

**Using Light Underwater:
Devices, Algorithms and Systems
for Maritime Persistent Surveillance**

by

Iuliu Vasilescu

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2009

© Massachusetts Institute of Technology 2009. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
January 30, 2009

Certified by
Daniela Rus
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
Terry P. Orlando
Chairman, Department Committee on Graduate Students

**Using Light Underwater:
Devices, Algorithms and Systems
for Maritime Persistent Surveillance**

by

Iuliu Vasilescu

Submitted to the Department of Electrical Engineering and Computer Science
on January 30, 2009, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

This thesis presents a novel approach to long-term marine data collection and monitoring. Long-term marine data collection is a key component for understanding planetary scale physical processes and for studying and understanding marine life. Marine monitoring is an important activity for border protection, port security and offshore oil field operations. However, monitoring is not easy because salt water is a harsh environment for humans and for instruments. Radio communication and remote sensing are difficult below ocean surface.

Our approach to ocean data collection relies on the integration of (1) a network of underwater sensor nodes with acoustic and optical communication, (2) an autonomous underwater vehicle (AUV) and (3) a novel sensing device. A key characteristic is the extensive use of visible light for information transfer underwater. We use light for sensing, communication and control.

We envision a system composed of sensor nodes that are deployed at static locations for data collection. Using acoustic signaling and pairwise ranging the sensor nodes can compute their positions (self-localize) and track mobile objects (e.g., AUVs). The AUV can visit the sensor nodes periodically and download their data using the high speed, low power optical communication. One consequence of using optical communication for the bulk of the data transfer is that less data needs to be transferred over the acoustic links, thus enabling the use of low power, low data rate techniques. For navigation, the AUV can rely on the tracking information provided by the sensor network. In addition, the AUV can dock and transport sensor nodes efficiently, enabling their autonomous relocation and recovery. The main application of our system is coral reef ecosystem research and health monitoring. In this application the robot and the sensor nodes can be fitted with our novel imaging sensor, capable of taking underwater color-accurate photographs for reef health assessment and species identification.

Compared to existing techniques, our approach: (1) simplifies the deployment of

sensors through sensor self-localization, (2) provides sensor status information and thus enables the user to capture rare events or to react to sensor failure, (3) provides the user real time data and thus enables adaptive sampling, (4) simplifies mobile sensing underwater by providing position information to underwater robots, (5) collects new types of data (accurate color images) through the use of new sensors.

We present several innovations that enable our approach: (1) an adaptive illumination approach to underwater imaging, (2) an underwater optical communication system using green light, (3) a low power modulation and medium access protocol for underwater acoustic telemetry, (4) a new AUV design capable of hovering and of efficiently transporting dynamic payloads.

We present the design, fabrication and evaluation of a hardware platform to validate our approach. Our platform includes: (1) AQUANET, a wireless underwater sensor network composed of AQUANODES, (2) AMOUR, an underwater vehicle capable of autonomous navigation, data muling, docking and efficient transport of dynamic payloads and (3) AQUALIGHT an underwater variable-spectrum Xenon strobe which enables underwater color accurate photography. We use this platform to implement and experimentally evaluate our algorithms and protocols.

Thesis Supervisor: Daniela Rus

Title: Professor of Electrical Engineering and Computer Science

Acknowledgments

This thesis only exists because of many wonderful people I met during my PhD. I will be forever in debt to my advisor, Daniela Rus, for her help, kindness, continuous support and infinite patience. She gave me the right amount of freedom and guidance, took me places I would not have hoped to visit and was always present to encourage me and advise me. It was great working with you, Daniela!

I thank you, Dina and Rodney, for accepting to be in my thesis committee, for taking the time to look at my work, being forgiving, and for our excellent discussions. I'd like to thank the excellent professors that taught and advised me during my PhD: Michael Sipser, John Leonard, Henrik Schmidt, Michael Ernst, Jimm Glass and Luca Daniel. I was lucky to have great professors in Romania: Maria Burz, Mihaita Draghici, Aurelia Mirauta, Emilia Mititelu, Irina Athanasiu, and Adrian Surpateanu. I thank you for all the heart and education you gave me. I miss those years so much!

I would like to thank my close collaborators and friends, we've made a great team together. Keith: I'll always remember our discussions and your help; you jump-started my robotics skills. Carrick: we build so many things and we went on so many great trips together - those were some of the best times of my PhD. I'd be very worried to go on a research trip without you. Paulina: our first robot exploded, but it was quite fun to build it and play with it. Alex: I still can't believe we are not in the same building anymore! Marty: I'll always remember your elegant mechanical designs. Daniel, Jan and Stefan: precise, collected Germans - the robot got so much better with you around. Marek and Nikolaus: my other very adorable German friends! Katya: no problem was too hard for us, together. Talking to you has always been very refreshing. Matt: are you at the gym? Kyle, Mac, David, Philipp, Elizabeth, Albert, Olivier, Lijin, Jessie, Sejoon, Seung Kook, Johnny, Yoni, Maria, Ann, Marcia, and Mieke - a true family in which I felt very welcomed!

I thank all my good friends! Florin: too bad grad school ended, we could have continue to be great roommates. Alex: I always enjoyed talking to you - we are

simply on the same wavelength. Eduardo: thank you for singing my name, always smiling and providing energy boosts. Irina: I hope we will dance again, the music will still be the same. Johnna: we started and finished together, so let's stay in sync! Adrian, Radu, Laura, Alina, Mihai, Monica, Emanuel, Karola, Bogdan, Evdokia, Tilke, Evelyn, and Tanya: I was so happy to talk to you every single time - I hope we'll always be in touch!

I thank my parents, Zamfir and Cornelia, for their more than thirty years of unconditional love, care, support and patience. I thank my sister, Corina, for being patient with me and not complaining for my long months of silence.

Thank you Anya for being with me, so gentle and open-hearted! You gave me all the motivation I needed to finish what seemed never to end!

Contents

1	Introduction	19
1.1	Motivation	20
1.1.1	Physical Oceanography	21
1.1.2	Marine Biology	21
1.1.3	Ship Hull Inspection	21
1.1.4	Offshore Oil Exploitation	21
1.1.5	Harbor and border protection	22
1.2	Challenges	22
1.3	Problem Statement	23
1.4	Related Work	24
1.5	Approach	26
1.5.1	Contributions	31
1.5.2	Application Limits	32
1.6	Outline	33
2	Underwater Perception: AQUALIGHT	35
2.1	Motivation	35
2.2	Related work	39
2.3	Approach	42
2.3.1	Optical Properties of Water	43
2.3.2	The Naive Approach to Addaptive Illumination	45
2.3.3	The Human Visual System	48
2.3.4	The Color Rendering Index	50

2.3.5	Perceptual Adaptive Illumination	52
2.4	Hardware Description	55
2.5	Experiments and Evaluation	57
2.5.1	Output calibration	58
2.5.2	Optimization	60
2.5.3	Pool experiments	62
2.5.4	Ocean experiments	67
2.6	Summary	67
3	Underwater Sensor Network: AQUANET	75
3.1	Related Work	76
3.2	Why Optical?	78
3.3	Hardware Description	80
3.4	The Optical Networking	83
3.4.1	Why Green Light?	84
3.4.2	Hardware Description	86
3.4.3	Physical Layer	88
3.5	Acoustic networking	90
3.5.1	Hardware Description	90
3.5.2	Physical Layer	93
3.5.3	Medium Access Control	96
3.6	Experiments	103
3.6.1	Optical Communication and Data Muling	103
3.6.2	Acoustic Communication and Ranging	106
3.6.3	Localization and Tracking	107
3.7	Summary	110
4	Underwater Autonomous Vehicle: AMOUR	113
4.1	Related Work	115
4.2	Hardware Description	117
4.2.1	Propulsion	119

4.2.2	Docking	120
4.2.3	Buoyancy and Balance	121
4.2.4	Inertial Measurement Unit	123
4.2.5	Central Controller Board	125
4.2.6	Sensor Node	126
4.3	Algorithms and Control	126
4.3.1	Pose estimation	126
4.3.2	Hovering and Motion	128
4.3.3	Docking	130
4.3.4	Buoyancy and Balance	131
4.4	Experiments	133
4.4.1	Control	133
4.4.2	Navigation	135
4.4.3	Docking	138
4.4.4	Buoyancy and Balance	140
4.5	Summary	141
5	Conclusions	149
5.1	Contributions and Lessons Learned	149
5.2	Near Future Goals	152
5.3	Future Research Directions	153
A	Adaptive Light Optimization	155
B	Flash Optimization Results	159
C	Adaptive Flash Schematics	163
D	Flash Timing Control	169
E	AQUANET Schematics	173
F	TDMA Implementation	187

G AMOUR Motor Controllers	191
H AMOUR Docking Controller	197

List of Figures

1-1	Water attenuation coefficient for electromagnetic waves	20
1-2	Approach concept picture	27
1-3	The hardware instantiation of our approach	28
1-4	Thesis contributions and experimental validation	30
2-1	The first underwater color photo	36
2-2	Simulated effect of water on colors	38
2-3	Light absorption coefficient $a_w(\lambda)$ for water	44
2-4	Spectral power distribution of sun light	45
2-5	Simulated apparent color of Munsell 5Y8/10 color sample in water	46
2-6	Spectral power distribution for a naive compensatory light source	47
2-7	Spectral power distribution of CIE standard illuminant D65	49
2-8	The reflection coefficient of Munsell 5Y8/10 color sample	50
2-9	Reflected spectral power distribution of Munsell 5Y8/10 color sample	51
2-10	Normalized sensitivities of the three types of cone cells present in human retina	52
2-11	The human visual system color response	53
2-12	The block diagram of the underwater flash	58
2-13	Picture of the underwater flash	59
2-14	Picture of the flash and camera assembly	60
2-15	The flash output calibration test-bench	61
2-16	The normalized output spectral energy distribution of the 6 flashes	62

2-17	The normalized output spectral energy distribution of the unfiltered flash vs flash time	63
2-18	The normalized output energy of the unfiltered flash	64
2-19	The normalized total output energy of the adaptive flash	65
2-20	The normalized total output energy of the adaptive flash in water	66
2-21	The pool setup for underwater flash experiments	67
2-22	Pictures of the color palette taken via 4 methods	69
2-23	Color accuracy analysis plot	70
2-24	The red color sample rendering	70
2-25	The mustard color sample rendering	71
2-26	The blue color sample rendering	71
2-27	Reef picture from Fiji	72
2-28	Reef picture from French Polynesia	73
2-29	Reef picture from Cayman Island	74
3-1	AQUANODE version 3	81
3-2	AQUANODE version 1 (AquaFleck)	82
3-3	The sensor nodes, version 2	83
3-4	Spectral response of PDB-C156 PIN photodiode	84
3-5	Total optical system's gain	86
3-6	Block schematic of the optical modem	87
3-7	The optical modem	88
3-8	Picture of the long range optical modem	89
3-9	The optical modulation	89
3-10	Block schematic of the acoustic modem ver. 3 (current)	91
3-11	Picture of the acoustic modem ver. 3 (current)	91
3-12	Picture of the acoustic modem and transducer ver. 1	93
3-13	Picture of the acoustic modem and transducer ver. 2 (current)	94
3-14	Acoustic modulation: sparse symbols	95

3-15	Plot of the maximum theoretical data rate vs the transmitted energy per bit	97
3-16	The time lines of three nodes during TDMA communication	100
3-17	Starbug AUV acting as data mule	106
3-18	Acoustic range measurement test	108
3-19	Acoustic tracking test	110
3-20	The acoustic packet success rate experiment.	111
3-21	Tracking AMOUR	112
4-1	AMOUR and a few sensor nodes	114
4-2	Picture of the three generations of AMOUR	117
4-3	The docking mechanism	121
4-4	Picture of the docking cone	122
4-5	A docking sensor design	123
4-6	AMOURwith the buoyancy and balance mechanism installed	124
4-7	The Inertial Navigation Unit	125
4-8	The balance control mechanism in operation	131
4-9	Energy analysis plot: thrusters vs buoyancy engine	132
4-10	The control loops for the buoyancy and balance systems	132
4-11	Impulse response of the pitch axis	135
4-12	Impulse response of the roll axis	136
4-13	Impulse response of the yaw axis	137
4-14	Impulse response of the depth controller	138
4-15	AMOUR during a daylight mission and a night mission	142
4-16	Plot of an autonomous mission path	143
4-17	GPS path during endurance mission	144
4-18	Autonomous docking experiment	145
4-19	AMOUR docking StarBug	146
4-20	Thrusters' output during buoyancy experiment	146
4-21	Thrusters' output during battery movement experiment	147

4-22	Thrusters' output during balance experiment	147
C-1	Flash Timing Controller 1	165
C-2	Flash Timing Controller 2	166
C-3	Flash Timing Controller 3	167
C-4	Flash Voltage Inverter	168
E-1	AQUANODE Interface	175
E-2	AQUANODE Battery Management	176
E-3	AQUANODE Pressure Sensors	177
E-4	AQUANODE A/D Converter	178
E-5	AQUANODE Power Management CPU	179
E-6	AQUANODE Winch Control	180
E-7	AQUANODE Power Management	181
E-8	Optical Modem Schematic	182
E-9	AQUANODE CPU Power	183
E-10	AQUANODE Radio and GPS	184
E-11	AQUANODE Accelerometer and Compass	185
E-12	AQUANODE CPU Flash Storage	186
G-1	AMOUR Motor controller power board	193
G-2	AMOUR Motor controller regulator	194
G-3	AMOUR Motor controller feedback	195
G-4	AMOUR Motor controller CPU	196

List of Tables

2.1	The fire duration for the 6 flashes for subject to camera distances between 1m and 5m and their expected CRI.	62
3.1	Reception rates for optical communication in clear water. The reception rate drops sharply at the modem's range limit.	105
3.2	Reception rates for optical communication in clear water for the long range modem. No reception was observed beyond 7m.	105
3.3	Acoustically measured ranges between node A and node E compared to the ground truth.	107
3.4	Acoustic packet success rate between node E and nodes A-D. The one way packet rates success rates were computed considering equal that the two packets in a round trip communication have an equal probability of success.	109
3.5	The ranges between the static sensor nodes A-D as measured and exchanged by the nodes. Each node collected a copy of this matrix. . .	109
3.6	The position of the sensor node in the locally constructed coordinate system, based on pairwise ranges.	109
4.1	Control output mapping for vertical orientation of the robot	128
4.2	Control output mapping for horizontal orientation of the robot	129
4.3	Waypoints used for the autonomous drive test	139

List of Algorithms

1	Light optimization	55
2	Flash calibration and optimization	63
3	Self-Synchronizing TDMA	102
4	Slot Synchronization	103
5	Slots' positions initialization	104
6	Pose Estimation	127
7	Switch between horizontal and vertical orientation	130

Chapter 1

Introduction

This thesis proposes a new, automated approach to persistent data collection and data retrieval in maritime domains. Maritime data collection consists of measuring the spatial and temporal distribution of physical, chemical and biological parameters of the water and of submerged objects. Data retrieval consists of relaying the measured parameters back to the user.

Our approach integrates new sensors (underwater accurate color imagers), traditional sensors (chemistry, physical properties), wireless underwater sensor networks (WUSN) and autonomous underwater vehicles (AUV). We use the sensors for collecting the data. We use the WUSN operating in synergy with the AUV for storing the data and for retrieving the data in real time.

The central novelty of our approach is the extensive use of visible light for information transfer underwater. We use light for (1) data collection, (2) wireless data transmission, (3) autonomous control.

