the highest frequency component, provided the signal is compressible and the measurements satisfy the proper conditions. Random measurements are shown to satisfy these conditions (with high probability), however a random sampling strategy may not be optimal when sampling with mobile agents, most directly from the high cost of movement on the limited energy budget of a ve-

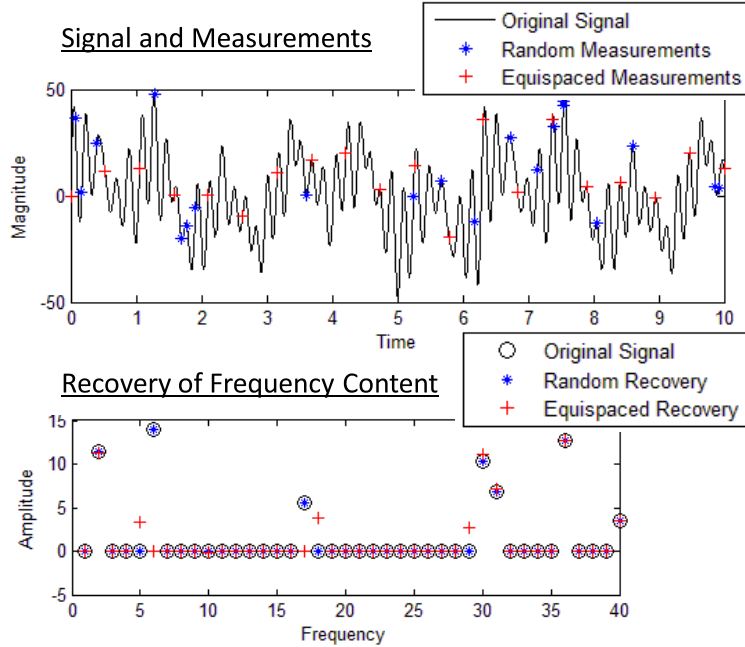


Figure 4-1: Example of compressive sensing (CS) in one dimension. The signal, shown in the top plot, is the sum of seven sinusoidal signals, with the maximum frequency equal to 40 Hz. Twenty measurements are taken within ten seconds, either randomly chosen or equi-spaced. CS reconstruction assumes the signal is sparse, that is, many of the frequency components are zero. The CS reconstruction of the frequency coefficients, shown in the bottom plot, is exact, while that from the equi-spaced measurements is incorrect. Without exploiting sparsity, the signal would need to be measured above the Nyquist rate of 80 Hz.

hicle. Figure 4-1 presents a one-dimensional example showing the potential benefits of compressive sensing.

A nonzero cost of transit between samples would argue for a traveling salesman path visiting a fixed set of random points, but there is also a practical desire to be robust against premature termination of the CS mission. This argues instead for suitably-sized random steps taken incrementally across the space. Conventional regular paths, such as the cellular lawnmower in two-space [9] are analogous to the TSP in the sense that premature termination is likely to destroy any chances of a meaningful reconstruction from the data in hand. We note that although samples can often be taken in transit, from the CS point of view these are by definition spatially correlated with the endpoints, and therefore in some sense redundant and less valuable than the vertices. It therefore remains an open question how to best design missions

for CS reconstruction.

Our contributions are as follows. We explicitly combine the constraints of energy costs with compressed sensing for the development of practical path planning. We further provide an optimized method for designing the randomized path which both satisfies the compressed sensing measurement requirements and explicitly includes navigation costs. We also introduce a sampling-based path planning technique that can incorporate kinematic constraints such as obstacles. We then present results from numerical simulations and experiments with the vehicles described in Chapter 2.

This chapter is organized as follows. Section 4.2 presents the CS background relevant to our work. Section 4.3 gives an overview of the related work in the CS field. Section 4.4 describes energy-constrained sampling strategies with numerical simulations presented for evaluation. Section 4.5 introduces a kinematically-constrained greedy algorithm to minimize the coherence between the field and the measurements, which improves the suitability of the selected points for CS reconstruction. We again provide numerical simulation results for evaluation. Section 4.6 provides experimental results involving environmental sampling and CS reconstruction using the vehicles described in Chapter 2. Section 4.7 concludes the chapter.

4.2 Compressive Sensing Theory

We present some key results from compressive sensing [3, 4, 6, 18], and lay out our particular scenario.

Let $\Psi = [\psi_1, \psi_2, \dots, \psi_n] \in \mathbb{R}^{n \times n}$ be an orthonormal basis in which a signal, y , has a concise or *sparse* representation. Suppose $y = \Psi x$ for $x \in \mathbb{R}^n$. Formally, x is *S-sparse* if it has at most S non-zero elements. Further, x is *compressible* if for x_S obtained from x by keeping only the S largest terms x^i , the loss in signal energy, $\|x - x_S\|$ is “small”. Many natural phenomenon are compressible in domains such as frequency, wavelets or total variation [12]. Let $\Phi = [\phi_1, \phi_2, \dots, \phi_n]$ be another orthonormal basis of \mathbb{R}^n and let z^1, z^2, \dots, z^m be the measurements obtained by projecting y on vectors $\phi_1, \phi_2, \dots, \phi_m \in \mathbb{R}^n$ i.e. $z^i = \langle y, \phi_i \rangle$. In the usual robot sampling scenario, sensing

is performed in the spatial domain, so that the measurement basis functions are delta functions in space. The domain grid is vectorized into a $n \times 1$ column, and each measurement is an $n \times 1$ column of zeros, with a single 1 at the measurement location. Φ is orthogonal, a permutation of the $n \times n$ identity matrix. We are interested in cases where the number of samples is significantly less than the dimension of the field i.e. $m < n$; this is encoded into an $m \times n$ matrix, R , which pre-multiplies Φ to select which points to sample.

For an S -sparse field, the reconstruction problem can now be formulated as a convex optimization program

$$\min_{x \in \mathbb{R}^n} \|x\|_{\ell_1} \text{ subject to } z = R\Phi\Psi x$$

This optimization finds the sparsest signal, x , which is consistent with the measurements, z . A key theorem is [6]:

Theorem 4.2.1 *Fix $y \in \mathbb{R}^n$ and suppose that the coefficient sequence x of y in the basis Ψ is S -sparse. Select m measurements in the Φ domain uniformly at random. Then if $m \geq C \cdot \mu^2(\Phi, \Psi)S \cdot \log n$ for some positive constant C , the reconstruction of f is exact with high probability.*

This implies that if the incoherence is constant, then the number of measurements needed for exact reconstruction is $O(\log n)$. Note that in the best case scenario, if we knew *which* elements of x are non-zero, we would require only S measurements and in comparison, the oversampling here is $O(\log n)$. However, in general, knowledge about non-zero elements is not available, and a naive method would make n measurements.

4.3 Related Work

CS algorithms are generally applied with a class of random projection matrices constructed using Gaussian and Bernoulli variables that are known to satisfy the restricted isometry property and therefore lead to efficient reconstruction [6, 18]. Such random projections are not suited for our problem of sampling using robots because

the motion constraints make gathering of independent random measurements infeasible. Therefore it is desirable to find new projection matrices that 1) result in efficient CS reconstruction and 2) can be constructed efficiently by mobile robots. However this turns out to be a challenging problem and in general, it is not easy to show that structured projections satisfy the RIP. There have been recent efforts directed towards the design of purely deterministic projection matrices. A technique based on hash functions to explicitly construct a binary sensing matrix that has provably good reconstruction properties is provided in [31]. Similarly, finite field algebra is employed to design deterministic projections that satisfy the RIP [17]. Finally, a small number of chirp signals to form the projection matrix is employed in [1] and empirically shown to satisfy the RIP. The above approaches directly control the structure of the sensing matrix $A = R\Phi\Psi$ whereas in our case, we can only control the measurement matrix $R\Phi$. To the best of our knowledge, there are no techniques for structured random construction of $R\Phi$, especially given the motion constraints of robots.

We are aware of a few works treating the design of optimal projection matrices for a given dictionary. Elad [20] addressed the problem by minimizing mutual coherence of the effective dictionary A . The method iteratively reduces large elements in the Gram matrix $A^T A$ directly, and then computes a new projection satisfying rank conditions associated with the desired number of measurements. Duarte-Carvajalino and Sapiro [19] addressed the same objective shortly thereafter as an eigenstructure problem, and achieved faster computations; Xu *et al.* [79] followed with an approach based on the equiangular tight frame, and reported further improvements, although computing times are not given. These methods put no constraints on $R\Phi$ and the outcome of optimization is dense in the general case, because implicit is the assumption that all of y is available. We also note that very specific constraints in measurement matrix design have been developed in the MIMO radar community; here the task is to choose transmission waveforms and array parameters so as to create incoherence in the sensing matrix [7, 71, 81].

4.4 Energy-constrained Compressive Sensing

This section will investigate several mission planning strategies that explicitly incorporate energy costs and constraints. The goal of these missions is to achieve the best possible field reconstruction from under-sampled data. We will first lay out the energy cost function and then introduce several novel mission planning techniques and discuss their inherent benefits. Numerical simulation results are provided for a variety of mission scenarios. This work was presented at ICRA 2011 in [29].

4.4.1 Energy Budget and Costs

Energy-constrained compressive sensing allows us to incorporate vehicle- and mission-specific parameters into the selection of measurement locations. Our energy cost function includes costs for forward motion, direction changes, and, and taking measurements. The equation is given by

$$E = \alpha \sum_{i=1}^m L_i + \beta \sum_{i=1}^m (\Delta\theta_i)^2 + \gamma m$$

where E is the expended vehicle energy, α is the forward motion cost, β is the direction change cost, γ is the measurement cost, L_i is the length of step i , $\Delta\theta_i$ is the magnitude of the direction change before step i , and m is the number of measurements. We forego a separate cost for communication; it can be assumed that communication cost is either (a) a constant hotel load and can be included in movement cost, or (b) the communication of data and can be included with the measurement cost.

Several scenarios were investigated. One scenario is that of expensive movement and cheap measurement, such as when measuring a temperature field. Another scenario is that of measurement costs on a par with movement costs, such as with biological sampling [60] or measurements that require active station-keeping for an extended period. Also, direction change costs were varied to simulate differences in vehicle dynamics and domain size; sharp turns with streamlined vehicles will greatly reduce speed, thereby increasing the amount of energy consumed. These scenarios

are discussed in more detail in Section 4.4.5.