Visible light constitutes a narrow band of the electromagnetic spectrum that can travel through water with significantly lower attenuation than the rest of the electromagnetic spectrum ([Figure 1-1](#)). We use green light for data transmission and optical signaling as it provides the highest overall figure of merit. We use an adaptive light for underwater imaging.

Compared with previous approaches, our system is characterized by several novel



Figure 1-1: Water attenuation coefficient for electromagnetic waves [57]

and desirable properties (Section 1.4): (1) it provides temporal and spatial coverage needed for marine studies, (2) it provides the user with real-time data, (3) it is characterized by a high degree of autonomy in deployment, operation and retrieval and (4) it is easy to deploy from small boats at a low cost.

To evaluate our approach we built hardware, developed supporting algorithms and performed ocean experiments.

1.1 Motivation

Our work is motivated by the importance of ocean studies to humans. The ocean studies are important for many applications including understanding natural global processes (e.g., weather phenomena, oxygen and carbon dioxide cycles, global warming, tsunamis), supporting commercial activities (e.g., water transportation, off shore oil drilling, harbor and border security) and advancing marine biology (e.g., understanding life origin, evolution, species behavior).

1.1.1 Physical Oceanography

Oceanographers look at large scale phenomena that involve the ocean (e.g., currents, underwater tectonic activity). They need spatially and temporarily dense physical data samples to create and verify models for such phenomena. A system that provides a continuous stream of data would contribute greatly toward developing accurate models of the ocean and would lead to meaningful adaptive sampling strategies.

1.1.2 Marine Biology

Most biological and behavioral marine studies require in situ long term data collection. Such studies include determining the fragile balance and interdependencies between species, determining the effects of natural and human generated pollutants on various species' development, tracking marine mammals and understanding their life cycle, understanding the impact of global warming on marine flora and implicitly the future of the oxygen and carbon dioxide cycles.

1.1.3 Ship Hull Inspection

Ships require regular inspection of their hull. This is performed for maintenance and, in some cases, before entry in ports. The inspection is typically performed by divers or by taking the boat out of the water. An automated inspection would save time and cost. Such a system requires precise navigation relative to the ship hull, while acoustic and visual imaging is performed.

1.1.4 Offshore Oil Exploitation

The offshore oil industry relies heavily on marine sensing and data collection for the prospection of new oil fields and for daily operations. Finding oil fields requires large scale ocean floor surveys for chemical and seismic profiling. For chemical surveys, chemical sensors are towed across the area of interest by boats. For seismic surveys, vibration sensors are manually deployed across the area of interest and controlled

explosions are used to measure the shock wave propagation through the rock under the ocean floor. Subsequently, during the exploitation of the oil field, constant monitoring of the wells, pipes and other underwater equipment is required. Today, this is typically done by divers and by human operated instruments and vehicles.

1.1.5 Harbor and border protection

Harbor and oceanic border protection is a daunting task due to its wide geographical extent. Today, it is performed by human operated boats, planes, and radar. Alternatively, permanently deployed underwater sensors could improve coverage and the quality of the surveillance information. For example, engine and propeller noise could be recorded and used to identify boats. Through triangulation, the boats can be localized and investigation crews can be properly dispatched.

1.2 Challenges

Water covers approximately 70% of our planet. The volume of water in the oceans is far bigger than the volume of land above ocean level. The sheer size of the ocean poses challenging problems to modeling its physical, chemical and biological state, in a comprehensive way. On land and in air, remote sensing has mitigated most of the scale problems. Today, satellites provide precise measurements of many physical parameters (e.g., temperature) of the Earth surface and atmospheric conditions. However, this is not the case yet for the ocean. The impedance mismatch at the water surface leaves very little information “escape” from below the surface. Moreover, the water column is a highly heterogeneous environment. Understanding the phenomena that takes place in water depends on understanding the distribution and mapping of water’s properties. Measuring water properties thus requires wide geographical presence of measuring instruments. Even local studies (e.g., reef biology studies around the costal water of a island) require a significant coverage area, which is typically beyond what can be done manually.

Another obstacle in obtaining meaningful data about the ocean is the time scale.

Some phenomena (e.g., temperature or chemical variations) happen very slowly. Other phenomena (e.g., tsunamis or coral reproductive cycles) happen very quickly and their timing is hard to predict. In both instances what is required is long term presence and the ability to react quickly to changing conditions.

The large geographical and temporal scale problems on land have been addressed with wireless sensor networks (WSN). WSNs are composed of many static or mobile sensor nodes distributed over an area of interest. Each WSN sensor node is capable taking measurements, storing data and communicating wirelessly with other nodes in the network. Communication is important for tasks such as localization, data retrieval and network status updates. However, communication is very challenging in water. Electromagnetic waves are heavily attenuated by water and they cannot be used for communication beyond distances on the order of meters. Acoustic waves are typically used for communication underwater but they have orders of magnitude lower data rates and much higher energy per bit requirements than electromagnetic communication in air. This adds strain to the limited available power of the sensor nodes. Relative to the speed of light, acoustic waves are very slow, inducing significant delays in data links that cannot be ignored when designing network protocols. Additionally, the attenuation of electromagnetic waves renders GPS unavailable underwater. Without GPS, localization and navigation of autonomous robots are difficult problems.

Last but not least, water is a very harsh environment. Most metals are corroded heavily by salt water. Salt water is a good electric conductor, so all electronics have to be protected in water-tight cases. In situ experiments are difficult to perform and incur the overhead of cleaning the equipment after each experiment. Since instruments cannot be opened and debugged in situ, debugging is difficult.

1.3 Problem Statement

We describe the requirements for an ideal system that would enable pervasive underwater data collection and the creation of ocean observatories:

- The system should be capable of measuring physical and chemical properties of water and underwater objects (e.g., coral heads, fish, wrecks, pipes, boats). Typical water properties include water temperature, conductivity and dissolved gases. For submerged objects imaging is the typical data collected.
- It should be able to collect the data over a distributed geographical area. A single statically placed sensor has too narrow of a footprint. Multiple sensors have to be distributed across the area of interest. When the temporal component is not a significant constraint, we can instead use a smaller number of mobile nodes.
- The data collection should be persistent. We require our system to be able to collect data over extended periods of time (e.g., on the scale of months) without human intervention.
- The data should be available to the user with low or no latency. An ideal system reports the data to the user in a timely fashion. This enables the user to react in real time to sensor failure to adjust the sensing strategy in order to better capture the relevant data.
- The system should require minimum human effort for tasks such as: deploying the sensors, determining their location, collecting data, retrieving the data during operation and recovering the sensors after the experiment.

1.4 Related Work

There are several classes of data collection techniques currently in use.

Boat-deployed instruments are still the most often used method for covering large areas of ocean. In a typical scenario a human operated boat travels in a raster fashion, covering the area of interest . The boat can either tow instruments such as CTD (conductivity, temperature, depth), magnetometers, imaging sonars or stop and lower instruments for sampling the water column. The main advantages of this

approach are the immediate availability of the data and the wide areas that can be covered due to the mobility of the boat. The disadvantages include high cost and high time consumption, especially if the measurements have to be repeated over time.

A different technique uses statically deployed sensors connected to data loggers. These sensors are deployed for long periods of time (months) and are often used in coastal waters where they can easily be retrieved. They require no human intervention except for deployment and recovery. The main advantage is the simplicity of use. The disadvantages include the lack of spatial distribution (i.e., one instrument can collect data in a single spot) and lack of real time feedback (i.e., if the sensor fails, the user learns about the failure only after the sensor is retrieved). Recent work proposes the use of acoustic communication for creating a wireless sensor network in order to address some of the disadvantages of independent sensors (see [Section 3.1](#)).

Finally, remotely operated vehicles (ROV) and autonomous underwater vehicles (AUV) are currently used for collecting data in cases where boat deployed instruments cannot be used effectively (e.g., deep ocean floor surveys). ROVs are operated from the surface ship and have the advantage of the human operator intelligence and unlimited energy supplied by the ship. ROVs have the disadvantage of limited mobility due to the umbilical cable. AUVs perform autonomous missions and require little or no supervision during their mission (see [Section 4.1](#)). An important problem limiting AUV operations is navigation. Due to the lack of GPS, AUVs rely on external acoustic localization systems (notably LBL¹) or inertial-acoustic dead reckoning (which drifts and is expensive), for localization. Most AUVs are also limited by their energy supply. Current AUVs can execute missions on the order of hours which is not always enough for ocean time scales. A special class of AUVs that addresses the endurance problem are the sea gliders [[102](#), [39](#)]. Sea gliders change their buoyancy and glide between the ocean's surface and ocean's bottom, instead of using propellers for locomotion. This zig-zag motion requires very little energy, enabling the gliders to execute missions of up to 6 months duration. The main disadvantage of gliders

¹Long base-line navigation (LBL) relies on acoustic beacons that respond to acoustic interrogation for round-trip range measurements. The beacons are manually deployed and localized, making the system extremely cumbersome

is the limited speed in the XY plane (i.e., to move against the currents) and their inability to follow goal-directed trajectories.

1.5 Approach

We propose a system composed of a network of static and mobile sensor nodes. The static sensor nodes (our prototype is called AQUANODE) have no means of locomotion. They are deployed to static locations but can be moved by the mobile nodes. The mobile nodes (i.e., the robots, our prototype is called AMOUR) are capable of locomotion and autonomous navigation. All nodes are fitted with sensors (which are application dependent) and are capable of reading the sensors and storing the data. All the nodes are capable of acoustic and optical wireless communication. Acoustic communication is used for signaling, ranging and status updates. Using pairwise range estimates the AQUANODES are capable of creating a coordinate system and track their locations in this system of coordinates. Given the estimated ranges between the sensor nodes and the robots, the robots compute their location in the same system of coordinates and use this information for navigation. The optical communication is used as a point to point link between a sensor node and a robot. This is used for real time data retrieval. At regular intervals, a robot visits the sensor nodes one by one and downloads their data optically. The robot is used as a data mule. The robots can also carry more advanced sensors (e.g., accurate color imaging AQUALIGHT) that might be too expensive to deploy in each sensor node or might not need to be continuously deployed. The robots can use these sensors to carry out periodic surveys. In addition the robots are able to transport, deploy and recover sensor nodes, facilitating network management. The user can direct the robots to re-deploy nodes, enabling adaptive sampling. [Figure 1-2](#) presents our vision graphically. [Figure 1-3](#) presents the hardware we built to evaluate our approach.

Consider the scenario of studying the effect of phosphates and nitrates on a coastal reef's health. The phosphate and nitrate salts are carried by rain from land surface

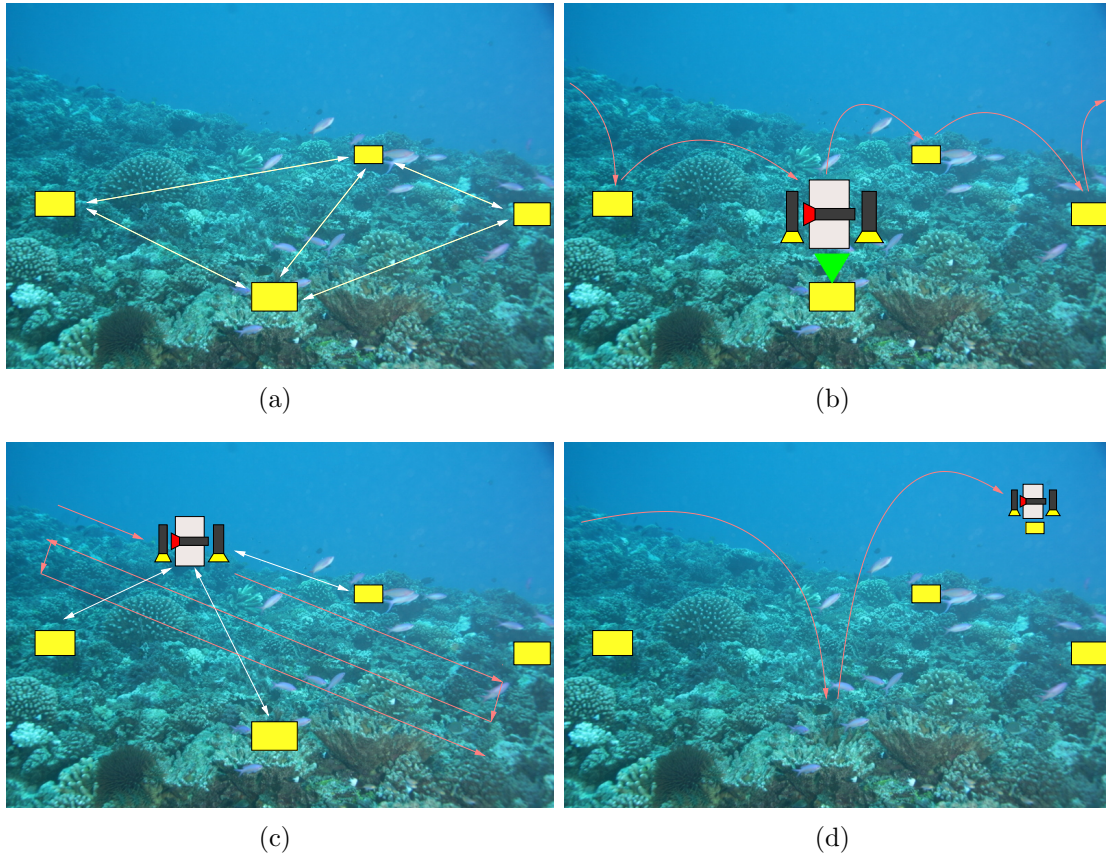


Figure 1-2: System concept. (a) The AQUANODES are deployed in the area of interest. The AQUANODES collect data, measure pairwise distances and establish a local coordinate system. (b) AMOUR acts as a data mule, visiting the nodes and downloading the data optically. (c) AMOUR carries out a visual survey of the area. It uses the AQUANODES as beacons for position information. (d) AMOUR repositions the AQUANODES for adaptive sampling.

to coastal waters. The sensor nodes (fitted with the relevant chemical sensors) are deployed in the area of interest (the reef lagoon along the coast) by manually dropping them from a boat. The nodes establish and ad-hoc network (using acoustic communication), self-localize (using inter-node ranging) and start collecting the data. When significant chemical activity is detected, the network signals the base through acoustic multi-hop telemetry. A robot is guided by the network to the event location for a complete, high density chemical survey of the area and for imaging of the coral heads. Using acoustics, the robot uses the network nodes as beacons for position estimation and navigation. This mission can be repeated at regular intervals (e.g.,




Figure 1-3: The hardware instantiation of our approach. It includes the robot AMOUR having attached the underwater accurate color imager AQUA-LIGHT, and the static underwater sensor nodes AQUANODES.

every day). The robot carries an accurate color imager, and reaches and photographs the same coral area in each mission, using its relative position to the sensor nodes. The evolution of the coral health can be assessed from color images. Additionally the robot can visit each sensor node in the network to download the collected data using the short-range high speed optical link. The hovering ability and maneuverability of the robot are essential for this task. The optical link assure a fast, low power data transfer, without requiring the nodes to be removed from their locations. Following the analysis of the muled data, the researchers may decide to reposition some of the sensor nodes for enhanced sensing. In a second mission the robot can dock with and pick up the sensor nodes from the chemically inactive areas to relocate them to the more dynamic areas, thus increasing the sensing density. These activities can be repeated over many months with very little human interaction. For the duration of the study, the scientists are provided with a quasi real time stream of data. This enables quick detection of sensor failure, procedural errors and adaptive data collection. At the end of the study the robot can autonomously retrieve the sensor nodes and return the sensor nodes to the mission headquarters.

This thesis develops three components that are critical to such maritime surveillance scenarios: an underwater sensor network and supporting communication and localization algorithms, an autonomous robot and supporting algorithms and a color-accurate imaging device. The components have been prototyped and deployed in several field experiments. The duration of the experiments was on the order of hours. During these experiments we: (1) tested the ability of the sensor nodes to establish an ad-hoc acoustic network, measure pair-wise distance, self-localize and signal, (2) tested the effectiveness of optical communication for underwater data download, (3) tested the ability of the robot to navigate while being tracked by the sensor network, (4) evaluated the color accuracy and visual quality of the pictures taken with the underwater imaging device, and (5) tested the ability of the robot to carry and use the underwater imaging device. We evaluated the acoustic network’s performance by measuring the packet-success rates and the maximum communication range, the self-localization and tracking algorithms’ precision against GPS ground-truth and the power requirements of our acoustic modem. We measured our optical modem’s communication range, data rate and compared its power requirements with that of traditional acoustic communication. To evaluate the performance of our robot we measured its speed, power consumption, single mission endurance, ability dock, pick up and carry efficiently sensor nodes, ability to carry our underwater imager. To evaluate our color-accurate imager’s performance we took pictures of a color palette and of real-world underwater scenes. We compared these pictures quantitatively and perceptually with the ground truth pictures (taken in air) and with picture taken using current techniques.

The central element in our system is the use of light for data sensing, data retrieval and underwater manipulation of the sensor node. We developed an underwater adaptive illumination system for taking accurate color images of underwater objects. This active illumination approach enables any standard camera to take accurate color pictures without any post-processing. This is not possible with previously existing techniques. Light combined with the mobility of our robot is also used for data

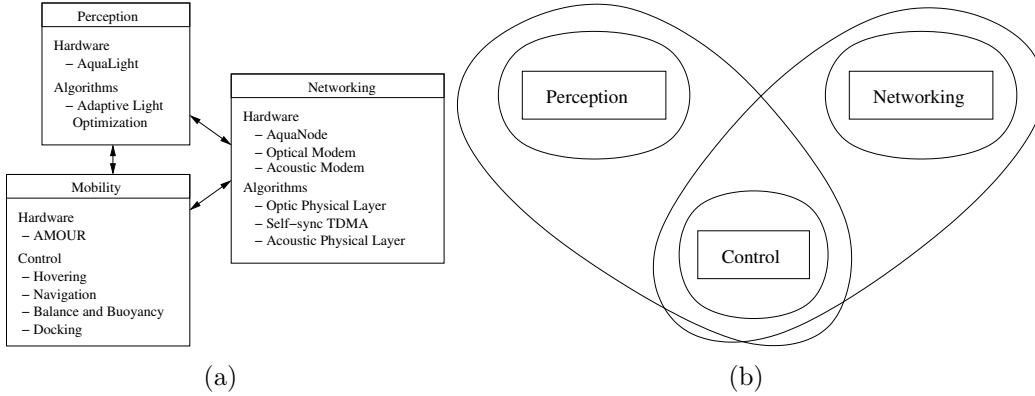


Figure 1-4: (a) Summary of the thesis contributions. (b) Experimental validation of the individual components of the surveillance system and the integration between the components.

downloading. Our optical communication system provides three orders of magnitudes higher data rates than acoustic systems, and enables scalable, low power and real time data downloads without removing the sensor nodes from their location. This is simply not possible with acoustic communication (see [Section 3.2](#)). At the same time, since data is downloaded optically, the bulk of the acoustic bandwidth is freed enabling power efficient acoustic data transmission techniques and medium access control protocols. Our acoustic modems trade data rates for power efficiency, thus reducing the strain on the sensor nodes energy supply. Finally AMOUR uses optical guidance for underwater manipulations. Our sensor nodes are fitted with optical beacons that enable the robots to autonomously dock to them, pick them up and transport them.

In summary, the use of optical communication and signaling is the catalyst that enabled the synergy between the complementary features of AUVs and WUSN. The advantage of WUSN is the ability to remain deployed and collect data over long periods of time. But, underwater acoustic communication is very inefficient (in terms of data rate and energy consumption), hindering the potential of WUSNs if used alone. By using optical communication, the energy allocated for communication can be dramatically reduced and instead be used for data collection. AUVs which have a renewable source of energy (e.g., can recharge after each mission) are the ideal

candidate for data muling. At the same time, AUVs benefit from the network by receiving localization information which greatly simplifies their navigation.

1.5.1 Contributions

1. A novel approach to underwater persistent sensing, surveillance and monitoring based on integration of new sensors, wireless underwater sensor networks (WUSN) and autonomous underwater vehicles (AUVs). Light is used as a core element for sensing, communication and manipulation.
2. A novel approach to short range underwater wireless communication based on green light. This approach is characterized by by three orders of magnitude higher data rates and three orders of magnitude lower energy per bit compared with traditional acoustic data transmissions. In addition, optical communication requires significantly lower computational power and can be implemented on much simpler electronic circuitry.
3. A novel approach to underwater color imaging. This approach is based on distance dependent spectrum-adaptive illumination, providing a solution for the long standing problem of obtaining color-accurate images underwater.
4. An analysis of underwater acoustic communication in shallow waters (i.e., in a severe multi-path environment). Based on this analysis we developed a novel physical layer (sparse symbols) and a medium access protocol (self-synchronizing TDMA) for low power acoustic communication in shallow waters. This approach trades off data rates for lower energy per bit requirements for acoustic communication.
5. The design, fabrication and experiments with an underwater Xenon flash, capable of an adjustable output spectrum. This device was used to test the adaptive illumination approach to color accurate imaging underwater.
6. The design, fabrication and experiments with an underwater optical modem, based on 532nm green light.

7. The design, fabrication and experiments with an power efficient, flexible underwater acoustic modem.
8. The design, fabrication and experiments with an underwater sensor unit capable of sensing, computation, data storage, acoustic and optic communication.
9. The design, fabrication and experiments with an underwater robot (AMOUR) capable of autonomous navigation, hovering, acoustic and optic communication with a sensor network, autonomous docking and pick-up of sensor modules, efficient transport of dynamic payloads (through internal buoyancy and balance change).

The thesis contributions and experimental validation are summarized in [Figure 1-4](#).

1.5.2 Application Limits

We chose the design and the scale of our system based on our primary application: reef monitoring. In general our system can be used in any monitoring application within shallow water (less than 100m deep), moderate areas (on the order of few km).

Maximum depth is limited by the ability to build and deploy the static sensor nodes. As the depth increases, the length of the mooring line increases as well, requiring the nodes to have increased buoyancy (and consequently bulk). If long mooring lines are used, currents are important, since they affect the position of nodes and this has to be explicitly taken into account for localization. As the depth increases the ocean gliders become a better alternative as they can cover very well the water column without mooring lines.

The maximum practical area for our system is limited by the acoustic communication range and by the number of nodes that need to be deployed. The communication range puts a lower bound density of nodes, in order to have a network. This constrains the number of nodes to grow proportionally with the area to be surveyed, which may

be too costly. As the area of interest increases, AUVs or boat-operated instruments become a better alternative, especially if the required temporal data-density is not very high.

1.6 Outline

The rest of this thesis is organized as follows. [Chapter 2](#) presents our novel approach to underwater accurate color imaging AQUALIGHT, its implementation and evaluation. [Chapter 3](#) presents the wireless underwater sensor network AQUANET, the acoustical and optical networking protocols and their experimental validation. [Chapter 4](#) presents the design, control and experiments with our autonomous underwater vehicle, AMOUR. Finally, [Chapter 5](#) presents an overview of what has been achieved, our conclusions, and future directions of research.

Chapter 2

Underwater Perception:

AQUALIGHT

This chapter presents a novel approach to underwater color-accurate imaging, based on adaptive illumination. Taking pictures underwater with accurate colors is a long standing problem, since the first underwater color picture ([Figure 2-1](#)). Water heavily distorts color, by absorbing long light wavelengths more than shorter wavelengths ([Figure 2-2](#)). Adaptive illumination uses a controllable artificial light source to illuminate the subject of interest. The spectrum of this device is adjusted to compensate for the color loss as the light travels through water. The compensation is computed based on known water optical properties and distance between the camera and the subject. We define the accurate color of an object as the color captured by a camera if the picture were taken in air, under sunlight.

2.1 Motivation

Humans use the visual system as the main information source. We gather more information (and faster) from a picture than from any other data source. Our visual system is so efficient that we try to present graphically most of the data we collect — we use graphs, maps or pictures. Studying the underwater environment is no different. There are many sensors available to measure water’s physical and chemical

properties. However, actual images of the area or object of interest are often the most desirable data.



Figure 2-1: Color underwater photography was born with this picture by Charles Martin and W.H. Longley and published by National Geographic in 1926

For marine biologists color is a very important dimension of images. Accurate color imaging enables easier species identification [21] for population count or biodiversity studies. Color is also a very reliable indicator for the state of health of coral habitats. It is used as such in many studies [62], enabling large geographical scale environmental assessments [105]. Recent studies analyze the color in the context of genetic, evolutionary or ecological characterization of coral reef fish [71, 72]. It is also known that color plays an important role as means of communication in some species of fish [12] and that color is a good indicator of behavioral and development patterns [28].

Underwater archeology can also benefit significantly from color imaging. Recent expeditions have used color photography and photo mosaicking to add completely

new visual dimensions to famous wrecks such as the 4000m deep Titanic [41, 42], or 2000 years old Roman archaeological sites in the Mediterranean Sea [104].

Marine geology is another field which can rely on color for assessing rock and sea floor composition, structure and age [48].

Accurate and consistent colors are also a significant aid to vision-based servoing and navigation [41]. Water is so distortive that even black and white pictures are affected — for example red or warm colored landmarks are perceived increasingly darker as the distance between the camera and the object increases while green or blue colored items are perceived as constant levels of gray as distance increases. This perceived color change affects negatively the performance of a system based on visual feature tracking or landmark based localization.

Another interesting application of underwater photography is automated ship hull inspection. All ships must have their hulls inspected periodically or before entry in many US ports. The operation is typically performed by in water human divers, or on the ground, by taking the boat out of the water, and it is expensive and time consuming. Recent work proposes the use of ROVs and AUVs and remote visual and acoustic imaging [76, 94, 118]. True color imaging would simplify the task of the human operator, making it easier to identify potential problems like bio-fouling, threatening devices and rusted or leaky areas.

Last but not least, there is the unquestionable beauty of the underwater landscape and wild life, which has been captured by many photographers and cameramen over time [25, 33]. Most photographers either use special techniques to capture some of the color of underwater scenes or, in many cases, they use the heavy hue shift as an artistic advantage. However, a method for reliably capturing the true vivid colors of sub sea scenes and life could open new doors to artistic expression.

Taking pictures underwater is not a trivial task. While air is an almost ideal medium for visible spectrum propagation (which makes possible high resolution pictures from even hundreds of kilometers away, e.g., imaging satellites), water strongly affects light propagation, drastically limiting the scope of underwater imaging [20, 19,

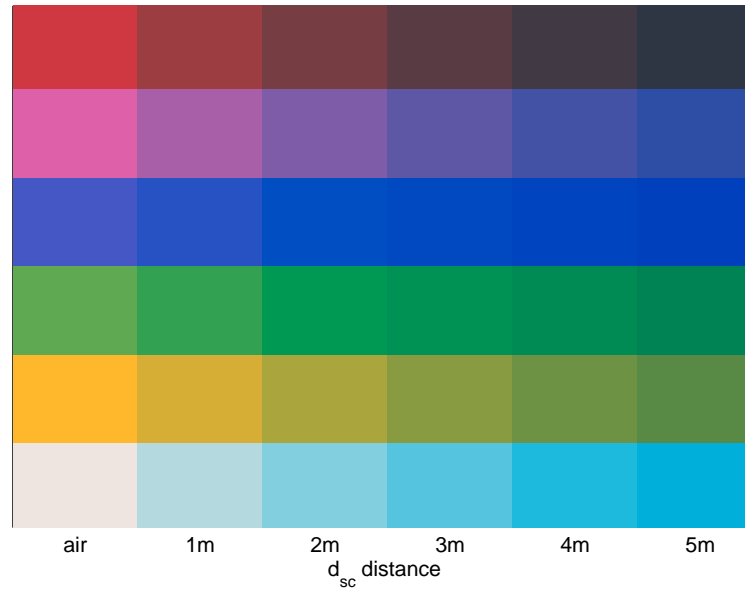


Figure 2-2: Simulated effect of water on colors. The first column presents 6 colors from the X-Rite ColorChecker Chart [2] as they appear in air. The subsequent columns show the same colors, how they would appear in water as the distance between the camera and the subjects increases. The blue and the green are not affected much, while white, yellow, pink or red are significantly shifted.

[106]. The two most important effects of water on light propagation are absorption and scattering.

The molecules of pure water and substances dissolved in salt water absorb different wavelengths (colors) of light at very different levels [106, 103, 19, 107]. This absorption is expressed as a strong hue shift in any underwater picture. Water absorbs warm colors — long wavelengths (red, yellow) more strongly than cold colors — short wavelengths (blue, green). This is why images shift toward blue/green. Moreover, since the absorption law is exponential, the longer the distance light travels through water the stronger the hue shift. Images of objects further than $1m$ from the camera have significantly desaturated colors. Images of objects further than $2m$ have very little red content remaining.

The second effect — scattering — is caused mainly by the particles suspended or dissolved in water. Scattering is responsible for the fuzziness and the “milky” look of

underwater pictures [74]. Our method does not compensate for this effect, but can be combined with techniques described in previous work [70, 74, 75] to obtain accurate color pictures in turbid waters.

2.2 Related work

Underwater photography is an old field that started with the famous photograph by Marin and Longley in 1926 (Figure 2-1) [45], and was widely expanded and brought to the public by the work of the famous oceanographer Jacques-Yves Cousteau (1950s — 1970s). Despite this flurry of activity, reproducing accurate colors at distances beyond $1m$ has remained an elusive goal.

When colors are important, photographers take the picture from a very short distance to the subject ($0.5m$). The picture is either taken very close to the ocean surface or by keeping a large light source just as close to the subject as the camera [38]. This is a very limiting factor as not all subjects will let the photographer come in such close proximity (e.g., fish). Despite very wide angle lenses, some subjects are too big (e.g., wrecks) for close proximity photography. To improve the range at which colors are rendered correctly, color compensating filters have been developed for underwater photography [38, 1]. Typically they are placed in front of the camera lens. While filters generate significantly better pictures than the no-filters setup, they are limited in scope by two factors. A single filter can only compensate for a specific depth and distance to the subject. Pictures taken closer than the optimal distance will be too red. Pictures taken too far away will be too blue (Section 2.3.1). In addition it is hard to manufacture filters with the exact transmission function required (Section 2.3.2), so filters will only approximately compensate for the color shift.

Another approach, that was made possible by the digital revolution, is post-processing. In this approach, pictures are taken under standard illumination (ambient or flash) and then post-processed with software filters [126, 115, 42, 10, 21, 105, 104, 51, 56, 55]. The process is similar to the process of achieving constant white balance

in pictures taken under various illuminants (daylight, fluorescent, incandescent). Several post-processing methods have been proposed to eliminate the typically blue cast. Most of these methods rely on assumptions on properties of the picture and ignore the phenomenon that led to color shift.

The simplest assumption is retinex [65]. It assumes that in each picture there are objects that reflect the maximal amount of red, blue and green, so that each of the three color channels can be normalized by the maximum pixel value of that respective channel contained in the picture. This is believed to be very similar to what the human brain does when evaluating colors under different illuminants. There is an entire class of retinex algorithms that attempt to achieve color constancy for computer vision. A significant part of the effort deals with the weakness of the assumption and the imperfection in the image — noise can affect dramatically the normalization. In underwater photography this assumption has been used under several names including “intensity stretching” [10] or “histogram stretching” [56].