The initial vehicle energy was based on the size of the domain. The energy was set to allow for a fixed number ($O(10)$) of traversals of the space, along with the associated 90° turns required for a complete lawnmower pattern, if no measurements were taken. The mission design must then trade off movement for measurement.

4.4.2 Mobility Models

Random Walk

The random walk mobility model is designed to approximate the compressive sensing desire for uniform random measurement locations while, at the same time, enforcing vehicle kinematic constraints. The vehicle begins in a randomly chosen location and a measurement is taken. After each measurement i , a joint probability mass function (PMF) $P(\Delta\theta, L)$ is sampled for a direction change $\Delta\theta_i$ and step length L_i . The vehicle follows this trajectory and takes a measurement at the end of the step. It continues in this manner, successively sampling the PMF, moving to the desired location, and taking a measurement. Fig. 4-2a shows a typical random walk tour. The step-by-step model is illustrated in Fig. 4-3. We assume instantaneous direction changes; a fair assumption for vehicles that are slow compared to the size of the domain. The result is a continuous path with discontinuities in heading. Domain boundary constraints are enforced by projecting the desired path back into the domain, with the angle of reflection equal to the angle of incidence. McNish [53] investigates reflection angles that achieve probabilistically-uniform coverage; however, that work applies only to circular domains with direction changes at the boundaries, with continuous sampling along the paths. Further work in reflection strategies may be warranted.

The PMF is optimized as described in Section 4.4.4. To reduce computational effort, the PMF contains only the magnitude of $\Delta\theta$; a positive or negative direction is chosen randomly with equal probability. An important feature is that no time or measurement horizon is required for the random walk. It simply continues until the

available energy is expended or some other mission parameter is satisfied. The use of an optimized PMF to generate the random walk is one of the major contributions of this work.

Random TSP

The random TSP mobility model is designed to use uniform random measurements that, with high probability, fully satisfy the CS incoherence requirement. A specific number, m , of measurement locations are randomly chosen from the domain. A tour of these points is then generated using a TSP approximation. The expected energy consumption is computed, and if it is "close" to the available energy, the path is used. Otherwise m is modified to better match the available energy, and another set of points and a tour are generated. For our comparisons, we use a genetic algorithm with nearest neighbor initialization to compute the TSP tour. Fig. 4-2b shows a typical random TSP tour.

Randomized Lawnmower

The randomized lawnmower mobility model is comprised of back-and-forth motions traversing the space. The deterministic lawnmower is a widely used nonadaptive measurement strategy for coverage of a space [10]; here we introduce randomness for the purposes of employing compressive sensing. Without this randomness, we would be subject to aliasing and the Nyquist-rate limit. The path is determined based on the the desired number of measurements, with the location of the lengthwise traversals chosen randomly. The measurement locations are then chosen randomly along this path. Additionally, the ratio of energy spent on movement and measurement is configured to match that of the random walk determine the number of traversals and measurements. Fig. 4-2c shows a typical random lawnmower path.

Mobility Models - Algorithm Complexity

Complexity of mission planning algorithms must be taken into account when the number of vehicles or measurements is large. The random walk and random lawnmower

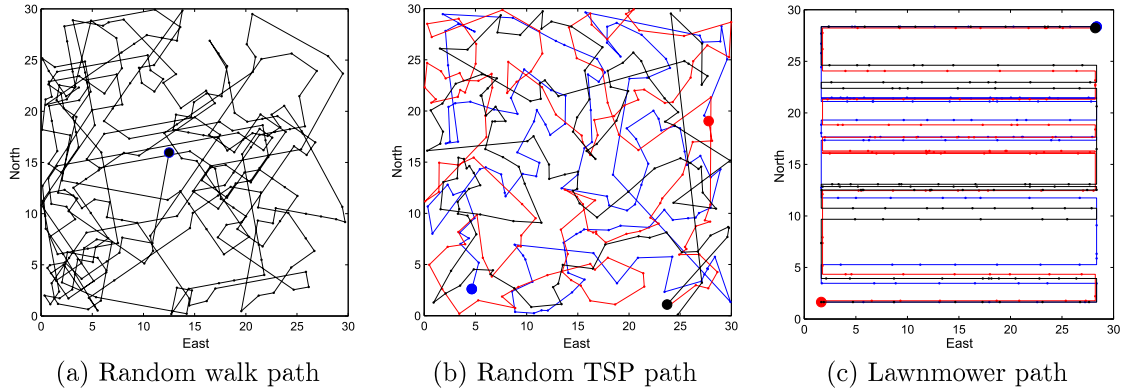


Figure 4-2: Typical robot paths. The random walk path begins at a random location and samples a joint PMF for heading change and distance after each measurement. The random TSP path is designed in three phases, each using some fraction of the available vehicle energy. A nearest neighbor tour connects randomly chosen points for each phase. The lawnmower path is designed by choosing the number of traversals and measurements to result in a desired ratio of measurement and movement energy expenditure. The traversal and measurement locations are then chosen randomly. The lawnmower path is also designed as a set of independent phases.




Figure 4-3: Illustration of the random walk. At each measurement location, a new step length L_i and a heading change $\Delta\theta_i$ are chosen from a joint PMF.

calculations have a cost which is linear in m . For the random TSP, the approximation we use here is also linear in m ; however, there is no approximation factor for this heuristic solution. Solutions with performance guarantees include a minimum spanning tree that gives an approximation factor of two with $O(m^2)$ computational cost, or a matching step that gives an approximation factor of $3/2$ with $O(m^3)$ cost [11].

4.4.3 Horizon Planning

Generating a pre-planned mission relies on predictions of vehicle energy expenditure. Stochastic variables such as ocean currents and moving or unknown obstacles can

cause unexpected changes in energy expenditure. Both energy surpluses and deficits result in inefficient use of resources. These inefficiencies include wasted support vessel time, personnel expenditures, and lost opportunity for the robot to perform other tasks. The unplanned random walk model does not suffer from this problem; its one-step horizon enables the use of all available energy. Differences in expected and actual energy expenditure, however, do affect the random TSP and lawnmower planning strategies. A common practice is to split a mission into a set of phases, each with a shorter horizon, as a hedge against unexpected energy consumption. A TSP path planned in this manner will suffer from a lower efficiency, as there are fewer points to connect with each phase of the mission, but will be more likely to cover the entire domain at least once or twice. With this in mind, we design the random TSP and lawnmower missions as a series of phases that combine into the complete path. In the specific example below we use three phases which each consume 33% of the vehicle energy. An alternative to this approach could be to plan a mission with, say, 80% of the available energy. In either case, the online-planning aspect of the random walk provides for a more complete and reliable use of mission resources.

4.4.4 Random Walk Optimization

To design the PMF for the random walk, a particle swarm optimization (PSO) algorithm is used. The PSO is performed with the method given in [30], which is a comprehensive learning variant [42]. PSO was used because it performs well with multi-modal functions where local optima may exist. The PSO is performed as follows: each position in the joint PMF is represented by several independent ‘particles.’ Each particle moves throughout the design space in response to the current best evaluation of the fitness objective function, influenced also by random noise. Our objective function is the reconstruction error achieved with a random walk tour using the PMF. Movement and measurement costs will affect this performance, so a separate optimization was run for each set of costs. Results are shown in Fig. 4-4 for the case of expensive and cheap movement. In the case of expensive movement 4-4a the optimization returns a PMF which prefers shorter travel distances. In contrast,

with cheap movement, shown in Fig. 4-4b the optimized PMF prefers longer travel distances.

4.4.5 Simulation Results


We have tested the proposed sampling strategies through extensive simulations in MATLAB. The procedure is as follows. The domain is a 2D grid with n grid points. The field is taken from measurements of sea surface temperature. Each of the sampling strategies mentioned above was compared under three different scenarios of vehicle movement and measurement costs, as described in 4.4.2. Each of the scenarios is meant to simulate differences in real-world operating conditions, such as vehicle mobility, type of measurement, and domain size. The random walk PMF optimization was performed for each scenario.

Scenario 1: Inexpensive Measurements and Large Domain


For this scenario, the forward motion cost, α , is set to one, and the direction change cost, β , and measurement cost, γ , are set to 0.01. These settings simulate a vehicle which is taking cheap measurements, such as temperature and salinity, or simple optical imaging (real-time transmission of images would likely be considered a high measurement cost). Additionally, the small turning cost simulates a domain that is large in comparison to the vehicle turning radius, or the vehicle is highly maneuverable. Fig. 4-5 (left) shows the simulation results for Scenario 1. The top plot shows energy expended as a function of measurement number. The bottom plot shows reconstruction error. For this plot, reconstruction error is computed successively during the mission as measurements are added.

Scenario 2: Inexpensive Measurements and Small Domain

For this scenario, α and β are set to one, and γ is set to 0.01. Here, direction changes play a larger role, and we refer to this as the small domain effect; it comes from having a less maneuverable vehicle or a domain which is small in comparison to the



(a) Expensive movement (cheap measurement)



(b) Cheap movement (expensive measurement)

Figure 4-4: Optimized PMF for the case of expensive (a) and cheap (b) movement relative to measurement cost. Sampling the PMF gives a non-dimensional length of travel L/L_0 , where L is the distance to move and L_0 is the largest dimension of the field, and a direction change $\Delta\theta$. In the case of expensive movement the optimization returns a PMF which prefers shorter travel distances. In contrast, with cheap measurement the optimized PMF prefers longer travel distances. For clarity, only two magnitudes of direction change are allowed, while in the full simulation a wider range of values was included in the optimization.

vehicle turning radius. The small domain imposes a larger turning cost because large direction changes will significantly slow the vehicle. High turning cost can also be embodied as the expense of depth or height adjustments when sampling a vertical domain. Additionally, a small domain will impose more severe limitations on the vehicle path that may decrease the efficiency of path planning algorithms. Fig. 4-5 (center) shows the simulation results for Scenario 2.

Scenario 3: Expensive Measurements and Small Domain

For this scenario, α , β and γ are set to one. These settings simulate a vehicle taking expensive measurements, such as biological samples in a small domain. Fig. 4-5 (right) shows the simulation results for Scenario 3.

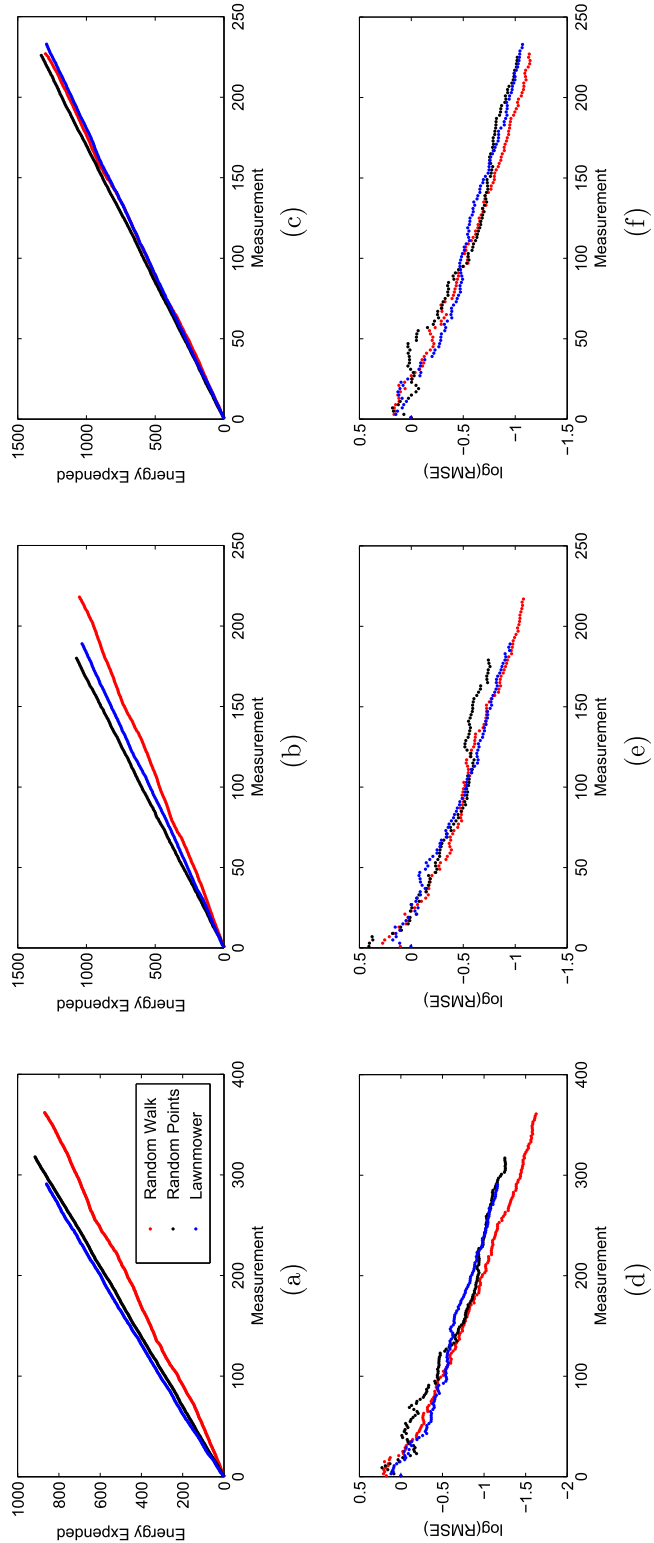


Figure 4-5: Results for Scenario 1 (a) and (d), Scenario 2 (b) and (e), and Scenario 3 (c) and (f). Plots (a)-(c) show energy expenditure as a function of the number of measurements taken. Plots (d)-(f) show the logRMSE as a function of the number of measurements. Scenario 1 - Inexpensive Measurements and Large Domain. Scenario 2 - Inexpensive Measurements and Small Domain. Scenario 3 - Expensive Measurements and Small Domain. In Scenario 1, the random TSP is more efficient, taking more measurements and also achieving a better reconstruction error. In Scenario 2, the random TSP is again more efficient and achieves a better final reconstruction error; however, for a given number of measurements, the random walk performs equally well. In Scenario 3, all strategies are similarly efficient, and here the random TSP is superior given the same number of measurements. The lawnmower performs similarly to the random walk in Scenarios 1 and 3. In Scenario 2 the random walk is superior to the lawnmower. *Legend: Blue-random walk; Black-lawnmower; Red-random TSP*

4.4.6 Discussion

Comparing results in each scenario shows that for a given number of measurements the random TSP strategy marginally outperforms the lawnmower and random walk. Likely, this performance difference can be attributed to the uniformly random measurements taken by the random TSP, as opposed to the quasi-random measurements of the other strategies. For the random walk, this may include the corner clustering evident in 4-2a, which is a result of our reflection model.


A larger influence on final reconstruction error seems to come from differences in the number of measurements taken with each mission type. The lawnmower mission was designed to match the random walk's ratio of movement energy to measurement energy, so a similar number of measurements is taken with the two strategies, thus resulting in a similar reconstruction error. The random TSP strategy, however, becomes more efficient as more measurements are taken. Fig. 4-6 shows the number of measurements taken as a function of energy expenditure for each of the measurement and movement cost scenarios for the random walk and the random TSP strategies, within a single planning phase. As expected, the random TSP mission has an increasing slope, becoming more efficient with more energy. At low energies, however, the random walk outperforms the random TSP. Thus, the decision of which strategy to use depends on the mission horizon – short missions would do best with a random walk, while long missions do best with a tour of random points.

As an example of a mission design decision, consider one vehicle operating under Scenario 1, that is, cheap measurements and expensive movement, with minimal direction change costs. Assume that a pre-planned mission must consist of at least four distinct phases, for reasons discussed in Section 4.4.3. If one desires to measure roughly 25% of the domain with $n = 900$ (in the case of Fig. 4-6, domain size is 30×30), about 200 measurements are required. Each of the four phases would then require 50 measurements. Looking at Fig. 4-6a, we can see that the random walk mission can achieve this number of measurements at lower cost than the random TSP mission. Conversely, if one desires a mission with phases consisting of greater


than 100 measurements each, the random TSP mission is more efficient. In the example above, 100 measurements will allow roughly four traversals of the space, a reasonable number from a field reconstruction point of view. Specific vehicle and field parameters will of course affect this tradeoff, so Fig. 4-6 must be recreated for each design scenario.

A competing design parameter is complexity of the planning algorithm. While random points will likely outperform the random walk at large numbers of measurements, the TSP approximation becomes increasingly expensive with more measurements. On the other hand, the random walk has no horizon. Referring again to Fig. 4-6, we can see that for some scenarios (namely, Scenario 3), the random TSP only marginally outperforms the random walk. In this situation the increased complexity may make the random points strategy undesirable, so that the actual crossover point is further right on the graph - i.e., at higher energy levels.


We have investigated three mission design strategies for compressive sensing of spatial environmental features using mobile robots with energy constraints. Our random walk strategy includes a novel addition to classical compressed sensing: we optimize the probability mass function of the walk in order to consider motion constraints. The random walk strategy is lightweight and scalable with regards to path planning and is efficient at low numbers of measurements. The random TSP strategy outperforms the random walk at larger numbers of measurements, but an additional computational cost must be incurred to design a TSP tour. The lawnmower strategy performs similarly to the random walk and is more computationally manageable than the random TSP method; however, it is still more expensive than the random walk. We establish that these measurement strategies combined with compressive sensing provide an efficient way to reconstruct compressible natural fields.



(a) Scenario 1: Inexpensive measurements and large domain



(b) Scenario 2: Inexpensive measurements and small domain



(c) Scenario 3: Expensive measurements and small domain

Figure 4-6: Plot of measurements and energy budget. The number of measurements taken versus initial vehicle energy is shown for random walk (+) and random points with TSP approximation (•). More measurements will likely result in better reconstruction, so, neglecting computational differences, a designer would want to choose the sampling strategy that maximizes the number of measurements possible for the energy available in the vehicle.

4.5 One-step-ahead Kinematic Compressive Sensing

A large portion of work on compressive sampling and sensing has focused on reconstructions from a given measurement set. When the individual samples are expensive and optional, as is the case with autonomous agents operating in a physical domain and under specific energy limits, the CS problem takes on a new aspect because the projection is column-sparse, and the number of samples is not necessarily large. As a result, the underlying incoherence properties in l_0 reconstruction, can motivate the purposeful design of samples in planning for CS with one or more agents. We develop here a greedy and computationally tractable sampling rule that will improve errors relative to random points. Several example cases illustrate that the approach is effective and robust.

Our goal in this section is to describe and develop this kinematic compressive sensing problem. In the following section we state it more fully and discuss how operational constraints affect compressive sensing. We propose a simple greedy algorithm - called RIPD - for ranking sample locations in a way that balances computational and operational constraints with reconstruction error. Modest but consistent improvements are demonstrated with discrete cosine and Haar wavelet transforms. This work will be presented at GlobeCom in 2011 [28].

4.5.1 Motivation

In the canonical case of CS, an under-determined system $Ax = z$ is solved while minimizing $\|x\|_0$, in many cases equivalent to minimizing $\|x\|_1$, for which fast algorithms exist. Here A is the effective dictionary, x the vector of modal coefficients for the dictionary, and z the set of m CS measurements. Denoting N as the number of degrees of freedom in the system, if x is known to have $S < N$ components that are nonzero, the signal is said to be strictly S -sparse, and exact reconstruction is possible with a small number of incoherent measurements.

CS methods carry the general recommendation of random samples, for the purpose of reducing their coherence, in probability, with respect to the reconstruction basis. We can write $A = P\Psi$, where Ψ is the underlying, possibly redundant dictionary matrix and $P = R\Phi$ is a projection matrix that is randomly chosen, or could be designed (above in Section 4.4 we compare three methods for choosing R). We thus differentiate between $z = Py$, the CS measurements, and $y = \Psi x$, the set of physical measurements that are possible. The physical measurements are the complete set of pixel values in an image, for instance, whereas the CS measurements are a smaller set of their linear combinations, as encoded in P . Thus if Ψ is complete (square), then P is $m \times N$, with $N \geq m$.