Another common assumption is the “gray world”. It assumes that the average color in a photograph should be neutral (a shade of gray). Given this assumption the picture’s pixels are normalized such that the picture’s average color is gray. This is the method favored in many underwater visual surveys [41, 42, 105, 104, 21]. The advantage of the gray-world assumption over the white retinex method is the higher noise tolerance (due to the higher number of pixels averaged together). However, the validity of the assumption remains highly dependent on the scene and subject, therefore yielding variable degrees of success. The best application for this method is in ocean floor surveying, where the large areas of sand and rock make the assumption plausible.

For specialized purposes, where a calibration color palette or objects of known color are present in the picture, various statistical or learning approaches have been applied [115]. In this case colors can be restored more faithfully, but at an additional cost: a calibration color palette has to be present in every picture close to the subject (the subject and the calibration palette must be at the same distance to the camera). In addition, manual identification of the calibration palette in every picture taken

makes surveying with large numbers of pictures labor intensive.

Finally, a recent paper by Yamashita et al. [126] considers the problem from a fundamental perspective. Instead of ignoring the knowledge about water’s optical properties and trying to correct using only the captured image, the authors model the water effects on the color using the wavelength-dependent absorption coefficients and the distance between the camera and the subject. They apply the computed inverse function to the captured image. The function they compute is a very coarse approximation that takes into account only the effect of water on three particular light wavelengths (instead of the continuous spectrum the camera captures). This approximation limits the performance and scalability (to significant distances) of this approach.

In general, there are some fundamental limits to all methods based on post-processing. First, the non-linear effect of water on color (Section 2.3.1) combined with the limited sampling of the light spectrum by the current cameras (Section 2.3.3) makes accurate color correction impossible. For example colors that are significantly different in air can be collapsed onto the same color in underwater pictures. This puts a hard limit on how accurately we can restore the color by post-processing. In addition, attenuation in the red part of the spectrum is so strong that a bit of information is lost for approximately every one meter of distance between the camera and the subject. This translates into a lower signal to noise ratio or, in the case of a processed picture, a higher noise especially in the areas with high red content. This effect is exacerbated by the relatively limited light available underwater (as it must be carried by a diver or a vehicle) and also by the limited dynamic range of current cameras.

In contrast, our method for color-accurate imaging by adaptive illumination does not have any of these disadvantages. It compensates for the light loss by using an illuminant with a radiation spectrum roughly the inverse of the water transfer function, therefore presenting to the camera’s CCD an already corrected image (e.g., as if the image were taken in air). The full dynamic range of the camera is used and post-processing is not necessary.

A recent body of research in underwater imaging, that is orthogonal to understanding color shift correction, considers the effect of light scattering by the particles suspended in water [55]. The two main approaches are structured lighting [74, 75] and range gating [70, 23]. In the case of structured lighting, the scene is illuminated by scanning with a narrow line of light, thus limiting the amount of back-scattering. Range-gating relies on illuminating the scene with a very short pulse of light combined with a precise activation of the camera. This results in capturing just the light that traveled for the desired distance — from the light source to the object and back to the camera. The back-scattered light, which travels less, is not captured.

2.3 Approach

In contrast to previous solutions, the approach discussed in this chapter is based on active, adaptive illumination. The subject is illuminated with a spectrum-controllable light source. The light composition is calculated so that the water between the camera and the subject transforms it into white light. For example, since water attenuates the red side of the spectrum the most, we can engineer light to contain more red energy at the source. As it travels through water, the red wavelengths will be attenuated, resulting in the optimal levels of red upon arrival at the camera. The exact energy content of the light source is calculated using water optical properties and the distance between the camera/light source assembly and the subject. The light is generated by a source composed of several Xenon light bulbs. The bulbs are filtered with different long-pass filters. Therefore each bulb becomes a light source with a particular spectrum. By adjusting the relative power of the Xenon flashes, we can generate a light source with variable spectrum. Sections 2.3.1-2.3.4 present the physics behind our approach, and Section 2.3.5 presents our solution to underwater accurate color imaging.

2.3.1 Optical Properties of Water

Our approach to underwater accurate color imaging relies on actively compensating for the wavelength-dependent absorption in water. This section presents optical properties of water.

The optical properties of water have been studied extensively in the past [5, 7, 9, 13, 14, 19, 20, 32, 43, 49, 54, 59, 60, 61, 73, 78, 81, 83, 84, 88, 90, 91, 92, 93, 97, 98, 101, 103, 106, 107, 108, 113, 114, 123, 124, 130, 131]. There are two important processes that happen while light travels through water: scattering and absorption.

Scattering is the physical process whereby light is forced to deviate from a straight trajectory as it travels through water. The effect is caused predominantly by solid particles suspended in water, but also by the water molecules and substances dissolved in water. Scattering is wavelength independent and does not affect the color balance of an underwater image. However, it can severely limit the practical range at which underwater pictures can be taken. In clear natural waters, visibilities of 10-15m are possible, but in turbid water, visibility can be reduced to less than 1m due to scattering.

If the light rays are deviated back to the light source, the process is called backscatter. Backscatter is a significant effect for all underwater imaging applications using an artificial light source. If the light source is close to the camera, the backscattered light is captured by the camera as an overlay, thus reducing the contrast and color saturation of the subject. The typical method to reduce the backscatter effects is to have as much separation as possible between the light source and the camera [55, 105].

The second physical process that happens as light travel through water is absorption. Water molecules and dissolved substances are the main contributors to absorption. The absorption coefficient, a_w , is wavelength dependent (Figure 2-3). Within the visible spectrum, longer wavelengths are attenuated more strongly than shorter wavelengths. This property is the main cause of color shift in underwater imaging.

The absorption law is exponential. The light energy transmitted through water




Figure 2-3: Light total absorption coefficient $a_w(\lambda)$ for water as a function of wavelength [106], plotted on a log scale. Notice the almost two orders of magnitude difference between the attenuation coefficient for blue (approx. $450nm$) and red (approx. $700nm$)

can be expressed as:

$$I(\lambda, d) = I(\lambda, 0)e^{-a_w(\lambda)d}$$

where:

$I(\lambda, 0)$ spectral power distribution at source

$a_w(\lambda)$ water absorption coefficient

$I(\lambda, d)$ spectral power distribution at distance d from the source, in water

Examples of the absorption law's effects are presented in [Figure 2-4](#).

When modeling the water effects in underwater imaging, both the light attenuation between the light source and the subject and the attenuation between the subject and the camera have to be considered. The total distance traveled by the light is:

$$d = d_{ls} + d_{sc}$$

where

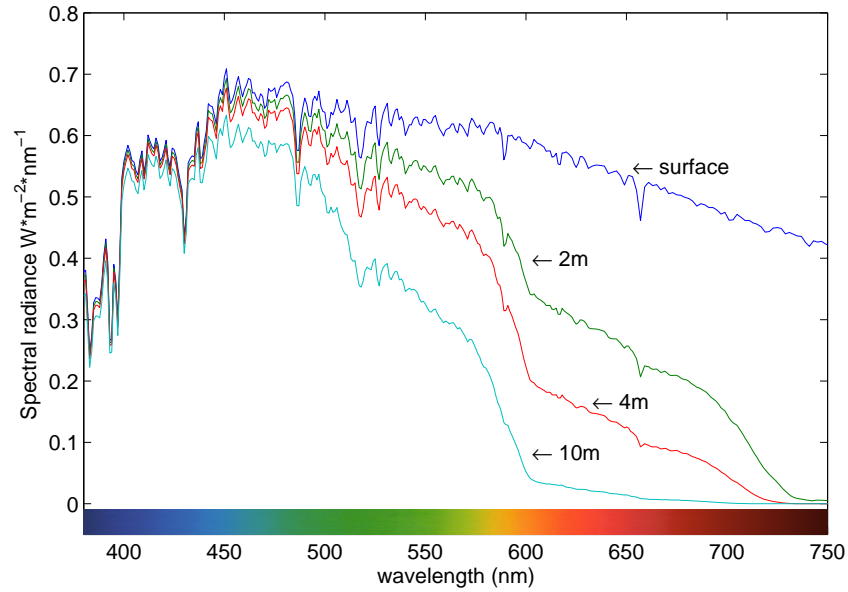


Figure 2-4: The spectral power distribution of sun light at water surface and various depth (simulated).

d_{lo} is the distance between the light source and the subject

d_{oc} is the distance between the subject and the camera

Our approach uses an artificial light source positioned relatively close to the camera, therefore:

$$d_{ls} = d_{sc} \Rightarrow d = 2d_{sc}$$

So the distance light travels through water (and is attenuated) is double the distance between subject and camera d_{sc} . Figure 2-5 presents the simulated apparent color of Munsell 5Y8/10 color sample at various distances from the camera, in water.

2.3.2 The Naive Approach to Addaptive Illumination

One direct solution to the color shift caused by light absorption in water is to create a light source that can compensate exactly for the light lost. For example, suppose we are given a distance between the subject and the camera. Suppose over this distance the 650nm red light is attenuated 5 times more than the 530nm green light. Therefore, the light source should output 5 times more power at 650nm than at

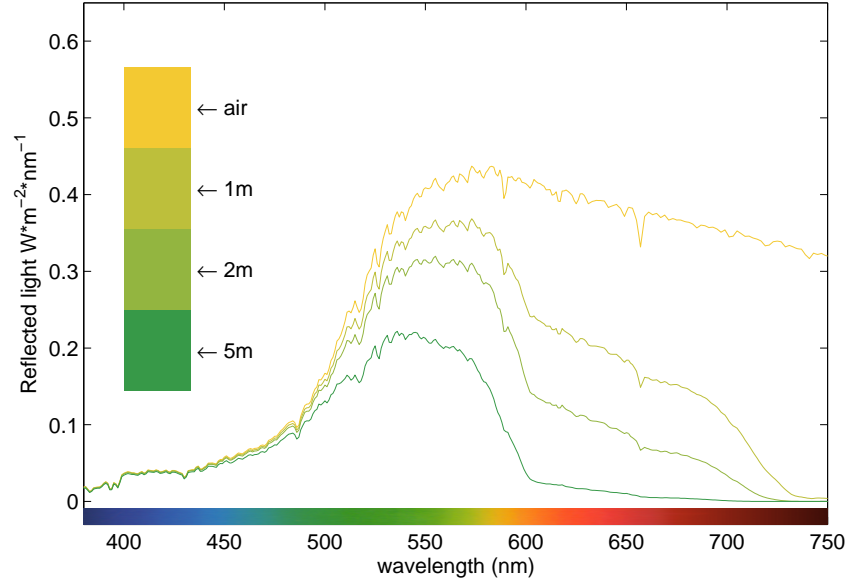


Figure 2-5: Simulated apparent color of Munsell 5Y8/10 color sample in water, at various distances d_{sc} from the camera. The color sample is assumed illuminated with a D65 (mid-day sunlight) illuminant positioned at the same distance from the sample as the camera. The color is based on the spectral power distribution of the reflected light at the point where it enters the camera lens.

530nm. In general the spectral output power of the light source should be:

$$I_{ec}(\lambda, d_{sc}) = \frac{I_{D65}(\lambda)}{e^{-2a_w(\lambda)d_{sc}}} = I_{D65}(\lambda)e^{2a_w(\lambda)d_{sc}}$$

Such a light source compensates exactly for the light loss at the specified distance d_{sc} . The subject appears to be illuminated with an ideal D65 (mid-day sunlight) illuminant in air. Figure 2-6 shows the required spectral power distribution for d_{sc} up to 3m (distance between the camera/light source assembly and the subject).

The usefulness of $I_{ec}(\lambda, d_{sc})$ in practice is very limited due to several important drawbacks. Fabricating a light source with the required spectral power distribution is challenging. Just illuminating scenes at distances of 0.5 – 3m from the camera requires light control of over 5 orders of magnitude of dynamic range. The amount of required optical power rises very sharply with distance (e.g., at a distance of 3m




Figure 2-6: Spectral power distribution for a light source that would compensate exactly for the water absorption (naive approach). The spectra are plotted on a semi-logarithmic scale due to the 5 orders of magnitude of dynamic range needed

more than 10^5 times more power is required in water than in air). All the extra power will be transformed in heat by the water it travels through.

In addition current light sources have a fixed spectrum, given by the technology and the physical properties of the materials used for their fabrication (e.g., incandescent bulbs and xenon discharge lamps have a broad spectrum uniform power output, LEDs can be fabricated for a few discrete $10 - 20nm$ wide bands). The limits in the output spectrum of existing light sources lead to the use of filters for generating the exact spectral distribution required. This increases the power requirements by another order of magnitude. Creating filters with the exact optical properties is very challenging.

These technical problems make the direct, brute force method not practical for accurate color imaging underwater.

2.3.3 The Human Visual System

Creating a light source that can compensate exactly for the light loss in water would result in accurate-color pictures (at least for relatively flat scenes, where the distance between the subject and the camera or light source is relatively constant). However this method is both wasteful in terms of power and resources, and not necessary as the human vision system is not a “perfect” light sampler. This section discusses the physics of the color vision system of humans. We use the properties of this system to create a low complexity and power efficient light source for color-accurate underwater imaging.

Creating the exact inverse of water attenuation will concentrate the majority of the energy in the deep red part of the spectrum where the water attenuation is the highest. However, it turns out that the human eye sensitivity is very low in this part of the spectrum. Very little extra accuracy would be achieved in this fashion.

Exact compensation is not necessary since the eyes (or cameras) sample the spectrum with only three discrete sensors therefore collapsing the the infinite-dimensional color space into a three dimensional space. To preserve the human-perceived color accuracy, it is only necessary to preserve the color coordinates for naturally occurring colors in the three dimensional color space of humans. This allows for some tolerance in the light source spectrum.

Under typical conditions the objects around us are illuminated by a white light source $I_0(\lambda)$, which contains energy at all visible spectrum wavelengths (Figure 2-7). The colors of an object is given by its wavelength-dependent reflection coefficient $R(\lambda)$ which represent the fraction of incident power that is reflected. For example, Figure 2-8 plots the reflection coefficient of the Munsell 5Y8/10 color sample (a saturated red). Our eyes receive the reflected light, as shown in (Figure 2-9). The spectral distribution of the reflected light is given by the following equation:

$$E(\lambda) = R(\lambda)I_0(\lambda)$$

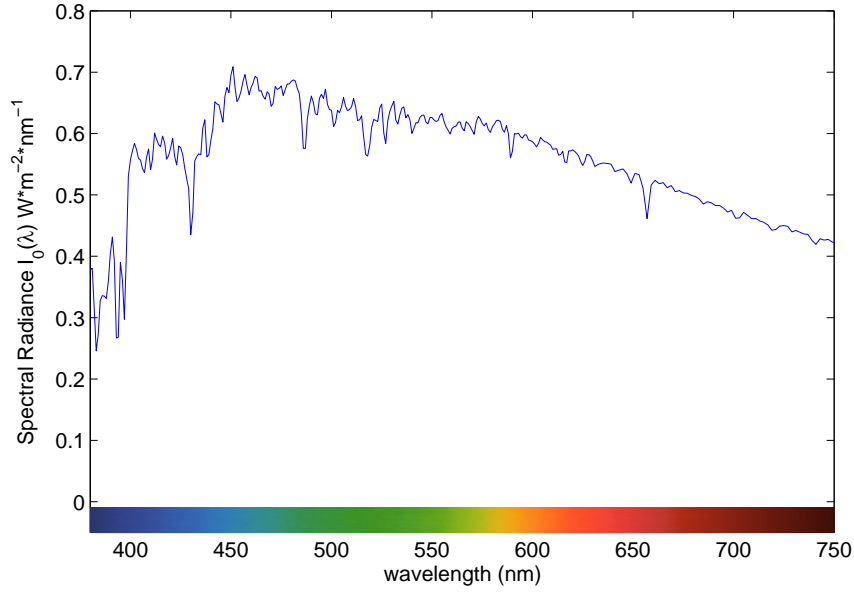


Figure 2-7: Spectral power distribution of CIE standard illuminant D65, having the approximative spectrum of mid-day sun in Western Europe.