When some elements of y are not used or available, the projection P has $N - m$ zero columns. This is a significant specification, because in contrast with the unconstrained versions above, the problem is now combinatorial and likely to be exponential in cost. Column-sparsity of P is the major consequence when point data is collected by mobile agents in a physical domain, and projections obtained by the means above remain an unhelpful relaxation. Additional factors can come into play. Motions of the agent may be subject to speed, attitude rate, energy, communication, and collision constraints. If the vehicle moves too slowly, the field may change. Thus we see that the collection of such data for compressive sensing is a rich planning problem; the insertion of transit costs into CS is only the first of many possible aspects.

As shown above in Sec. 4.4, assuming that the mission length allows multiple traversals of the space, the leading factor in reconstruction performance is simply the number of points. A key question then becomes the planning horizon, because both random-point TSP and randomized lawnmowing approaches are efficient only with long horizons, whereas a random walk has no particular horizon, beyond the implicit assumption that the agent crosses the space adequately. More specifically, we distinguish among three main paradigms for collection of data:

- A suitably long planning horizon such that $m \geq 4S$ is guaranteed. If a full set of points is chosen and they are not located trivially, they can be visited using a traveling salesman path, whose distance is approximately $0.75\sqrt{m}$ in the unit

square [68]. Spatial partitioning can be used for multiple agents, e.g. [33].

- A one-step ahead horizon, to generate the best CS product possible from points sampled so far, plus one additional.
 - Transit is cheap so step lengths are unconstrained.
 - Transit is expensive and so it is desired to collect samples frequently - this is accomplished by limiting the transit length.

A significant factor also is that the number of collected samples may not be large in kinematic CS, numbering perhaps in the hundreds or even tens for certain applications. This calls into question the suitability of fully random point selections, which are justified only in probability.

We present and demonstrate below a simple greedy method for actively minimizing coherence. As we show in examples, reconstruction errors appear to be systematically improved over random point selection, but two caveats are noted. The method incurs a significant computation load, although it is well within the capability of modern microprocessors during transit time for undersea and many other types of mobile agents. Second, the method requires that the reconstruction basis be specified at the time of sampling, although it can be changed on the fly with an added computation. We will not consider in this method learning of dictionaries, or non-square dictionaries.

4.5.2 Formulation

The specific condition for exact reconstruction in large-scale CS is that $m \geq C\mu^2(P, \Psi)S \log(N)$ [3]. Here, C is a constant, and $\mu(P, \Psi)$ captures the incoherence of the reconstruction basis P and the measurement basis Ψ . As noted above, most compressive sensing work to date advocates choosing samples randomly, because this almost surely minimizes μ . Related to incoherence, the bases are said to satisfy the restricted isometry property if each column of the effective dictionary A is nearly orthogonal to every other column. A scenario satisfying RIP will be

amenable to reconstruction with noise and when the signal is not strictly sparse, but still compressible.

For the purpose of describing our construction, let us say that i measurements have been taken already. Hence there is an $i \times N$ A -matrix, which we denote A_i , so that $A_{i,j}$ is the j 'th column of A_i . The last measurement is z_i , taken at location r_i . The decision is where to take the next measurement. The best we can do in a greedy algorithm is to choose the feasible measurement location r_{i+1} so as to minimize the resultant coherence of the set, which is

$$q(r_i, A_i, r_{i+1}) = \max_{j,k} |\langle A_{i+1,j}, A_{i+1,k} \rangle|, \\ \forall r_{i+1} \in R(p, r_1, r_2, \dots, r_i).$$

Here the function R represents the feasible set, given all the previous locations and parameters p . This encodes for example the constraints of no repeated points, limited stepsize, the domain boundaries, as well as others. For any non-trivial feasible region defined by R , the function q is complicated, however, with multiple local minima and discontinuities depending on the transform in use, and this leads to a sampling-based strategy. The idea is to select n_t trial locations $r_j^- \in R(r_i, p)$, $1 \leq j \leq n_t$, and set $r_{i+1} = r_j^-$ satisfying $\min_j q(r_i, A_i, r_j^-)$. This is the greedy method we are proposing, referred to as RIP design, or RIPD.

The ensemble of n_t candidates can be chosen by any method desired. For instance, if it is desired to cross the domain quickly, then one would put them on the boundaries of R , either deterministically or randomly. They can be placed uniformly random in the interior of R , as in our examples below. In the case of obstacles, feasible paths can be determined using a fast, and (typically) sub-optimal scheme such as the RRT [37]. Herein, we see that a new optimization problem arises that would weigh the long-term statistical improvement in reconstruction error (via greedily reducing coherence) against the very specific costs and risks of a complicated path through obstacles. This is outside our scope at present.

The above procedure of assessing n_t candidates and choosing one at each cycle

is constant-time and has cost that scales with $n_t N^2$. To see this, let Q_i be the symmetric $N \times N$ matrix whose diagonal is zero and whose off-diagonal terms are the inner products induced by A_i . This is the Gram matrix. Assessing Q_{i+1}^- for a candidate r^- is simple:

$$Q_{i+1}^- = Q_i + s(r^-)s(r^-)^T,$$

where $s(r^-)$ is the projection onto the reconstruction basis; $s(r^-)^T$ is the row that would be added onto A_i were the point r^- accepted.

We make a few implementation notes and observations. As written, Q^- will take on diagonal elements, but these are meaningless and should not be used for the purpose of selecting from the candidate points. Changing basis on the fly presents no difficulty, with the understanding that A_i and Q_i have to be recomputed each time it is changed. The greedy technique approaches a deterministic sequence as n_t grows, because the candidate points will identify the global minimum of q . Such design locations would depend only on the initial locations and therefore become correlated, but it is important to note that it is not in the reconstruction basis. As a result, increasing n_t as computational resources permit will tend to improve the reconstruction performance.

4.5.3 Examples

We now address several example computations; these are meant to represent typical effects of our greedy, RIP-based design (RIPD), but not to systematically characterize its behavior. In the first we consider the cosine basis, which is supported across the continuous measurement domain, and the effects of sensor noise estimation errors. In the second we study the Haar wavelets, involving unsupported basis functions on a discrete domain; there is no sensor noise in this case, but we vary the allowable transit distance. Our statistics are taken over five hundred trials, and our standard for comparison is a reconstruction based on randomly chosen points, termed RP.

Discrete-Cosine Transform in One and Two Dimensions with Noise

We use the discrete-cosine transform (DCT) in our first two examples, over a unit domain; the space is considered continuous. The specific problem parameters are $S = 15$ (sparsity), $N = 196$ (degrees of freedom), and $n_t = 100$ (candidates). We sample one point randomly and then proceed to design another forty-nine, so that $m = 10S/3$, a slightly lower value than is typical for exact reconstruction. We apply normally distributed, zero-mean noise to all measurements, with standard deviation ν equal to one percent of the RMS signal level, and solve the modified problem of l_1 minimization as given by [5]: $\min_x \|x\|_1$ s.t. $\|Ax - b\|_2 < \hat{\nu}\sqrt{m + \sqrt{2m}}$. We do not include kinematic constraints in the one-dimensional case.

As seen in Figure 4-7, the random and RIPD-selected points are indistinguishable viewed directly. Figure 4-8 shows the normalized reconstruction errors as a function of estimated noise level divided by true noise level; the zero log value on the horizontal axis corresponds to exact matching. Normalized reconstruction error is taken as the Euclidean norm of the error in y divided by the norm of the signal variation from its mean; we refer to the error simply as RMSE. We observe that the random-points have lowest mean error when $\hat{\nu}/\nu$ is around 0.3, an underestimate. This mean grows mildly with $\hat{\nu}$, and the scatter across the trials is significant but steady. Toward lower $\hat{\nu}$, however, the mean RMSE grows quickly to twice its minimum value and ultimately hits nearly three times that level. Coincident with this is an explosive growth of variability in the ensemble - a closer look at the distribution of cases shows there is a fraction of utterly failed reconstructions, with normalized errors greater than one(!). The RIPD points give a more consistent outcome in both mean and variability for the RMSE across the entire range of $\hat{\nu}/\nu$. Most notably, the mean RMSE is always lower than that of the random-points and thus the RIPD strategy on average improves the reconstruction error. The lower plot in Figure 4-8 confirms that for higher $\hat{\nu}/\nu$ a 15-20% reduction in RMSE is expected, whereas for low $\hat{\nu}/\nu$ this improvement is close to fifty percent. The variability is high however, and it must be kept in mind that only 60-70% of trials show an improvement. Overall, this result argues for setting $\hat{\nu}$




Figure 4-7: A representative signal of measurement point selection in the 1D case with discrete cosine transform. Fifty points are selected randomly, or designed according to the RIP, with one-hundred candidates considered for each addition.

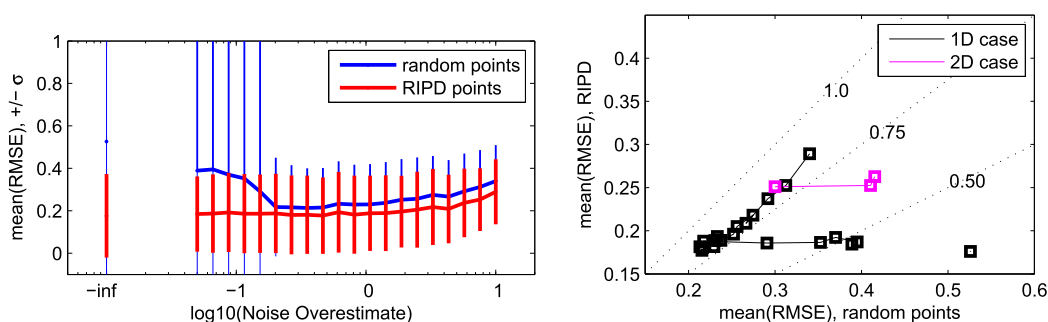


Figure 4-8: Summary of DCT results. *Top*: 1D case, mean values of the random-point and RIPD ensembles, along with lines showing plus and minus one standard deviation for the ensemble, as a function of $\hat{\nu}/\nu$; the leftmost point represents $\hat{\nu} = 0$. *Bottom*: Alternate view of the mean values above for one dimension, and three similar points for two-dimensional case with $\hat{\nu}/\nu = [0, 0.05, 1]$, right to left.

safely low - even zero when it is known that there is nonzero sensor noise - and using the RIPD strategy.

In two dimensions, for the DCT we assume no moves can be greater than 0.3 units Manhattan distance. The lower part of Figure 4-8 indicates a similar trend as for one dimension; as the sensor noise is underestimated, the RP paths deteriorate while the RIPD maintains a steady ensemble mean and standard deviation (not shown).

Haar Wavelet Transform in Two Dimensions without Noise

Unlike the DCT, the Haar transform has a discrete domain, and almost all modes are zero on a fraction of it. This makes D very sparse and therefore amenable to speed

improvements through proper programming. For instance in our example 32×32 image, $N = 1024$, but a given row of D has only 36 nonzero values, so that the number of nonzero scalar products in the row is 630, compared to the 523,776 products for a dense row. By maximizing re-use of the fixed part of the inner product set among a large group of candidates, we have been able to evaluate 200 candidates per second on a standard laptop computer. One might be tempted to apply this speed to a dynamic programming routine on the discrete state-space, but this attraction may be illusory because the inner product set for each given path-to-go still has to be maintained.

This example has no sensor noise, but we consider different step lengths. Figure 4-9 shows a test image derived from sea-surface temperature data, that has been explicitly sparsified. The image that is to be reconstructed has fifty Haar elements and contains both large-scale and small-scale features, most of which come directly over from the full-resolution base image. Paths associated with the RP and the RIPD strategy illustrate that the RIPD is regularizing the points, in the classical sense of reducing discrepancy and dispersion, along the lines of quasi-Monte Carlo sequences such as the Halton and Hammersley.

For $m = 4S$, RMSE values are lowest when the kinematics are unconstrained; this is a natural analog to the standard CS recommendation of random points. Even in this case, however, Figure 4-10 indicates that design via RIPD will offer a modest but significant improvement over RP of about ten percent. When the stepsize constraints are turned on, the improvement offered by RIPD increases to about fifteen percent, with improvement occurring in seventy percent of trials. Looking at $m = 2S$, we expect and observe worse RMSE levels, with RIPD offering no significant improvement (points along the line of slope 1.0 on the graph). An exception occurs at stepsize five, however, where the random-point paths fail badly because sometimes the agent becomes localized in a corner and so does not visit or even cross the space. RIPD helps the agent to escape these conditions; improvements here occur in sixty percent of trials. At the other extreme, setting $m = 10S$ for the stepsize of ten reduces RMSE values, the reconstructions are very good, and RIPD again gives a mild improvement. We cannot achieve an exact reconstruction in general for the Haar wavelets because




Figure 4-9: *Top left*: Discretized image of a sea-surface temperature map. *Top right*: The sparsified SST map with $S=50$. *Middle left*: Typical reconstruction result using 200 points randomly selected, with maximum move 10 pixels. *Middle right*: Same for 200 RIPD points; the RMSE is 15% less than for random points. *Bottom left*: Random path associated with the reconstruction above. *Bottom right*: Same for RIPD path. *Note*: the bounding box drawn is outside the feasible set of points.




Figure 4-10: Summary of 2D Haar results. Each color corresponds to a maximum step length (in pixels), with number of measurements m in square brackets. Mean values center on lines of ± 1 ensemble standard deviation. In all cases, increasing m reduces both random-point and RIPD reconstruction errors, as expected.

of limited support of the modes.

We have described a variation on standard compressive sensing, wherein we acknowledge the cost of transit and the desirability of interruptible CS missions. These and other aspects of CS data collection with mobile agents are interesting because they motivate new schemes for selecting points. Our main contribution in this paper was to apply a cheap rule for the incremental selection of points. A TSP path can efficiently visit a large set of such points, or they can be visited as generated; limited step lengths are then a simple way to reflect transit cost. Although the improvements in reconstruction error obtained through our RIPD points are modest, at 10-20% compared with random points, they appear to be robust; we encountered no conditions where the RIPD was outperformed.

4.6 Experimental Results

As discussed in Chapter 2, we performed experiments in Buzzards Bay, MA, in August 2011 using our kayak. One mission performed by the kayak was a compressive sensing



Figure 4-11: Photograph of TSP sampling mission in Buzzards Bay. The kayak measured the water depth near some small islands. The measurements were used for CS reconstruction of the sea floor.

environmental sampling mission. As a first investigation into compressive sensing with mobile agents, we performed a TSP tour of random points, measuring the depth at 100 locations in a square domain 100 meters wide. From the measurements, we performed an CS reconstruction to recover the basis coefficients of the field. We used the discrete sine and cosine basis. A photograph taken during the mission is shown in Fig. 4-11. The results are presented in Fig. 4-12.

When performing the reconstruction, we noticed that the final result depended heavily on two factors. The first factor is the expected measurement noise. This value tells the reconstruction how closely the final result must match the data. If the constraint is too tight, the basis coefficients returned are very large—the signal must move at extreme rates to match noisy measurement located near to each other. Obviously, the sea floor does not actually behave this way, so we must relax the noise estimate to allow for a smoother output. However, if the noise is estimated too high, the reconstruction will have enough freedom to almost completely ignore the measurements. Therefore, it is important to get an accurate estimate of the measurement noise for a particular mission. It should be noted that the sensor measurement noise tested in the lab may not be an accurate representation of the noise in the field, as waves and other real-world conditions will act to increase the noise level.

The noise levels of the sensor are related to those used in the reconstruction as follows. The reconstruction is performed by ℓ_1 minimization given by [5] as $\min_x \|x\|_1$ s.t. $\|Ax-b\|_2 < \hat{\nu}\sqrt{m} + \sqrt{2m}$. This constraint is derived so as to provide a 95 percent

confidence that the reconstruction will be feasible, assuming Gaussian measurement errors. In our sampling experiment, at each measurement location the kayak holds station for $N = 100$ measurements, and the results are averaged together to provide a single sample, therefore $\hat{\nu}$ is the variance of the sample mean. The value of $\hat{\nu}$ is related to the error of a single measurement by $\hat{\nu}^2 = \sigma^2/N$.

A second factor that affects the reconstruction is the number of basis coefficients used. Too few coefficients reduces the degrees of freedom available to the reconstruction. Whether there are negative effects to the reconstruction from too many coefficients is unknown at this time, however certainly using more coefficients slows down the computations.

Our results presented here give several combinations of number of basis coefficients and noise level. Certainly some of them seem unlikely to represent the actual sea floor measured by our kayaks. It remains to be seen whether there is a provable method for choosing these values, so we present the ensemble as a whole.

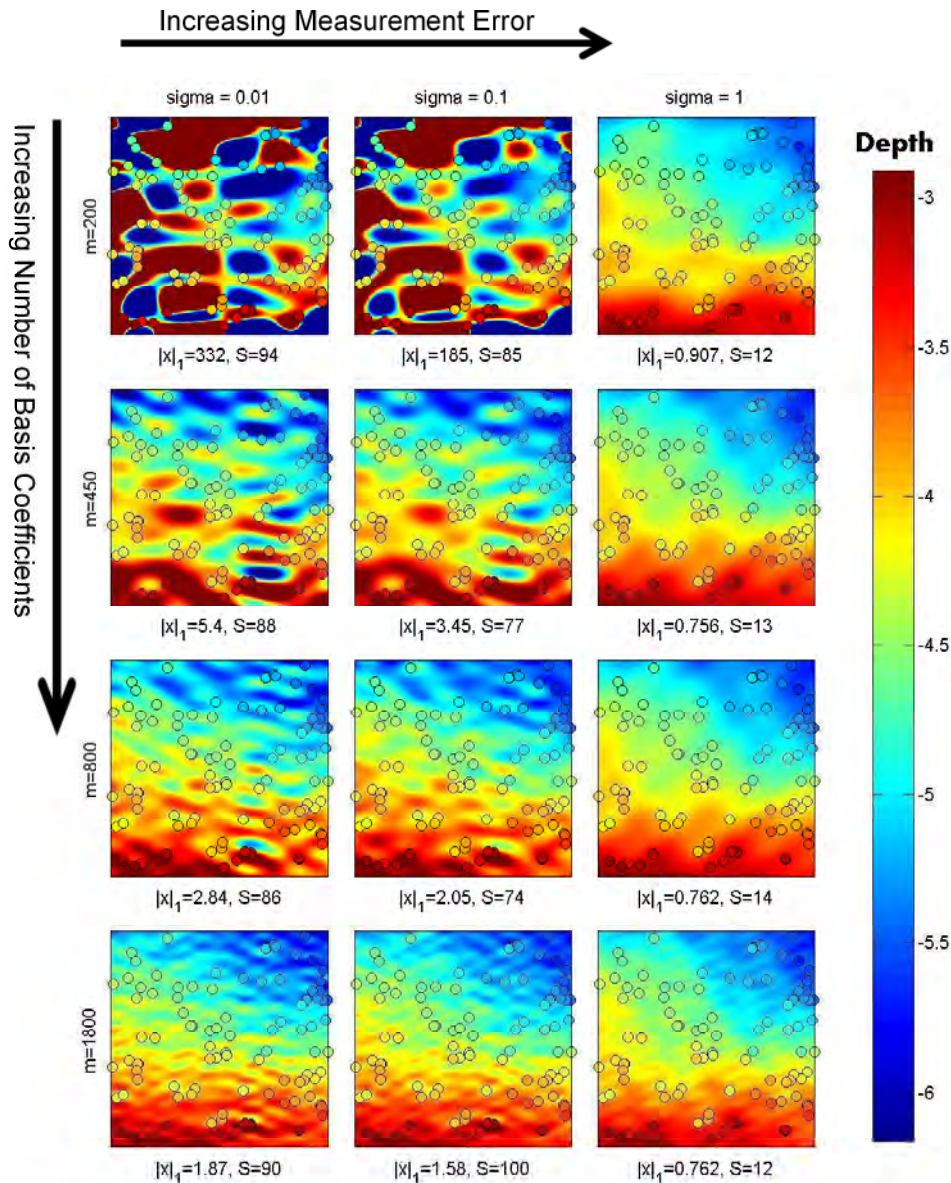


Figure 4-12: Experimental results from sea trials in Buzzards Bay, MA. Water depth was measured with an autonomous kayak performing a TSP tour of 100 random points. Sea floor field reconstruction was performed using the Fourier basis. In performing the reconstruction, sensor measurement error must be estimated, shown increasing from left to right. Additionally, the number of basis coefficients is shown increasing from top to bottom. Measurements are shown as circles. The ℓ_1 norm and sparsity, S , are given below each result. The trends can be inferred as follows. With low estimated measurement noise, the reconstruction is tightly constrained by the data, leading to a high sparsity and ℓ_1 norm and a reconstructed field that contains extreme values in places where no measurements were taken. Increasing the noise estimate relaxes this constraint, allowing for a lower sparsity and ℓ_1 norm and a more realistic reconstructed field. The number of basis coefficients determines the highest frequency component that is available for the reconstruction. Its effect on reconstruction is unclear, however it does act to decrease the ℓ_1 norm of the result.