The human vision system perceives colors by sampling the light spectrum with three different type of cone cells (ρ , γ and β). Each type of cell is sensitive to a different region of the visible spectrum — long, medium and short wavelengths. The three normalized sensitivity curves $S_\rho(\lambda)$, $S_\gamma(\lambda)$ and $S_\beta(\lambda)$ are plotted in [Figure 2-10](#). The responses of these three types of cells T_ρ , T_γ and T_β are used by the central nervous system to associate colors to objects. For example in the case of the Munsell 5Y8/10 color sample, illuminated with white light the responses of the three types of cone cells can be calculated as:

$$T_\rho^{5Y8/10} = \int_{380nm}^{750nm} S_\rho(\lambda)E(\lambda)d\lambda = \int_{380nm}^{750nm} S_\rho(\lambda)R_{5Y8/10}(\lambda)I_0(\lambda)d\lambda$$

$$T_\gamma^{5Y8/10} = \int_{380nm}^{750nm} S_\gamma(\lambda)E(\lambda)d\lambda = \int_{380nm}^{750nm} S_\gamma(\lambda)R_{5Y8/10}(\lambda)I_0(\lambda)d\lambda$$

$$T_\beta^{5Y8/10} = \int_{380nm}^{750nm} S_\beta(\lambda)E(\lambda)d\lambda = \int_{380nm}^{750nm} S_\beta(\lambda)R_{5Y8/10}(\lambda)I_0(\lambda)d\lambda$$

Therefore, the brain receives a tuple of three numbers $(T_\rho, T_\gamma, T_\beta)$ and associates




Figure 2-8: The reflection coefficient of Munsell 5Y8/10 color sample.

colors to objects based on this tuple. This results in an apparent three dimensional color space for humans. Any two objects that generate the same respective response in the three types of cone cells will appear to human observers as having the same color. This property can be exploited to reduce the complexity and power requirements of the underwater light source.

Using this observation, our key idea is to develop a light source that preserves the coordinates of most naturally occurring colors in the human three dimensional color space, when used underwater. With this approach, the coordinates are also preserved in the color space of the current cameras, since they closely model the human vision system color space.

We will use the Color Rendering Index (CRI) metric to evaluate the ability of a light source to render colors accurately.

2.3.4 The Color Rendering Index

The color rendering index is a metric for evaluating the quality of artificial and natural light sources. It measures the ability of a light source to render colors accurately,




Figure 2-9: Reflected spectral power distribution of Munsell 5Y8/10 color sample when illuminated by the CIE standard illuminant D65 (Figure 2-7)

relative to a reference light. We will use CRI as the goal function for an algorithm for underwater illumination spectrum optimization.

Artificial light sources do not have the same spectral composition as the sun light. Therefore they do not render objects' colors in the same way sunlight does. The Color Rendering Index (CRI) was developed by Commission Internationale de l'Éclairage (CIE) as a metric for comparing ability of light sources to render color accurately. The metric is based on comparing the color coordinates of 15 sample colors illuminated with the reference light and with the tested light. The color samples have been chosen as representative for natural occurring colors. The euclidean distance between the color coordinates of two colors represented in CIE 1964 $U^*V^*W^*$ color space [125] is used as a quantitative value for color difference. $U^*V^*W^*$ is a uniform color space, meaning the euclidean distance is a consistent measure of perceptual color difference, across the entire color space.

$$\text{CRI} = 100 - 4.6\Delta\bar{E}_{UVW}$$




Figure 2-10: Normalized sensitivities of the three types of cone cells present in human retina (ρ , γ and β) [110]

$$\Delta \bar{E}_{UVW} = \frac{1}{15} \sum_{i=1}^{15} \Delta E_i$$

$$\Delta E_i = \sqrt{\Delta U_i^2 + \Delta V_i^2 + \Delta W_i^2}, i = 1, \dots, 15$$

The Color Rendering Index has values between 0% and 100%. A CRI of 100% corresponds to a perfect light source which renders colors identically to the reference light source.

2.3.5 Perceptual Adaptive Illumination

Our method for taking color accurate pictures underwater is called perceptual adaptive illumination. It consists of compensating for the color shift underwater by illuminating the subject with a light source with higher power at wavelengths where water absorbs light more. The spectral power distribution depends on the distance between the camera and the subject (thus the method is adaptive), and is computed based on




Figure 2-11: The human visual system color response as modeled by the CIE 1931 2° Standard Observer [52, 53]. The 3 curves model the response of the three type of cone cell in human retina.

a priori knowledge of the optical properties of water and the color perception system of humans. The spectral power distribution is computed such that when filtered by the water between the camera and the subject, it will render the subject's colors to the human perception system as well as natural light in air. The Color Rendering Index is used as a metric for measuring the quality of the generated light. The required light spectra is generated by a variable spectrum light source composed of several basic light sources, with fixed but distinct spectral power distribution. The variable spectrum is achieved by varying the relative power of the composing light sources.

The perceptual adaptive illumination differs from the naive approach described in Section 2.3.2 in that it uses only the light needed to keep the colors unchanged for the human observer, and not all the light needed to fully compensate for the water attenuation. This decreases significantly the power requirements and complexity of the light source, while keeping the subject's colors unchanged (relative to when illuminated by a natural light source, in air), for the human observer and cameras. The power savings come mainly from a reduced power output in the deep red region

where the human eye is not very sensitive and the color rendering is not affected. The complexity reduction comes from exploiting the limited spectrum sampling by human color vision system, knowledge of which allows for less precision in the required spectral power distribution.

Our approach relies on generating the subject-camera distance dependent spectra with a source composed of n light sources with fixed output spectra $I_i(\lambda)$, but adjustable output power p_i . Thus, the resulting output spectral power distribution can be any linear combination of the composing light sources:

$$I(\lambda) = \sum_{i=1}^n p_i I_i(\lambda)$$

The component light sources' spectra can be viewed as basis functions, used to generate the required spectral power distribution. The scaling factors, p_i , correspond to the output power setting of the component light sources.

The generated light is filtered by the water between the light source and the subject and by the water between the subject and the camera. Given the assumption that the light source is positioned close to the camera, light travels through $2d_{sc}$ of water, where d_{sc} is the distance between the subject and the camera. Water attenuation makes the light that illuminates the subject equivalent with a light source in air:

$$\hat{I}(\lambda) = e^{-2a_w(\lambda)d_{sc}} I(\lambda) = e^{-2a_w(\lambda)d_{sc}} \sum_{i=1}^n p_i I_i(\lambda)$$

From the perspective of the camera (and of the observer), the subject looks as though it is illuminated by $\hat{I}(\lambda)$ in air. The problem can be posed as adjusting $\hat{I}(\lambda)$ so that colors are rendered accurately in the observer's color space. CRI ([Section 2.3.4](#)) can be used as a metric for evaluating the color rendering accuracy of $\hat{I}(\lambda)$. The $\hat{I}(\lambda)$ spectrum can be varied by adjusting the output power settings of its component lights p_i . The power settings can be found by solving the following optimization problem:

$$p_1..p_n = \arg \max_{p_1..p_n \in [0..1]} \text{CRI} \left(\hat{E}(\lambda) \right) = \arg \max_{p_1..p_n \in [0..1]} \text{CRI} \left(e^{-2a_w(\lambda)d_{sc}} \sum_{i=1}^n p_i I_i(\lambda) \right)$$

This optimization problem can be solved by any off-the-shelf numerical optimization software. It has to be solved for the given subject to camera distance.

CRI is independent of light intensity, so any scaled version of a particular set of coefficients $p_i..p_n$, by a constant factor, $kp_1..kp_n$ has the same CRI. Therefore one of the power factors can be fixed prior to the optimization, and will determine the absolute power of the light source. Algorithm 1 summarizes the computation of the optimal light for each distance in the range of interest. The Matlab implementation is included in [Appendix A](#).

Algorithm 1 Light optimization

Require: $I_i(\lambda, t_i)$ Output characteristics of the component light sources
 d_{max} Maximum distance for which the light will be used

- 1: **for** $d_{sc} = 0\text{m}$ to d_{max} in 0.1m steps **do**
- 2: compute $t_1(d_{sc})$ for optimal illumination at distance d_{sc} in air
- 3: compute $t_2(d_{sc})..t_n(d_{sc}) = \arg \max_{t_2..t_n \in [0..1]} \text{CRI} \left(e^{-2a_w(\lambda)d_{sc}} \sum_{i=1}^6 I_i(\lambda, t_i) \right)$
- 4: **end for**

The range of computable Color Rendering Index for $\hat{I}(\lambda)$ depends on the choice of the source light components $I_1(\lambda)..I_n(\lambda)$. This choice is restricted by the available light source technologies and filters and is discussed in [Section 2.4](#). The aim is to obtain a CRI of at least 90%, which would be equivalent to the best artificial light sources in air.

2.4 Hardware Description

To evaluate the light composition optimization results , a light source with variable spectral power distribution has been built. It is based on 6 discrete light sources, each with a different spectral power distribution. The output power of each component

light source can be adjusted independently. This results in an apparent variable output power for the composite light source.

Figure 2-12 shows the block diagram of the underwater flash.

We used a xenon discharge lamp as the basic light source. It has been chosen for its uniform output power across the visible spectrum, which is very similar to sun light. In addition, it is capable of the highest instantaneous output power, for its size (and thus its widespread use in general photography). Each of the 6 component light sources has a Xenon discharge lamp and a long-pass filter. Filters are needed to generate different spectra for the 6 light sources, given that the Xenon light has a fixed output spectra. Long-pass filters are used in order to improve the overall efficiency of the device. Since water has an almost monotonic increase of attenuation with wavelength, at any distance more power is needed at longer wavelengths than the shorter wavelengths. One of the flashes is unfiltered providing enough power for the short-wavelength of the spectrum. The filters were chosen empirically based on experimental observations of the required spectrum. Their final choice was based on simulation of the achievable CRI and illumination power. The cut-off wavelengths for the 5 filters are 475nm, 575nm, 600nm, 600nm, 630nm (Figure 2-16).

The output power of the Xenon lamps is not easily controllable in practice. The power depends on the size of the capacitor that will be discharged in the lamp and the initial voltage across its pins, and it varies with time. However, the discharge is very short (less than 4ms) and the entire energy is discharged while the camera is capturing the image. Therefore, controlling the total energy discharge of the Xenon lamp is a way to control the apparent power during the picture exposure. This can be done by timing the discharge. When triggered, all flashes are started at the same time and stopped according to the required energy discharge. The circuit for stopping the flashes is based on fast, high instantaneous current IGBTs. During the discharges, the current through the circuit can reach 250A.

For the duration of their discharge the Xenon lamps are powered by a capacitor bank. Each lamp has associated a 1080 μ F capacitor, for a maximum discharge energy of 33J. The capacitors are charged to 350V by an inverter. The entire unit is

powered by a 11.7V, 2.1Ah Lithium-Polymer battery which provides enough power for approximatively 300 full power flashes. The complete schematics are included in [Appendix C](#).

An NXP LPC2148 CPU together with Xilinx XC2C256 CPLD control the timing of the 6 flashes. The Verilog code is included in [Appendix D](#). The timing is based on the distance to the subject and a precomputed lookup table of distance vs. power requirements. The distance to the subject can be input to the unit through a user interface (using magnetic switches). The distance can also be automatically set by an ultrasonic pinger. Most of the pictures were taken with manually set distance, as the ultrasonic pinger that was tested (Teledyne/Benthos PSA-916) performed poorly in swimming pool conditions as well as in typical reef environment (noisy measurements).

The flash electronics are housed in a water and pressure resistant enclosure ([Figure 2-13](#)) made out of clear acrylic. Two underwater connectors are placed on the back of the enclosure. One is used for the electrical connection between the camera and the flash (for synchronization). The other is used for interface with the ultrasonic pinger, programming and battery re-charging.

The flash is attached to the camera through a rigid fixture, which aligns them ([Figure 2-14](#)).

2.5 Experiments and Evaluation

To evaluate our method for adaptive illumination, several experiments have been performed. First, the output of each flash has been measured and calibrated. Second, using the available output from the flashes, the optimization algorithm was run, to determine the subject-camera distance dependent parameters and to evaluate the theoretical performance of the flash. And finally, to evaluate the real life performance of the method, pool and ocean experiments were performed. Pictures of a color palette and of coral reef areas were taken.

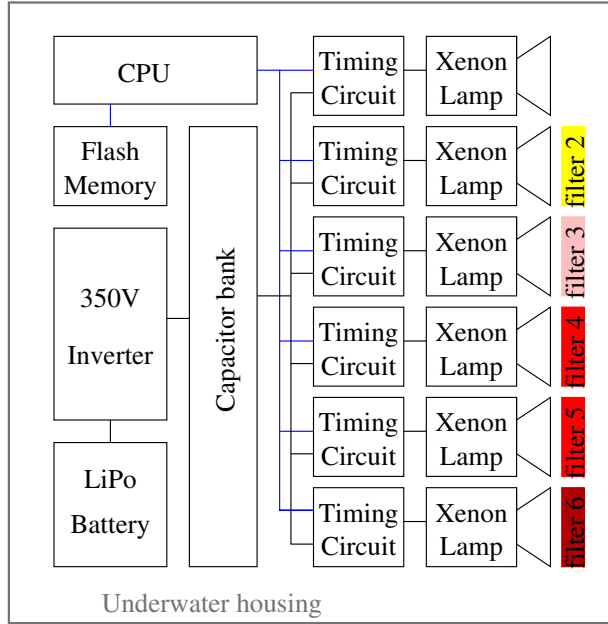


Figure 2-12: The block diagram of the underwater flash with adjustable spectral power distribution. The flash is composed by 6 independently timed Xenon lamps, filtered with different long-pass filters. One lamp has no filter, providing enough power for the blue and green part of the spectrum. A CPU adjusts the timing based on the distance to the subject and the calibration data stored in a non-volatile memory. Power is provided by a lithium-polymer cell which charges the capacitor bank at 350V through an inverter.

2.5.1 Output calibration

The goal of the output calibration is to determine $I_i(\lambda, t)$ (i.e. the spectral power distribution of the generated light as a function of the flash time) for each of the $i = \overline{1..6}$ composing lamps. The functions $I_i(\lambda, t)$ have to be determined experimentally since the characteristics of the filters used in the flash construction is only known approximately. In addition, the output energy has a hard to predict dependency of the flash duration.

The main instrument used for output calibration is a CCD based spectrometer: BWTek BRC111A. The spectrometer outputs the spectral power distribution of the light captured by its optical probe. Since the spectrometer's gain is not constant with wavelength (i.e. the ratio between optical power and spectrometer output varies with wavelength), the spectrometer output was calibrated using an output-calibrated




Figure 2-13: The underwater flash. The 6 filtered Xenon lamps are visible on the left side. The next layer is formed by the capacitor bank and the discharge control circuit. The right-most layer contains the CPU board and the Lithium-Polymer battery powering the flash. All the electronics are contained in a acrylic water tight case.

incandescent lamp Ocean Optics LS-1.