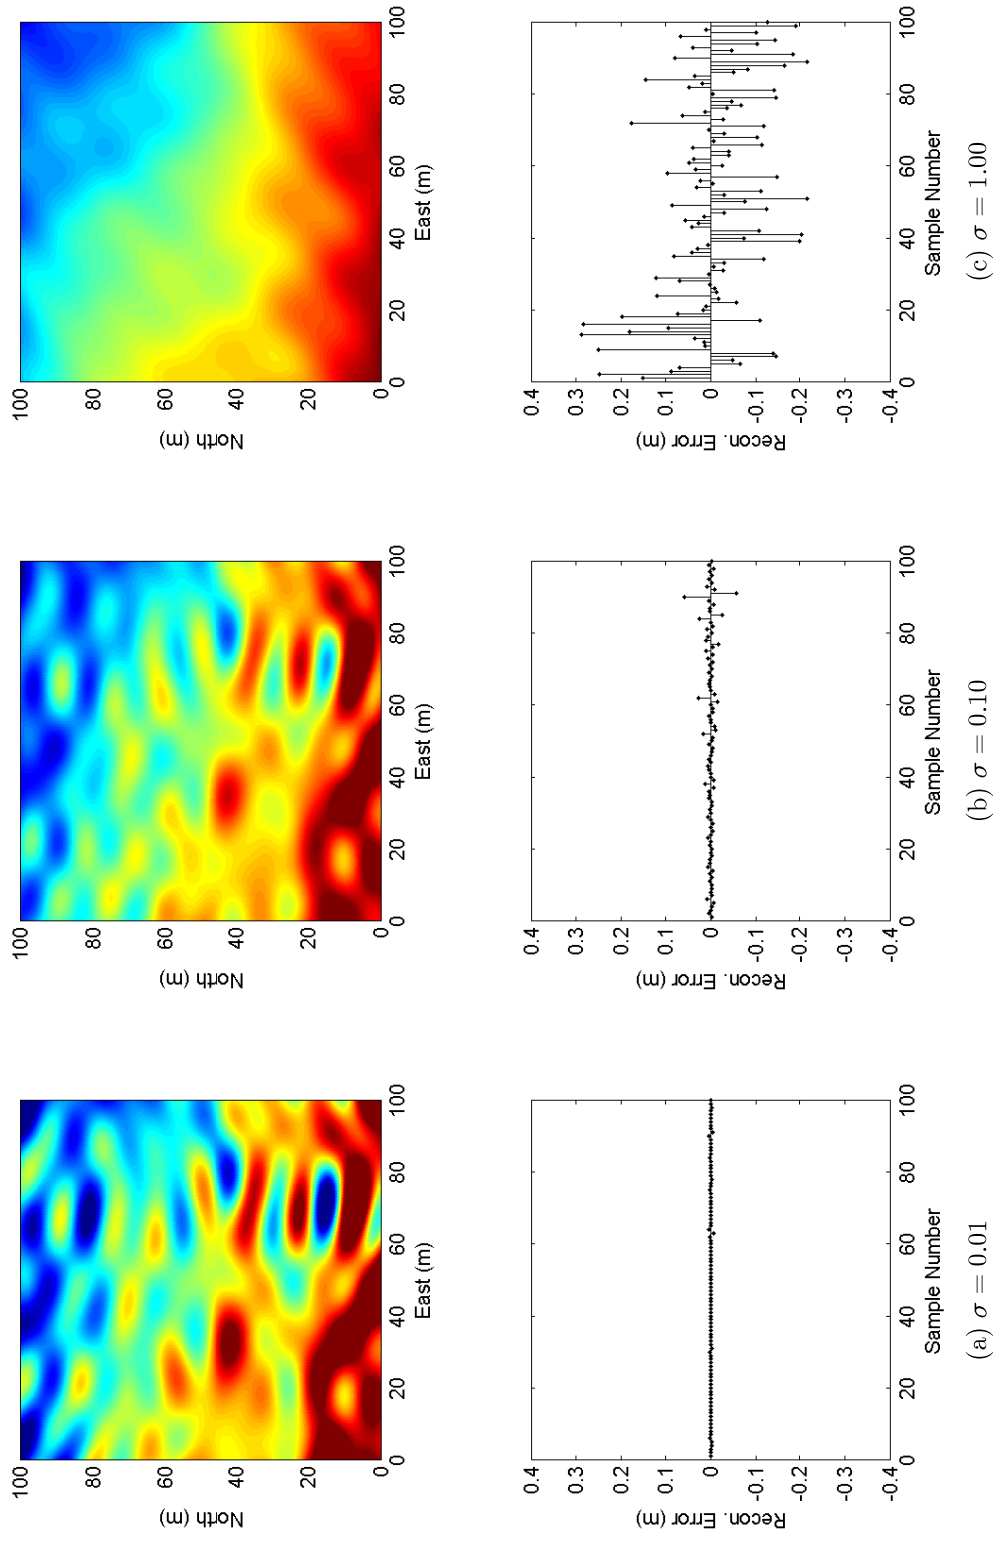


Figure 4-13: CS reconstruction error. These reconstructions correspond to the second row in Fig. 4-12, with the number of basis coefficients equal to 450. Here, the top row shows the reconstructed field, and the bottom row shows the difference between the measurements and the reconstructed field at the measurement locations. Columns (a), (b), and (c) give the reconstruction results for measurement errors (σ) equal to 0.01, 0.10, and 1.00, respectively.

4.7 Conclusion

Future work with the compressive sensing mission planning techniques presented here could involve multi-step look-ahead for the greedy approach, however this would greatly increase the computational cost. Also, various other heuristic optimization techniques could be explored for the random walk planning strategy. Future experiments could combine high-resolution field mapping with low-resolution compressive sensing missions, to provide a reliable ground truth for evaluation of the CS reconstruction.

We have presented here several new mission planning techniques for exploiting the sparsity of natural fields with the use of compressive sensing. The random measurement requirement for CS requires careful consideration when vehicles are used, as the high transit cost makes random movements throughout a domain impractical. We have shown several alternative point selection procedures that either outperform random points or are comparable with an improvement in mission flexibility. We have also presented results from experiments performed at sea that utilize compressive sensing. Clearly, in a sampling scenario where each measurement consumes a considerable share of mission resources, the benefits from reducing the sensing requirements warrant further study in compressive sensing applications for mobile agents.

Chapter 5

Conclusion

In this thesis we have developed several new pieces of infrastructure for research with autonomous marine networks. Additionally, we have presented novel mission planning techniques for compressive sensing which have the potential to dramatically reduce the measurement requirements for environmental sampling missions.

The infrastructure developments, two types of low-cost and scalable vehicles and a low-cost acoustic modem, are designed for the reduce the resource investment required for research in large-scale marine autonomy, both in terms of initial vehicle procurement costs and multi-vehicle deployment effort. The first vehicle, a small holonomic raft and high-resolution position system, is suitable for various laboratory experiments such as multi-vehicle sensor allocation problems, dynamic obstacle avoidance, heterogeneous fleet cooperation, and many others. The second vehicle is a five-foot-long kayak outfitted with a trolling motor for propulsion. The kayak has been shown to be seaworthy, and can carry a substantial payload of several tens of kilograms. At a cost of three thousand dollars, the autonomous kayak provides the capability of a large autonomous fleet with minimal resource investment.

The acoustic modems presented here are designed to be inexpensive to allow for multi-node deployments, especially in scenarios where the cost of each individual node is low, for example in distributed sensor networks or with the low-cost vehicles developed here.

The compressive sensing planning techniques provide improvements to reconstruc-

tion quality and mission flexibility, while taking into account the mobility and energy constraints present when sampling with mobile agents. This work has the potential to reduce the cost of environmental sampling missions by decreasing the number of samples required for accurate estimation of a scalar fields.

Future work with these infrastructure developments could include (a) building a new type of autonomous platform from the components already developed, such as an underwater or terrestrial vehicle, (b) building several more kayaks and/or rafts to expand our fleet, (c) with further modem performance improvements, creating a dense underwater acoustic network, or (d) using the vehicles for more experiments with compressive sensing.

Appendix A

Tables

Table A.1: Raft mechanical components

Description	Source / Fab. Method	Part Number	Qty	Total Cost
Foam hull	McMaster	3152T52	1	137
Frame	Waterjet Alum. Plate	N/A	4	32
Standoffs	McMaster	91145A190	16	1
Bilge motors	ITT Flow Control	Rule 500-25DC	8	120
Propellers	Dumas Products Inc.	3003	8	10
Prop coupling	Outsourced lathe work	N/A	8	24
Threaded rods	McMaster	95412A568	8	14
Nuts	McMaster	91845A029	16	1
PVC pipe	McMaster	48925K96	1	8
Clamping hangers	McMaster	3006T69	8	48
			Total	394

Table A.2: Raft electrical components

	Item	Cost
	Arduino	65
	OpenLog Data Logger	25
	XBee Series 2 (ZB)	50
	Garmin GPS	100
	Compass	150
	Rate gyro	25
	Temperature sensor	2
	Humidity sensor	17
	Circuit Board Connectors	10
	Printed Circuit Boards	20
	Waterproof Connectors	30
	Waterproof switch	14
	Electronics Enclosure	120
	Lithium Polymer Batteries - 3 cell	150
	NiMH Batteries, AA, Low self-discharge	10
	Voltage Regulator - 5v	30
	6-AA Battery Holder	2
	Motor Driver	80
	Schottky Diodes	3
	Solid State Relays	13
	Capacitors	2
	Misc	5
	Total	923

Table A.3: Kayak mechanical components, off-the-shelf

	Description	Source	Part #	qty	Price Ea.	Total Cost	
Kayak Thruster Assembly	Servo, High-torque, Digital	ServoCity	HS7940TH	1	139.99	140	
	Clamping Hub 1/4"	ServoCity	3162CH	1	7.99	8	
	Gear, Aluminum, 20 deg, 32P, 60T	ServoCity	NHL32-46-60	1	25.99	26	
	Servo Shaft, 1/4"	ServoCity	250HS	1	9.99	10	
	Gear, SS, 20 deg, 32P, 20T, 1/4" Bore	McMaster	6832K52	1	16.82	17	
	Gear Shaft, 1/4" x 6" Hardened SS	McMaster	6253K31	1	8.80	9	
	Bronze Sleeve Bearing, 1/4"	McMaster	1677k2	2	0.64	1	
	Spring-loaded PTFE Shaft Seal, 1"	McMaster	13125K85	1	21.30	21	
	Cap Screw, 1/4-20, 2-1/2", 316 SS	McMaster	92186A552	4	0.71	3	
	Shaft Collar, Nylon, 1/4"	McMaster	60485K69	2	14.79	30	
	Locknut, 1/4-20, 18-8 SS	McMaster	90715A125	4	0.17	1	
	Washer, 1/4", 316 SS	McMaster	91525A117	8	0.07	1	
	Socket Cap Screws, 1/4-20, 1", SS	McMaster	92196A542	4	0.18	1	
	Socket Cap Screws, 1/4-20, 5/8", SS	McMaster	92196A539	4	0.14	1	
	Trolling Motor, 12 lbs. thrust	Sevylor	N/A	1	80.00	80	
	Propeller	Sevylor	N/A	1	15.00	15	
	Frame	Drop-in T-slot Nut	McMaster	47065T151	4	2.27	9
		Vibration-damping Mount	McMaster	93945K34	4	2.10	8
	Bilge Pump	Bilge Pump, 650 GPH, Automatic	Whale	SS5212	1	60.00	60
		Inline Non-return Valve	Whale	LV1219	1	15.00	15
Bilge Flex Hose, 3/4", 3ft		Sierra Marine	116-120-0346W	3	1.00	3	
Other	Kayak	Wavesport	Fuse 35	1	749.00	749	
	Spray Skirt, Neoprene XS Rim	HarmonyGear	8024217	1	71.25	71	
	L-bracket for Skeg	McMaster	47065T223	4	3.98	16	
	Misc. screws, washers, nuts, etc.	McMaster	—	—	—	10	
					Total (USD)	1304	

Table A.4: Kayak mechanical components, fabricated. Labor cost is estimated at 30 USD/hr.

Assembly	Description	Fab. Methods	Material	Labor Cost	Stock Cost	Total Cost
Kayak Thruster Assembly	Servo Mounting Plate	3d Print	ABS	15	70	85
	Shaft Coupling	Machined	Aluminum 6061	30	2	32
	Bearing Block	Machined	Acetal Copolymer	30	8	38
	Sandwich Block	Machined	Acetal Copolymer	15	4	19
	Vertical Side Plate	Waterjet	Aluminum 6061	12	10	22
	Vertical Back Plate	Waterjet	Aluminum 6061	7.5	5	12.5
Skeg	Skeg Mount	Waterjet	Aluminum 6061	30	5	35
	Skeg	Cut, Drilled	UHMW PE	7.5	5	12.5
Frame	Frame rails	Cut, drilled	Aluminum T-slot Frame	7.5	30	37.5
	Frame Stand-off	Cut-to-length	Polyethylene	7.5	2	9.5
				Labor Cost	162	
					141	
				Labor Cost	162	303
					141	303

Table A.5: Kayak electrical components

	Item	Cost
	Arduino	65
	OpenLog Data Logger	25
	XBee Series 2 (ZB)	150
	Garmin GPS	100
	Compass	150
	Rate gyro	25
	Temperature sensor	2
	Humidity sensor	17
	Circuit Board Connectors	20
	Printed Circuit Boards	25
	Waterproof Connectors	60
	Waterproof switch	14
	Electronics Enclosure	110
	Lithium Polymer Batteries - 4 cell	430
	NiMH Batteries, AA, Low self-discharge	10
	Voltage Regulators - 5v, 12v	95
	6-AA Battery Holder	2
	Motor Driver	60
	Solid State Relays	13
	Capacitors	2
	Resistors	2
	Misc	5
	Total	1380

Bibliography

- [1] Applebaum, L., Howard, S., Searle, S., and Calderbank, R., Chirp sensing codes: Deterministic compressed sensing measurements for fast recovery. *Applied and Computational Harmonic Analysis*, 26(2):283–290, 2009.
- [2] Bhattacharya, S., Heidarrson, H., Sukhatme, G.S., and Kumar, V., Cooperative control of autonomous surface vehicles for oil skimming and cleanup. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 2374 – 2379, May 2011.
- [3] Candes, E. and Tao, T., Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12):4203 – 4215, December 2005.
- [4] Candes, E. and Wakin, M., An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):21 –30, March 2008.
- [5] Candes, E., Wakin, M., and Boyd, S., Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14:877–905, 2008.
- [6] Candes, E.J., Romberg, J., and Tao, T., Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.
- [7] Chen, C.Y. and Vaidyanathan, P.P., Compressed sensing in MIMO radar. In *2008 2nd Asilomar Conference on Signals, Systems and Computers*, pp. 41–44, 2008.
- [8] Chitre, M., Shahabudeen, S., and Stojanovic, M., Underwater acoustic communications and networking: Recent advances and future challenges. *Marine Technology Society Journal*, Spring:103–116, 2008.
- [9] Choset, H., Coverage for robotics - a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31(1-4):113–126, 2001.
- [10] Choset, H. and Pignon, P., Coverage path planning: the boustrophedon cellular decomposition. In *International Conference on Field and Service Robotics*, 1997.
- [11] Christofides, N., Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon University, 1976.

- [12] Crouse, M.S., Nowak, R.D., and Baraniuk, R.G., Wavelet-based statistical signal processing using hidden Markov models. *IEEE Transactions on Signal Processing*, 46(4):886–902, 1998.
- [13] Curcio, J., Leonard, J., and Patrikalakis, A., SCOUT - a low cost autonomous surface platform for research in cooperative autonomy. In *OCEANS, 2005. Proceedings of MTS/IEEE*, pp. 725 – 729 Vol. 1, 2005.
- [14] Curcio, J., Leonard, J., Vaganay, J., Patrikalakis, A., Bahr, A., Battle, D., Schmidt, H., and Grund, M., Experiments in moving baseline navigation using autonomous surface craft. In *OCEANS, 2005. Proceedings of MTS/IEEE*, pp. 730 – 735 Vol. 1, 2005.
- [15] Curtin, T., Bellingham, J., Catipovic, J., and Webb, D., Autonomous oceanographic sampling networks. *Oceanography*, 6(3):86–94, 1993.
- [16] Desert Star Systems, SAM-1 Miniature Acoustic Modem. http://www.desertstar.com/Products_product.aspx?intProductID=4, October 2011.
- [17] DeVore, R., Deterministic constructions of compressed sensing matrices. *Journal of Complexity*, 23(4-6):918–925, 2007.
- [18] Donoho, D.L., Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [19] Duarte-Carvajalino, J. and Sapiro, G., Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization. *Image Processing, IEEE Transactions on*, 18(7):1395 –1408, July 2009.
- [20] Elad, M., Optimized projections for compressed sensing. *Signal Processing, IEEE Transactions on*, 55(12):5695 –5702, December 2007.
- [21] Engel, R. and Kalwa, J., Relative positioning of multiple underwater vehicles in the GREX project. In *OCEANS 2009 - EUROPE*, pp. 1 –7, May 2009.
- [22] Evologics, R-Series. <http://www.evologics.de/en/products/acoustics/index.html>, October 2011.
- [23] Exar Corporation, *XR2211 FSK Demodulator/Tone Decoder, Datasheet*. June 1997.
- [24] Fiorelli, E., Bhatta, P., and Leonard, N.E., Adaptive sampling using feedback control of an autonomous underwater glider fleet. In *Proc. 13th Int. Symposium on Unmanned Untethered Submersible Tech*, pp. 1–16, 2003.
- [25] Fiorelli, E., Leonard, N.E., Bhatta, P., Paley, D.A., Bachmayer, R., and Fratantoni, D.M., Multi-AUV control and adaptive sampling in Monterey Bay. *Oceanic Engineering, IEEE Journal of*, 31(4):935 –948, October 2006.

- [26] Graeme, J., *Amplifier Applications*. The McGraw-Hill Companies Inc., Martinsburg, West Virginia, 1999.
- [27] Greytak, M.B., *High performance path following for marine vehicles using azimuthing podded propulsion*. Master's thesis, Massachusetts Institute of Technology, Dept. of Mechanical Engineering, February 2006.
- [28] Hover, F., Hummel, R., Mitra, U., and Sukhatme, G., One-step-ahead kinematic compressive sensing. In *IEEE GLOBECOM 2011 - Wireless Networking for Unmanned Autonomous Vehicles*, to appear, 2011.
- [29] Hummel, R., Poduri, S., Hover, F., Mitra, U., and Sukhatme, G., Mission design for compressive sensing with mobile robots. In *Proc. Int. Conf. Robotics and Automation*, 2011.
- [30] Hummel, R.A., Taylor, J.A., and Hover, F.S., Numerical optimization of generative network parameters. In *ASME International Mechanical Engineering Congress and Exposition*, 2010.
- [31] Indyk, P., Explicit constructions for compressed sensing of sparse signals. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 30–33, Society for Industrial and Applied Mathematics, 2008.
- [32] Iniewski, K., *Medical Imaging: Principles, Detectors, and Electronics*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2009.
- [33] Karp, R.M., Probabilistic analysis of partitioning algorithms for the traveling-salesman problem in the plane. *Mathematics of Operations Research*, 2(3):209 – 224, 1977.
- [34] Kholkin, A.L., Pertsev, N.A., and Goltsev, A.V., Piezoelectricity and crystal symmetry. In A. Safari and E.K. Akdogan, editors, *Piezoelectric and Acoustic Materials for Transducer Applications*, chapter 2, Springer, New York, 2008.
- [35] Kunz, C., Murphy, C., Camilli, R., Singh, H., Bailey, J., Eustice, R., Jakuba, M., Nakamura, K., Roman, C., Sato, T., Sohn, R., and Willis, C., Deep sea underwater robotic exploration in the ice-covered Arctic Ocean with AUVs. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pp. 3654 –3660, September 2008.
- [36] Labat, J., Lapierre, G., and Trubuil, J., Iterative equalization for underwater acoustic channels potentiality for the TPIDENT system. In *OCEANS 2003. Proceedings*, volume 3, pp. 1547 –1553 Vol.3, September 2003.
- [37] LaValle, S.M. and Kuffner, J.J., Randomized kinodynamic planning. *The International Journal of Robotics Research*, 20(5):378–400, 2001.