The flash was positioned facing the spectrometer's probe, $2m$ apart (Figure 2-15). Each filtered Xenon lamp was fired separately for time intervals of $0.06ms$, $0.125ms$, $0.25ms$, $0.5ms$, $0.7ms$, $1ms$, $1.5ms$, $2ms$, $3ms$, $4ms$. The output spectrum was captured for each of these cases. Figure 2-16 plots the measured relative spectra of the 6 flashes, when fired for $0.5ms$. The Xenon lamps generate approximately the same amount of power, and the same spectral distribution (observable at wavelengths longer than the filters cut-off wavelengths). The filters used have a very sharp transition band, very high attenuation in the stop-band, and high transmissivity in the pass-band.

Figure 2-17 plots the measured spectral energy distribution of flash F1 (unfiltered) as a function of flash time. Figure 2-18 plots the output energy of flash F1 as a function of flash time.




Figure 2-14: The flash and camera assembly. The camera is a standard digital SLR (Canon 30D) in an underwater case. The camera and the flash are connected electrically (to synchronize the flash firing), and are rigidly attached to an aluminum bar. A flashlight points in the same direction. It illuminate the subject to help the camera's auto-focus system in low ambient light.

Using the measured output spectra we construct the functions $I_i(\lambda, t)$ as piecewise polynomial interpolations of the measured data.

2.5.2 Optimization

Using the spectral energy distribution of the 6 unit flashes the optimal light composition has been determined with the algorithm described in [Section 2.3.5](#). We posed the optimization problem in terms of the flash firing durations $t_1..t_n$ eliminates the calculation the inverse of time to power dependency measured in [Section 2.5.1](#).

$$t_2..t_n = \arg \max_{t_2..t_n \in [0..1]} \text{CRI} \left(e^{-2a_w(\lambda)d_{sc}} \sum_{i=1}^6 I_i(\lambda, t_i) \right)$$

Only the firing durations $t_2..t_n$ are computed through the CRI optimization. The duration t_1 , which corresponds the the unfiltered flash, was used as a scale factor for the total output power of the flash. Its value is based on the necessary illumination for the camera to expose the picture properly. The necessary illumination level depends



Figure 2-15: The flash output calibration test-bench. The flash is triggered in front of spectrometer’s cosine corrected optical probe. An optical fiber relays the light captured by the probe to the radiometric spectrometer. The black background prevents reflections from nearby objects from altering the measurements.

on the CCD sensitivity, the gain factor, and the lens aperture. Since the flash angle is constant and the light spreading is constant, the power required to achieve the constant illumination level is proportional with the square of distance between the light source and the subject (d_{sc}^2). The illumination level necessary was determined experimentally, in air. The measurements remain valid for underwater pictures as well, as the flash F1 generates the blue light necessary for the picture, which propagates almost unattenuated in water. The flash calibration and optimization is summarized in Algorithm 2.

The optimization results are presented in Table 2.1. These settings result in the output spectra of the adaptive light source presented in Figure 2-19. The output power in the red region of the spectrum grows significantly with the distance. Figure 2-20 shows the spectra of the adaptive flash after being filtered by the water between the subject and the camera. The spectra at all distances between 0m and 5m are very similar – which means colors are rendered in similar manner.




Figure 2-16: The normalized output spectral energy distribution of the 6 flashes. F1 is unfiltered, with the approximative spectrum of daylight sun. F2-F6 have increasingly long cut-wavelength to compensate for the water absorption which increases with wavelength

d_{sc}	F_1	F_2	F_3	F_4	F_5	F_6	CRI
1m	0.0652	0.0227	0.0405	0.0285	0.0272	0.0188	99.26
2m	0.0652	0.0317	0.0789	0.0710	0.0517	0.0302	98.56
3m	0.0652	0.0421	0.1798	0.1518	0.1425	0.1090	97.97
4m	0.0652	0.0475	1.0000	1.0000	0.9802	0.6584	96.16

Table 2.1: The fire duration for the 6 flashes for subject to camera distances between 1m and 5m and their expected CRI.

The flash durations are stored in the flash memory and used according to the distance between the subject and the camera.

2.5.3 Pool experiments

We evaluated the real life performance of the adaptive illumination approach in a first set of experiments, in a swimming pool. Pictures of a color palette with a 15




Figure 2-17: The normalized output spectral energy distribution flash F1 (unfiltered) as a function of flash time

different colors spread over the color space were taken at $1m$, $2m$, $3m$, $4m$ and $5m$. For each distance, three pictures were taken: using ambient light, using the standard white flash and the adaptive light as computed in Section 2.5.2. Subsequently, the pictures taken with the white flash were post-processed by scaling the color channels in order to bring the white color patch to a neutral color. This method is similar in results with white retinex methods (Section 2.2), but exploiting the fact that it is known what part of the picture is supposed to be white. The resulted (cropped)

Algorithm 2 Flash calibration and optimization

- 1: **for all** $i = 1$ to 6 **do**
 - 2: **for all** t_m in 0.06, 0.125, 0.25, 0.5, 0.7, 1, 1.5, 2, 3, 4 ms **do**
 - 3: measure $I_i(\lambda, t_m)$ using a spectrometer
 - 4: **end for**
 - 5: create $I_i(\lambda, t)$ as piecewise interpolation of the data collected above
 - 6: **end for**
 - 7: **for** $d_{sc} = 0m$ to $5m$ in 0.1m steps **do**
 - 8: compute $t_1(d_{sc})$ for optimal illumination at distance d_{sc} in air
 - 9: compute $t_2(d_{sc})..t_n(d_{sc}) = \arg \max_{t_2..t_n \in [0..1]} \text{CRI} (e^{-2a_w(\lambda)d_{sc}} \sum_{i=1}^6 I_i(\lambda, t_i))$ using a non-linear solver
 - 10: **end for**
-




Figure 2-18: The normalized output energy of flash F1 (unfiltered). The rest of 5 flashes have very similar profiles

pictures are presented in [Figure 2-22](#).

The perceptual difference is very noticeable. The ambient light pictures have very little saturation and almost no content of warm (long wavelength) colors. Since they were taken at 5m depth, the ambient (sun) light was attenuated by 5m of water by the time it reached the subject. In general the colors in ambient light pictures depend on the sum between the depth at which the subject is located and the distance between the subject and the camera, as this the total distance the light travels through water.

In the pictures taken with the white flash there is a noticeable decrease in color saturation as the distance between camera and subject increases. Red becomes black and yellow becomes a shade of green. Blue and green however are not affected very strongly, as expected.

In the post-processed pictures, some of the colors are recovered even at longer distances, but the color saturation is reduced. Some colors, especially the shades of red are still far from the originals. This is due to the non uniform effect of water on




Figure 2-19: The normalized total output energy of the adaptive flash, set for subject-to-camera distances of $1m$, $2m$, $3m$, $4m$ and $5m$.

colors. Different colors are shifted by different amounts. Due to the limited sampling by camera sensor (just three areas of the spectrum, corresponding to RGB), recovering the original color is not possible.

The adaptive flash pictures keep their saturation and hue up to $5m$. There are small visible variations from picture to picture which are mainly due to variations in exposure level and imprecision in measuring the distance between the subject and camera flash assembly. Additionally, there is the backscattering caused by particles suspended in water which is not compensated for.

To do a quantitative analysis we measured the difference between the average color of each color patch in the pictures taken underwater and the picture of the palette taken in air. The difference between two colors was measured as the euclidean distance between the coordinates of those colors in a perceptually uniform color space $L^*a^*b^*$. In a perceptually uniform color space the distance between the coordinates of two colors is a good measure for the perceptual difference between those colors. $L^*a^*b^*$ is the currently best uniform color space. The pictures were instrumented manually with the location of the color patches, and the color was averaged over the




Figure 2-20: The normalized total output energy of the adaptive flash, as perceived by the camera when the flash is fired in water, for subject to camera distances of $1m$, $2m$, $3m$, $4m$ and $5m$.

instrumented areas. For each picture we computed the average over the differences between all 15 color patches and their corresponding patches in the reference picture (the picture taken in air). The results are presented in [Figure 2-23](#).

As expected, the pictures taken with the ambient light have a poor color accuracy at all distances (due to the 5m depth at which they were taken). The pictures taken with the white flash have a steady color degradation with the distance between the camera and the subject. After post-processing the accuracy is improved, however the accuracy still decreases with distance, due to the non-uniform and non-linear distortion water introduced by water, which cannot be compensated for through a linear transformation. Finally the colors captured with the adaptive flash maintain steady color accuracy for the calibration range (1m-5m).

Figures [2-24](#), [2-25](#) and [2-26](#) present the effect of water on individual colors. Red fades into a dark shade of blue or black, mustard shifts into a shade of green, by losing the red content, while blue is not affected.




Figure 2-21: Pool experimental setup. The color palette is imaged from $1m$, $2m$, $3m$, $4m$ and $5m$. In each case three pictures are taken: using ambient light, using the standard white flash, and using the adaptive flash.

2.5.4 Ocean experiments

Figures 2-27, 2-28 and 2-29 present examples of field tests of the underwater flash. In all cases two pictures were taken, using the standard white flash and the adaptive flash. The pictures were printed unprocessed, as captured by the camera. Colors of several details of the pictures were extracted for side by side comparison. Figure 2-27 shows the advantage of adaptive illumination in identifying the various types of algae in a single scene. Figure 2-28 and Figure 2-29 were taken by a biologist and respectively a photographer. Each user was given 5 minutes of training in using the system. This demonstrates the ease of use of this system.

2.6 Summary

This chapter presented the perception component of our underwater pervasive surveillance system. We demonstrated a device capable of distance-dependent adaptive illu-

mination to capture color accurate pictures of submerged objects. We also presented an algorithm for optimizing the light composition and its experimental evaluation. The next chapter presents AQUANET the sensor network component of our system which provides the infrastructure for data collection, data storage and communication. AQUANET can use AQUALIGHT for sensing.

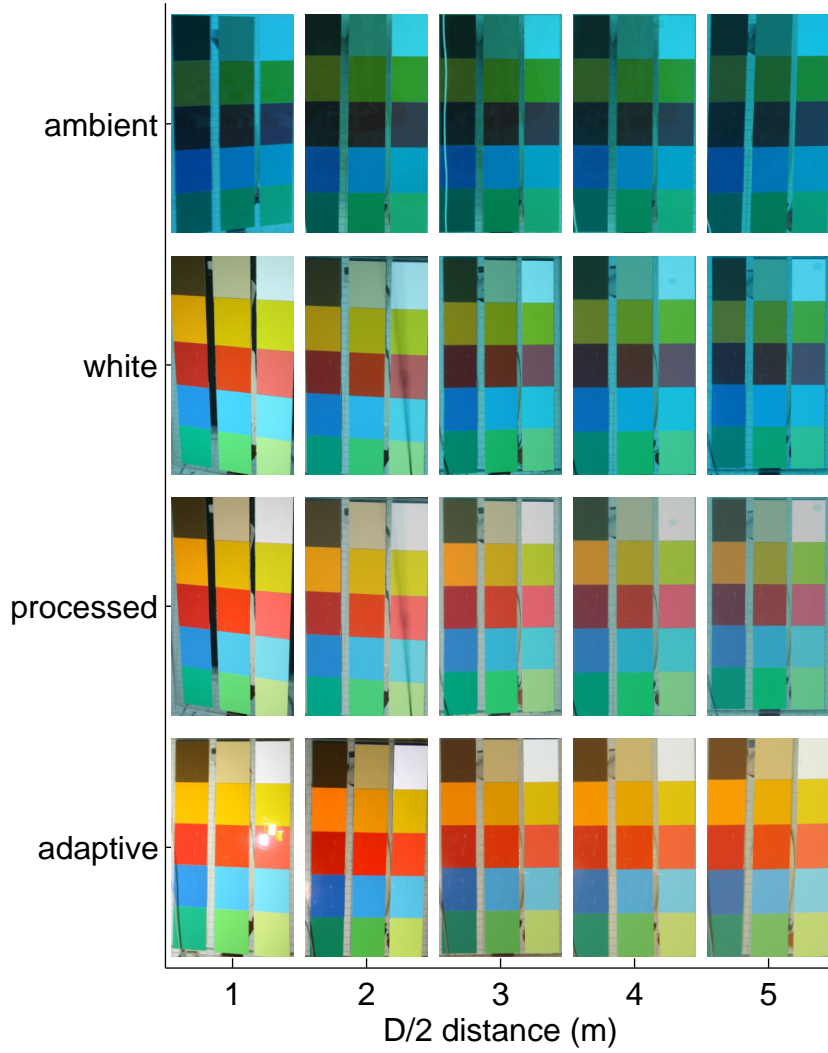


Figure 2-22: Pictures of the color palette, in MIT Alumni Pool. The pictures were taken at $1m$, $2m$, $3m$, $4m$ and $5m$ using ambient light, standard flash, standard flash and post processing and adaptive flash

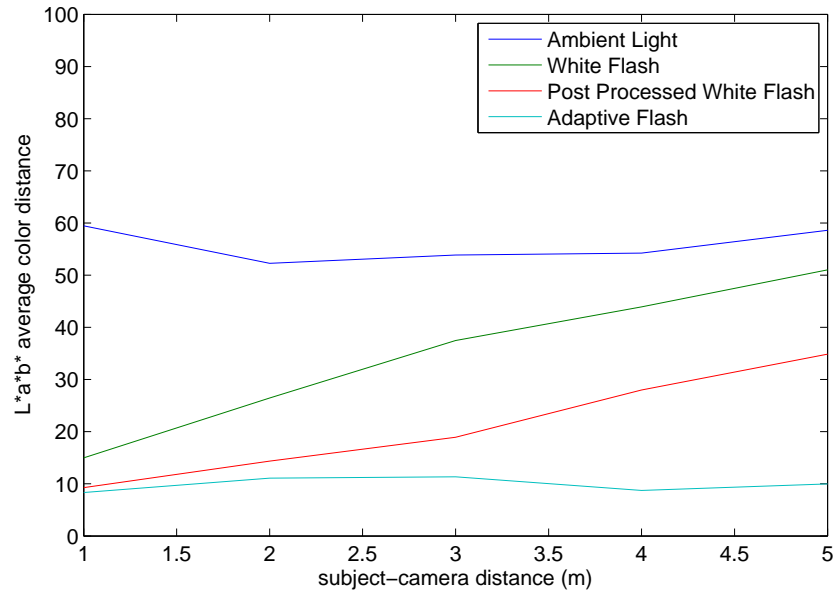


Figure 2-23: The average distance between all color samples of all pictures in Figure 2-22 and the reference picture taken in air. The distance was measured in $L^*a^*b^*$ perceptually uniform color space. Larger distances represent less color accuracy.

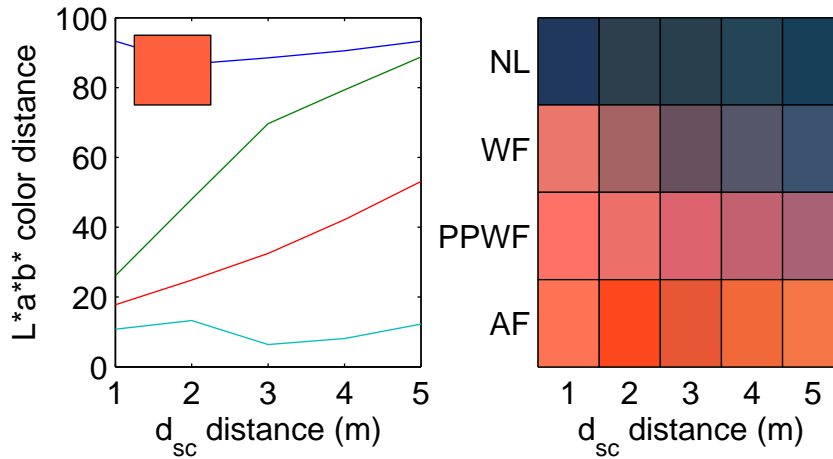


Figure 2-24: The red color sample rendered by ambient light, white flash, white flash post-processed and adaptive flash at 1m-5m. On the left graph is plotted the $L^*a^*b^*$ difference to the reference color.