- [38] Leonard, N., Paley, D., Lekien, F., Sepulchre, R., Fratantoni, D., and Davis, R., Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE*, 95(1):48–74, January 2007.
- [39] Leonard, N.E., Paley, D.A., Davis, R.E., Fratantoni, D.M., Lekien, F., and Zhang, F., Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay. *Journal of Field Robotics*, 27(6):718–740, 2010.
- [40] Lermusiaux, P.F., Adaptive modeling, adaptive data assimilation and adaptive sampling. *Physica D: Nonlinear Phenomena*, 230(1-2):172–196, 2007.
- [41] Li, W. and Preisig, J., Estimation of rapidly time-varying sparse channels. *Oceanic Engineering, IEEE Journal of*, 32(4):927–939, October 2007.
- [42] Liang, J.J., Qin, A.K., Suganthan, P.N., and Baskar, S., Comprehensive learning particle swarm optimizer for global optimization of multimodal functions. *IEEE Transactions on Evolutionary Computation*, 10(3):281–295, 2006.
- [43] Linkquest Inc., Underwater acoustic modem models. <http://www.link-quest.com/html/models1.htm>, October 2011.
- [44] Lochmatter, T., Roduit, P., Cianci, C., Correll, N., Jacot, J., and Martinoli, A., SwisTrack - a flexible open source tracking software for multi-agent systems. In *Intelligent Robots and Systems, IEEE/RSJ International Conference on*, pp. 4004–4010, September 2008.
- [45] Lopez, M. and Singer, A., A DFE coefficient placement algorithm for sparse reverberant channels. *Communications, IEEE Transactions on*, 49(8):1334–1338, August 2001.
- [46] Lurton, X., *An Introduction to Underwater Acoustics: Principles and Applications*. Springer-Praxis, New York, 2002.
- [47] Manley, J., Development of the autonomous surface craft ACES. In *OCEANS 1997 MTS/IEEE Conference Proceedings*, volume 2, pp. 827–832, October 1997.
- [48] Manley, J., Marsh, A., Cornforth, W., and Wiseman, C., Evolution of the autonomous surface craft AutoCat. In *OCEANS 2000 MTS/IEEE Conference and Exhibition*, volume 1, pp. 403–408, 2000.
- [49] Maxim Integrated Products, *MAX274/275 4th- and 8th-Order Continuous-time Active Filters, Datasheet*. October 1996.
- [50] Maxim Integrated Products, *MAX274 Filter Design Software*. Downloaded from <http://www.maxim-ic.com/products/filters/>, March 2011.
- [51] McCarthy, A., Heterogeneous raft sampling mission. Unpublished data, August 2011.

- [52] McCarthy, A., Raft motor calibration procedure and results. Unpublished data, June 2011.
- [53] McNish, M.J., *Effects of Uniform Target Density on Random Search*. Master's thesis, Naval Postgraduate School, 1987.
- [54] National Semiconductor Corporation, Op amp circuit collection. Application Note 31, September 2002.
- [55] Pon, R., Batalin, M.A., Gordon, J., Kansal, A., Liu, D., Rahimi, M., Shirachi, L., Yu, Y., Hansen, M., Kaiser, W.J., Srivastava, M., Sukhatme, G., and Estrin, D., Networked infomechanical systems: A mobile embedded networked sensor platform. *2005 Fourth International Symposium on Information Processing in Sensor Networks*, pp. 376–381, 2005.
- [56] Popa, D.O., Sanderson, A.C., Komerska, R.J., Mupparapu, S.S., Blidberg, D.R., and Chappel, S.G., Adaptive sampling algorithms for multiple autonomous underwater vehicles. *2004 IEEE/OES Autonomous Underwater Vehicles*, pp. 108–118, 2004.
- [57] Reinhart, R., Steil, J., Huntsberger, T., and Stoica, A., Tacking reduces bow-diving of high-speed unmanned sea surface vehicles. In *Emerging Security Technologies (EST), 2010 International Conference on*, pp. 177 –182, September 2010.
- [58] Rodriguez-Ortiz, C.D., *Automated bathymetry mapping using an autonomous surface craft*. Thesis (M. Eng.), Massachusetts Institute of Technology, Dept. of Ocean Engineering, September 1996.
- [59] Schmidt, H., Bellingham, J., Johnson, M., Herold, D., Farmer, D., and Pawlowicz, R., Real-time frontal mapping with AUVs in a coastal environment. In *OCEANS '96. MTS/IEEE: Prospects for the 21st Century*, volume 3, pp. 1094–1098 vol.3, September 1996.
- [60] Scholin, C., Jensen, S., Roman, B., Massion, E., Marin, R., Preston, C., Greenfield, D., Jones, W., and Wheeler, K., The Environmental Sample Processor (ESP) - an autonomous robotic device for detecting microorganisms remotely using molecular probe technology. *OCEANS 2006*, pp. 1–4, 2006.
- [61] Shannon, C.E., Communications in the presence of noise. In *Proceedings of the Institute of Radio Engineers*, volume 31, pp. 10–21, reprint as classic paper in: Proc. IEEE, vol. 86, no. 2, (Feb. 1998), January 1949.
- [62] Singh, A., Krause, A., Guestrin, C., Kaiser, W., and Batalin, M., Efficient planning of informative paths for multiple robots. In *IJCAI'07: Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 2007.

- [63] Singh, A., Krause, A., and Kaiser, W.J., Nonmyopic adaptive informative path planning for multiple robots. In *IJCAI'09: Proc. of the 21st International Joint Conference on Artificial Intelligence*, pp. 1843–1850, 2009.
- [64] Singh, A., Nowak, R., and Ramanathan, P., Active learning for adaptive mobile sensing networks. *IPSN 2006: The Fifth International Conference on Information Processing in Sensor Networks*, pp. 60–68, 2006.
- [65] Singh, S., Webster, S.E., Freitag, L., Whitcomb, L.L., Ball, K., Bailey, J., and Taylor, C., Acoustic communication performance of the WHOI micro-modem in sea trials of the *Nereus* vehicle to 11,000 m depth. In *Proc. IEEE/MTS OCEANS Conf. Exhib.*, October 2009.
- [66] Smith, R.N., Stauffer, B., Heidarrson, H., Das, J., Pereira, A., Chao, Y., Darjany, L., Cetinic, I., Estrin, D., Oberg, C., Ragan, M., Jones, B.H., Sukhatme, G.S., and Caron, D.A., Design and implementation of sensor networks for the observation and research of harmful algal blooms in Southern California coastal waters. In *Proceedings of the Conference on Coastal Environmental Sensing Networks*, 2009.
- [67] Sonardyne, LBL, USBL and LUSBL. <http://www.sonardyne.com/products/positioning.html>, October 2011.
- [68] Stein, D.M., Probabilistic analysis of a routing problem. *Mathematics of Operations Research*, 3(2):89 – 101, 1978.
- [69] Stojanovic, M., Catipovic, J., and Proakis, J., Adaptive multichannel combining and equalization for underwater acoustic communications. *Journal of the Acoustical Society of America*, 94(3):1621–1631, September 1993.
- [70] Stojanovic, M., Freitag, L., and Johnson, M., Channel-estimation-based adaptive equalization of underwater acoustic signals. In *OCEANS '99 MTS/IEEE. Riding the Crest into the 21st Century*, volume 2, pp. 590 –595 vol.2, 1999.
- [71] Strohmer, T. and Friedlander, B., Compressed sensing in MIMO radar - algorithms and performance. In *43rd Asilomar Conference on Signals, Systems and Computers*, pp. 464–468, 2009.
- [72] Sukhatme, G.S., Dhariwal, A., Zhang, B., Oberg, C., Stauffer, B., and Caron, D.A., Design and development of a wireless robotic networked aquatic microbial observing system. *Environmental Engineering Science*, 24(2):205–215, 2007.
- [73] Teledyne Benthos, Acoustic modems. <http://www.benthos.com/acoustic-teleonar-modem-product-comparison.asp>, October 2011.
- [74] Texas Instruments, *OPA227/228 High Speed, Low Noise Operational Amplifiers, Datasheet*. [Revised Jan 2005], May 1998.

- [75] Tritech International Ltd., Tritech Micron Data Modem. http://www.tritech.co.uk/products/products-micron_modem.htm, October 2011.
- [76] Williams, A.B. and Taylor, F.J., *Electronic Filter Design Handbook*. McGraw-Hill Professional, New York, fourth edition, 2006.
- [77] Woods Hole Oceanographic Institution, Micro-Modem. <http://acomms.whoi.edu/umodem/>, October 2011.
- [78] Xiong, F., *Digital Modulation Techniques*. Artech House Publishers, Boston, Massachusetts, second edition, April 2006.
- [79] Xu, J., Pi, Y., and Cao, Z., Optimized projection matrix for compressive sensing. *EURASIP Journal on Advances in Signal Processing*, 2010.
- [80] Yilmaz, N.K., Evangelinos, C., Lermusiaux, P.F.J., and Patrikalakis, N.A., Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. *IEEE Journal of Oceanic Engineering*, 33(4):522–537, 2008.
- [81] Yu, Y., Petropulu, A., and Poor, H., Measurement matrix design for compressive sensing based MIMO radar. *Signal Processing, IEEE Transactions on*, PP(99):1, 2011.
- [82] Zhang, B. and Sukhatme, G.S., Adaptive sampling for estimating a scalar field using a robotic boat and a sensor network. *Proceedings of the 2007 IEEE International Conference on Robotics and Automation, Vols 1-10*, pp. 3673–3680, 2007.
- [83] Zhang, F., Fiorelli, E., and Leonard, N., Exploring scalar fields using multiple sensor platforms: Tracking level curves. In *Decision and Control, 2007 46th IEEE Conference on*, pp. 3579–3584, December 2007.