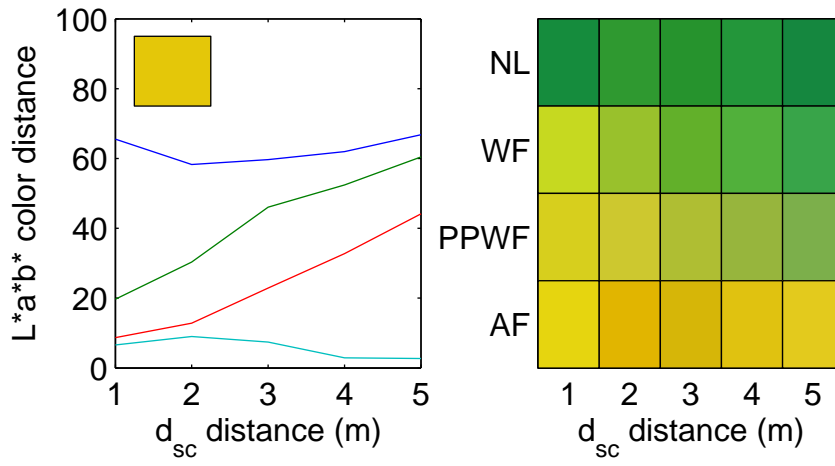


Figure 2-25: The mustard color sample rendered by ambient light, white flash, white flash post-processed and adaptive flash at 1m-5m. On the left graph is plotted the $L^*a^*b^*$ difference to the reference color.

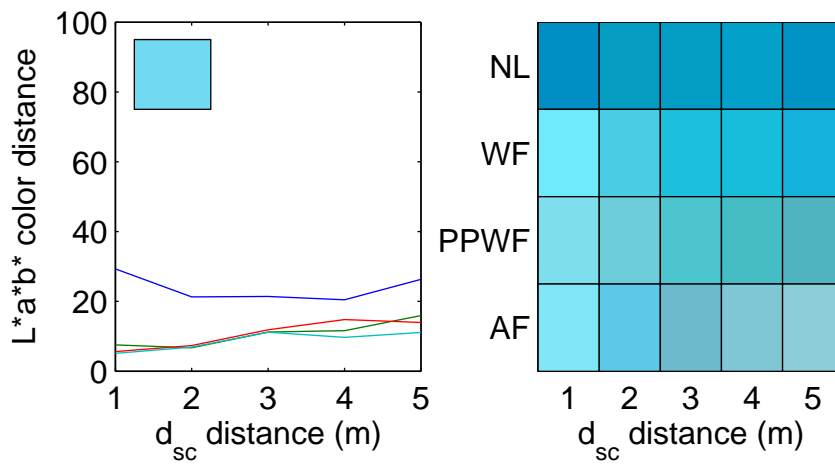


Figure 2-26: The blue color sample rendered by ambient light, white flash, white flash post-processed and adaptive flash at 1m-5m. On the left graph is plotted the $L^*a^*b^*$ difference to the reference color.




Figure 2-27: Pictures of a typical reef environment of Taveuni Island coast, Fiji, at 20m of depth, with 2.5m between camera and subject. The top picture was taken using the standard white flash. The lower picture was taken using the adaptive flash. Colors of the same regions of the pictures were extracted for comparison.




Figure 2-28: Pictures of a clown fish and its host anemone of Moorea Island coast, French Polynesia, at 18m of depth, with 3m between camera and subject. The top picture was taken with standard white flash, the lower picture was taken using the adaptive flash. Colors of the same regions of the pictures were extracted for comparison. Courtesy of Melissa Hollbruck.




Figure 2-29: Picture of a wreck section of Cayman Island coast, 15m of depth, 3m between the camera and the subject. The top picture was taken with standard white flash, the lower picture was taken using the adaptive flash. Colors of the same regions of the pictures were extracted for comparison. Courtesy of Cathy Church.

Chapter 3

Underwater Sensor Network: AQUANET

This chapter presents the design and operation of the underwater sensor network AQUANET¹. AQUANET is composed of sensor nodes (AQUANODES) that are statically deployed underwater, capable of sensing, data logging and communication with other nodes. The design goal of AQUANET is creating a testbed for experimenting with underwater data collection, communication and interaction with autonomous underwater vehicles (AUV such as AMOUR, described in [Chapter 4](#)).

We set the following requirements of the AQUANODE, the functional unit of AQUANET:

- The AQUANODES should be self-sufficient units containing sensing, computation, communication and energy supply. This enables the user to easily deploy the nodes at arbitrary locations.
- The AQUANODES should be capable of point-to-point optical communication. This enables in-situ data muling by an AUV, thus giving the user access to complete data sets without recovery and redeployment of the sensors.
- The AQUANODES should be capable of broadcast acoustic communication. This enables the nodes to create an ad-hoc network capable of relaying messages, sta-

¹The results in this chapter have been presented in [26, 31, 122, 119, 120]

tus information, measuring pair-wise ranges. This in turns enables the network to create a local coordinate system in which the network can track a mobile robot and help it navigate.

- Finally, we want the AQUANODES to be fitted with hardware that enables a mobile robot to pick them up and transport them. This enables autonomous deployment and recover of AQUANET by an AUV.

3.1 Related Work

An underwater sensor network system is a combination of wireless sensor network nodes, underwater communication technologies and algorithms addressing the particular characteristics of the sub-sea environment (e.g., long delay in acoustic communication).

Wireless sensor networks (WSN) emerged as a research field in the late 1990s through a set of reference papers by Pottie, Estrin and others [86, 85, 40, 117]. The availability of low power sensors, embedded processors, and radio technology sparked the interest in large scale distributed sensing for a variety of applications such as environmental monitoring, smart buildings and spaces, security, agriculture [40]. The hope was to automate monitoring applications where data and actions are distributed geographically by using a wireless network of distributed *sensors*, *computation*, and *communication*.

Distributed sensing promise the ability to get higher signal to noise ratio at a lower cost, for applications where the measured quantity is inherently distributed (e.g., soil moisture in a field) [40]. Instead of developing a highly sensitive sensor capable of remote sensing (and for example deploying it in an artificial satellite), the WSN approach is to deploy many local humidity sensors across the area of interest. In this example, the WSN approach provides consistent data regardless of clouds, foliage and other obstructions that may preclude the a satellite or plane carried sensor. In general the trade-off between WSN and remote or mobile sensing involves cost, data quality and data latency. The WSNs have the ability to collect and report a constant stream

of data. In contrast, the remote sensor approach might imply significant delays (e.g., the satellite's orbit determines the schedule of passes over the area of interest).

Wireless communication addresses some aspects of the cost and convenience of the system. Distributed wired sensors are acceptable in many cases. However, if the area to be covered is significant, the terrain is too hostile to wire deployment, or a high degree of mobility in the sensor positions is required, wireless communication is an effective alternative. In typical scenarios the sensors double as communication nodes in an ad-hoc network capable of routing data and messages and informing the user of interesting events. However, the removal of the wires turns the WSNs into fully distributed systems and brings all the challenges associated with distributed systems (time synchronization, reaching of an agreement, routing, etc.) as well as the benefits (fault tolerance and flexibility). Consequently much of the work over the last 20 years has been dedicated to solving these network related problems in the context of low power and minimal computation requirements for sensor networks. In some cases, communication is an area service that a WSN can provide to the user in addition to the sensing.

Distributed computation is needed to support the sensor reading operations, data storage and communication protocols. In addition there are some situations in which data can be processed directly by the network nodes, in order to take immediate decisions about the sensing process, to notify users of interesting events, to do compression, to provide services such as localization and tracking, etc. These additional uses of in-network computation have to be carefully weighted (in terms of power consumption, cost, time and fault tolerance) with the straight forward solution of routing the data to a central location (or a set of CPU locations) where it can be processed more effectively.

Wireless underwater sensor networks (WUSN) are a very new field. We are not aware of any working systems deployed and demonstrated up to now. What sets WUSN apart from air WSN (and hindered their development), is the vastly different communication environment of the undersea. Electromagnetic waves are attenuated very strongly in salt water making their use for communication impossible sub-sea.

This leaves the acoustic waves as the main communication carrier. A good presentation of the challenges posed by acoustic WUSN can be found in [3, 80].

Acoustic communications underwater have been developed over the last 30 years. The field is vast and sparked a lot of interest due to the very challenging problems posed by water as a communication channel. Very good overviews of the field of underwater telemetry can be found in [8, 58]. The water's acoustic properties and their implications for communication are discussed extensively in [67, 87]. The particular properties of the shallow water channel are discussed in [89]. Finally several medium access control and higher level protocols have been proposed in [82, 95, 111, 109]. They range from TDMA to CDMA but are limited to theoretical analysis. We reported our recent work in underwater acoustic localization and tracking in [29, 31, 122].

Optical communications underwater is a very young field. Optical communication in open air have been reported in [79]. An early underwater analog communication system was reported in [116]. It uses infrared light to transmit crayfish neuronal activity information from live crayfish in an aquarium. We reported the first use of optical networking underwater in [121, 119, 35, 120]. There are a few recent studies exploring possible techniques and systems for underwater optical communication [22, 46, 99, 109]. Recently the use of waveguide modulated optically lasers has been proposed for high speed optical communication [50]. The device is however directional, very bulky and expensive due to the difficulty in directly modulating green laser at high speed.

3.2 Why Optical?

Optical communication in water has been given very little consideration. Over the last 30 years only acoustic communication and very long wave radio communication have been considered. However, optical communication promises orders of magnitude higher data rates compared to their acoustic and radio counterparts at significantly lower power, size, and computational complexity. The drawback of the optical modems is their range limited by water turbulence, water light absorption, and power

loss due to propagation spherical spreading. In our approach we use the optical communication for the bulk sensor data downloading and we rely the mobility of our underwater vehicle AMOUR to mitigate the limited range of optical communication. There are two advantages to this approach compared to traditional acoustic data muling: (1) this solution provides a data bandwidth that scales well with the sensor network size and density and (2) it dramatically reduces the communication energy requirements on the sensor nodes (non-renewable), trading it for the vehicle's energy requirements (renewable). Current state of the art acoustic techniques provide data rates on the order of tens to hundreds bits per second over kilometers (shallow waters communication) while using tens of watts of power, and a single channel being available over the covered area. These data rates can be quickly overflowed as the amount of sensory data increases. The quantity of data increases with the sensor network density, the number of different sensors used and the amount of data generated by each sensor. For example even small images are beyond what can be practically transmitted over the acoustic channel. The power requirements for acoustic data downloading is incompatible with the long term goals of our sensor network, by requiring large energy storage for each node. The following example identifies the scaling factors governing the two approaches to data muling, acoustic and optical combined with mobility.

Consider an area of sea floor where sensor nodes are placed on regular 50m square grid over an area 2km x 2km, and thus 40 x 40 nodes. Each sensor node collects 16 bytes (salinity, temperature, pressure, nitrogen concentration and phosphate concentration with 24bit precision each) every 5 minutes. This results in a yield of 72 bytes/hour/node or 300 Kbytes/hour for the entire network, or 49.2 MBytes/week for the entire network. This data can be downloaded using the two methods: (1) using acoustic communication and node-to-node or node-to-sink communication or (2) using short range optical communication with an autonomous robot that visits the nodes in turn. A state of the art acoustic modem (WHOI Micromodem) has a data rate of 80 bits/sec at 5000m and 10W transmission power transmitting. Downloading one week worth of data using such a modem requires 8.96Wh of energy per node and

a total of 1433.6h (59.73 days). With an efficient AUV traveling at 1 knot, visiting the nodes in a raster fashion and downloading the data with our short range 300kbits/sec optical modem would require only 22.4 min of transmission and 0.002Wh of energy per node. The AUV total travel time would be of 43h (1.8 days).

In this example the acoustic network cannot relay data at the same rate as it is collected. As the data quantity increases, the only way to keep up with it is to deploy geographically multiple data sinks. These can be acoustic modems wired to the shore or to surface radios. Both solutions are hard to deploy and prone to failure and theft. In the case of the optical method, the download time depends very little on the amount of data, it is dominated by the AUV travel time. As this time increases (and thus the travel time) multiple AUVs can be easily deployed to speed up the process.

Our solution is a dually-networked system that uses acoustic communication for broadcast of small-sized data (e.g., an alert message) and optical communication for point-to-point communication with higher data rates. The next sections explain the hardware and algorithms that we developed for this kind of WUSN system.

3.3 Hardware Description

The functional unit of AQUANET is the AQUANODE a functional self-contained unit comprising sensing, computation, multi-modal communication, and energy storage (Figure 3-1). The sensor node is built around an 32bit ARM7 LPC2148 CPU running at 60Mhz.

The AQUANODE is equipped with four standard sensors: temperature, internal and external pressure, and GPS. Additional sensors can be attached to the internal A/D or on the serial communication lines. All the sensors are read with 24bit precision. The temperature and external pressure sensor are used for monitoring the water parameters. The internal pressure sensor is used for detecting leaks in the sensor node housing, before in-water deployment. The sensor node's internal pressure is

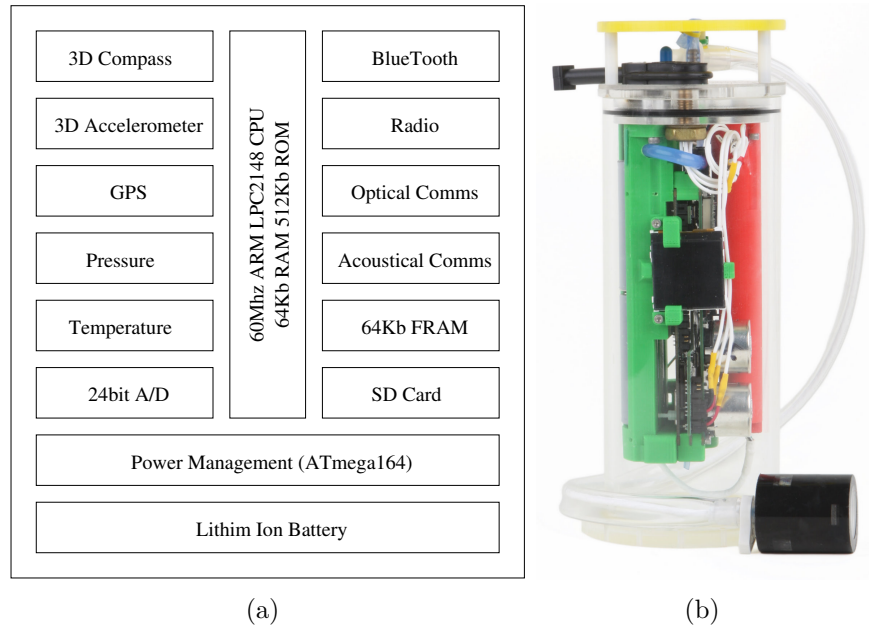


Figure 3-1: AQUANODE version 3. (current). (a) Block schematic. The board is built around a fast ARM7 processor with several sensors (left side), multi-modal communication (right side) and SD Card for data storage. The board’s power is managed by a low power AVR processor, which can turn off the rest of the sensor node and uses a charge counter to keep track of the battery state. (b) Picture of an actual node (approx 25cm of height). The lower right black component is the acoustic transducer

maintained around 30% less than the atmospheric pressure, in order to hold the top cap. The GPS is used when the node is at surface, before deployment for an estimate position. In the future, we envision having the sensor nodes actuated by an electric winch which will permit them to reach the surface for a GPS fix.

The sensor node is equipped with an SD card slot for data logging, a realtime clock, and a non-volatile FRAM module for storing calibration data. The power is managed on a separate board controlled by a low power AVR ATmega164 CPU. The AVR is connected to a charge counter which can precisely track the available power in the battery. The AVR has the ability to turn off all the sensor node subsystems, including the main board, reducing the overall power consumption of the sensor node to below 1mW.

The sensor node is powered by a 44Wh Lithium Ion battery that can be recharged





Figure 3-2: AQUANODE version 1 (AquaFleck). Designed for feasibility study, they contain an 8-bit processor, temperature sensor, camera and optical modems, all housed in a off-the-shelf watertight case. The top of the lower right node has the docking rod installed. AMOUR, the robot in the picture was able to pick up and relocate the node.

through AQUANODE's top connector. The battery supports the sensor node's operation for 6 months in sleeping mode, for 2 months with sensor readings every minute, or for 2 weeks with continuous acoustic communication.

For communication the AQUANODE contains a radio modem, Bluetooth, acoustic modem and optical modem. In software, the physical communication layers are abstracted, and therefore all types of communication are possible on each of the physical layers. In practice however, due to the vastly different performance and required conditions the 4 modes of communication are used for specialized purposes.

The radio modem (900Mhz Aerocomm AC4790-1000 1W) is used when a sensor node is carried by the robot on water surface ([Section 4.2.6](#)). It is used for uploading the missions, setting parameters or transmitting commands (e.g start mission). The main advantage is the relative long range (aprox. 200 meters max observed range) and good data rate (76kbs), which gives reliable communication while the robot is floating around the boat. The Bluetooth module is used for uploading new software or



(a)

(b)

Figure 3-3: The sensor nodes (version 2) shown deployed off the coast of Roatan, Honduras. The nodes are buoyant and they are attached to a mooring weight by a line. These nodes are equipped with the acoustic modems ([Figure 3-13](#)) and were used for the first self-localization and tracking experiments. Their dimensions are approx. 30cm high and 15cm diameter

software updates in the sensor node and for downloading data logs. It is convenient as it requires no wire, underwater connectors or opening the node, and supports a high data rate (1Mbps). Its disadvantage is the relative limited range (10meters). The acoustic modem is used for in-water communication and ranging and is described in [Section 3.5](#). The optical modem is used for in-water and in-air data downloading and sensor node reprogramming.

The AQUANODE's schematics are included in [Appendix E](#).

3.4 The Optical Networking

This section presents the design of optical networking hardware and protocols. Optical networking has been used extensively in the past in form of optical fiber communication and in open air optic communication (e.g remote controls, infrared data links – IrDa). The existing techniques concentrate on achieving high data rates. Typically air and optical fiber communication use red and infrared light which cannot be applied directly to WUSNs, due to optical properties of water and power constraints of

the sensor nodes. Water attenuates the red and the infrared regions of the spectrum significantly.

The design goals for our AQUANET's optical modem are:

- Low Power. Energy available to the sensor nodes is limited, implying modulation techniques that minimize energy per bit.
- High data rate. The data will be downloaded from the sensor nodes using an AUV which must maintain its position during the data transfer. This is a difficult task (e.g., in the presence of currents). High data rates reduce the amount of time necessary and simplify the AUV's task

3.4.1 Why Green Light?

The choice of light color reflects the need to maximize the total system gain — transmitter to receiver. The gain depends on the transmitter efficiency, optical properties of water, and receiver gain.




Figure 3-4: Spectral response of PDB-C156 PIN photodiode. The figure plots the light-to-current conversion efficiency as a function of wavelength.

For optical transmitter the only effective choice is an LED. LEDs come in several discreet colors (center wavelength) depending on the materials used in their construction. LEDs are available with a center emission wavelengths of 455nm, 470nm, 505nm, 530nm, 590nm, 617nm, 627nm. They are produced with very similar optical efficiency and maximum power. Light is attenuated by water depending on the wavelength [Section 2.3.1](#). Longer wavelengths (warm colors) are attenuated significantly stronger than shorter wavelengths (cold colors). On the receiver side there are several options for converting optical information into electric signals: phototransistors, photodiodes (PIN, avalanche), photoresistors. Among them PIN photodiodes combine a high sensitivity and a fast response (advantageous for low-power communication) at a low cost. Photodiodes have a wavelength dependent conversion efficiency plotted in [Figure 3-4](#). They are most sensitive at long wavelengths.

Choosing a wavelength for data transmissions is a trade-off between available LEDs, water optical properties (most transmissive at short wavelengths) and the spectral response of the PIN photodiode (most sensitive at long wavelengths). The system gain can be expressed as:

$$gain(d, \lambda) = LED(\lambda)PIN(\lambda)e^{-a_w(\lambda)d}$$

where

- d distance between transmitter and receiver
- $LED(\lambda)$ LED conversion efficiency
- $PIN(\lambda)$ PIN photodiode conversion efficiency
- $a_w(\lambda)$ water wavelength-dependent attenuation coefficient ([Section 2.3.1](#))

[Figure 3-5](#) plots the system gain for the available LED colors, for distances between 0m and 5m. The data is based on the LUXEON III family of high power LEDs and Advanced Photonix PDB-C156 PIN photodiode. The spherical spread was ignored and not plotted as it is not wavelength dependent and does not affect the results. At short distances the red LEDs provide higher gain while at longer distances blue-green

LEDs are more efficient. The 530nm green LED was chosen as the best compromise, with highest gain between 3m and 5m. The decision was also affected by the easy availability of high efficiency bandpass filters for green light. The filters are placed on the photodiode to improve the signal to noise ratio by restricting the received bandwidth.




Figure 3-5: Total optical gain including LED efficiency, water attenuation and photodiode efficiency

3.4.2 Hardware Description

The block schematic of the optical modem is presented in [Figure 3-6](#) and its hardware instantiation in [Figure 3-7](#). The modem is built around an 8bit microcontroller ATmega8 running at 16Mhz. The microcontroller implements the physical layer logic: packet formatting, pulse forming and decoding. For transmission the pulses generated by the processor are amplified by a current amplifier which drives the gate of a MOSFET transistor. The transistor turns the LED on and off. The role of the amplifier is to reduce the turn on and turn off times of the MOSFET. This reduces the switching losses and improves the light pulse edges which translate into better




Figure 3-6: Block schematic of the optical modem. For transmission the CPU generates and formats the packet into a train of pulses which are amplified and converted into light by the green LED and then concentrated into a 90° cone by the focuser. For reception, the light is passed through a green filter, converted into electrical signal by a photodiode, amplified and high-pass filtered, and then converted into digital pulses through a threshold circuit.

signal to noise ratio at the receiver side. The light from the LED is passed through an optical focuser which concentrates it in a 90° even cone.

On the receiver side, the light is passed through a green band-pass filter to a PIN photodiode. The role of the filter is to eliminate the ambient light outside the wavelength of interest, thus improving the signal to noise ratio and reducing the input dynamic range requirements for the receiver. When unfiltered, the ambient light can easily be 5 orders of magnitude stronger than the LED signal. The photodiode signal is amplified, high-pass filtered (to remove the ambient light) and thresholded for conversion into digital form. The digital signal is passed to the microcontroller which measures the width of the pulses and decodes the data.

A long range instance of the modem was created by placing a convex lens in front of the LED and photodiode (Figure 3-8). The long range modem trades off the receiving and transmission cone angle for range. It was used to study the influence of water turbidity on the data link range and quality.




Figure 3-7: The optical modem. (a) The the first version of receiving and transmitting board, the high power LED, its light focuser and the photodiode and it's green filter. (b) Our current version of the modem, integrating the encoding and decoding CPU.

3.4.3 Physical Layer

Our modulation encodes groups of 2 bits as the duration between two short light pulses (Figure 3-9). The duration of each pulse is determined by the ability to turn on and off the LED fast enough. In our case we can pulse the LED at 250ns. The distance between each pair of pulses is determined by the ability of the CPU to measure and distinguish inter-pulse durations. Our CPU needs a minimum of 4us between each pulses and it can distinguish between durations if they are different by at least 0.5us. The resulted timings are:

Type	Time Interval
Data bits 00	4.0us
Data bits 01	4.5us
Date bits 10	5.0us
Date bits 11	5.5us
End of byte	6.0us




Figure 3-8: The long range optical modem. The large convex lens on the right side concentrates the LED light. This modem extends the communication range by narrowing the communication cone angle.




Figure 3-9: The optical modulation. Each group of 2 bits is encoded as the duration between 2 consecutive pulses.

They result in an average of $(4 + 4.5 + 5 + 5.5 + 6)/8 = 3.125\mu\text{s}$ per bit or 320Kbits/second. The data is transmitted in numbered packets of up to 255 bytes, protected by the CRC. Error correcting codes are not necessary because the optical link transitions sharply between a perfect transmission and no transmission at all (as distance increases).

No medium access protocol was necessary as in our application only two nodes are within communication range at all times (i.e. the AUV and the sensor node). Additionally, the application protocol is master-slave, the AUV initiates all communications, while the node passively listens and replies to data requests and commands.

3.5 Acoustic networking

The requirements for our modem architecture were: (1) capability of operation in shallow waters (severe multi path propagation), (2) low power (low energy per transmitted bit and low receiver power), (3) capability of operation in a multi-node network.

Underwater acoustic communication is an established field. There are a few commercially available underwater modems ([Section 3.1](#)). However, we chose to design our own modem hardware for the following reasons:

- All commercially available modems are closed source and do not permit access to source code nor changes to modulation type, packet format, etc. They are designed for peer-to-peer connection therefore the implementation a custom the medium access control (MAC) protocol is difficult
- The only research open source modem is the WHOI Micromodem. This modem was not (and is still not) available in the quantity we needed, and parts of it remained proprietary, making the implementation and power management of the MAC protocol difficult. The Micromodem uses an inefficient linear amplifier which was incompatible with our requirements of low power.
- All available modems are bulky, beyond the size of our sensor nodes, and expensive (> \$5000).

3.5.1 Hardware Description

The block schematic of our acoustic modem is presented in [Figure 3-10](#). The modem is build around an Analog Devices Blackfin BF533 16bit DSP/CPU. The Black Fin processor does all the processing required for packet transmission, reception, medium access control, and the higher level algorithms.

For transmission the processor generates a square-wave PWM (pulse-with-modulation) signal which is amplified by a class D amplifier (switching) and is sent though a transformer to the transducer. The modem uses an FSK modulation, which permits the




Figure 3-10: Block schematic of the acoustic modem ver. 3 (current). For the transmission side the DSP assembles the packet, the FPGA converts it into PWM (pulse width modulation) signals than amplified and finally converted in acoustic waves by the transducer. On the receiving side the same transducer converts the acoustic wave into electric signals, that are then amplified by a low noise amplifier. The bandwidth of interest is selected by a filter. Subsequently the signals are amplified and fed to the A/D converter. The FPGA down-converts the signals to baseband after which they are demodulated by the DSP.

use of a non-linear amplifier. Also, the high Q of the transducer (resonant) permits it to be driven directly with a square wave signal. The transformer is used for optimal power transfer (the 1Ω output impedance of the amplifier has to be matched to the 60Ω input impedance of the transducer). The output power of the modem can be adjusted by adjusting the duty cycle of the PWM signal. The class D amplifier configuration used, which takes square wave signals as input and amplifies them, has the advantage of very high efficiency due to its switch mode operation. The only losses are the transistors' switching losses (which are negligible). The entire power is transferred into the transducer.




Figure 3-11: Acoustic modem version 3 (current) board. The size of the board is 125mm x 50mm

For reception, the signal is passed through a low-noise charge amplifier (LNA). The role of this first stage is to convert the charge generated in the transducer by the water vibrations into a proportional voltage. The LNA is characterized by an input noise density of $3nV/\sqrt{Hz}$ and an amplification A by 10. The input noise of the amplifier is particularly important as it will be added directly to the input signal. The amplification A is also important as it will reduce A times the noise influence of the following stages. Next there is a 2-pole band-pass filter with a pass-band of 6Khz centered in 30Khz. Without this filter, strong out of band signals limit the dynamic range of the signal of interest by limiting the maximum amplification possible. Following the filter there is a variable gain amplifier (VGA). The VGA gain can be set by the processor. The goal of the VGA is to amplify the signal to cover the entire input range of the analog-to-digital converter. This results in the optimal sampling dynamic range. The signal is then low-pass filtered (for anti-aliasing) and sampled and converted in digital by an ADC.

The modem board is presented in [Figure 3-11](#).

The transducer was also designed and fabricated in house (visible in [Figure 3-1](#)). The main reason for in-house fabrication was the high cost of commercially available models ($> \$1000$). The transducer is constructed by molding a piezo-ceramic cylinder in a water tight rubber case. Applying an AC voltage between the inner and outer face of the ceramic cylinder makes the cylinder deform. The cylinder is acoustically connected to the water through the rubber case. The cylinder is operated close to its mechanical resonance for maximum efficiency. The mechanical impedance of the transducer can be adjusted by adjusting the thickness of the ceramic cylinder. The mechanical impedance was chosen to match water's impedance, for maximum power transfer.

The development of the acoustic modem went through several versions ([Figure 3-12](#) and [Figure 3-13](#)). The first version used a phased-locked-loop (PLL) circuit for the FSK demodulation. The circuit did not perform well in the presence of multi-path. A standard sealed air-transducer was used, which was very limited in power due to impedance mismatch with water. The second version was very similar to the

current modem version except for the power amplifier. The amplifier used by the second version was a linear amplifier (class AB) with a maximum theoretical power efficiency of approx. 60% vs. the current version with close to 100% efficiency.




Figure 3-12: Acoustic modem and transducer ver. 1 used for feasibility study. The modem uses FSK modulation and the receiver is based on a Phase Locked Loop for demodulation integrated circuit. This scheme did not perform well due to the severe multi path propagation underwater. A sealed air transducer was used, which had very low efficiency underwater due to impedance mismatch with water.

3.5.2 Physical Layer

The goal of our system is to operate reliably in coastal and reef waters which from an acoustic perspective is considered shallow water. Shallow water is defined as a water environment where the distance between the water surface and the bottom is smaller than the distance between the transmitter and the receiver. The acoustic shallow water channel is a challenging environment for communication. Here sound reflects off the ocean floor and off the water surface and arrives at the destination via multiple paths of various lengths. The problem is aggravated by waves and features on the ocean floor (e.g., rocks) which create additional paths. Furthermore, the continuous water and wave motion creates a highly dynamic environment hard to predict




Figure 3-13: Acoustic modem and transducer version 2. From left to right the power amplifier, receiver (the DSP on the lower side) and transducer. The power amplifier is linear (class AB) and due to the low efficiency needs large heatsinks. The receiver is using digital signal processing for demodulation. The transducer was fabricated in house by molding rubber (for water tightness) around a piezo-ceramic cylinder.

with quick varying propagation conditions. Under such conditions phase modulations (and in general coherent modulations) are inefficient and hard to implement due to the constant need for equalizer training data transmission. Frequency-shift, spread spectrum and frequency hopping modulation schemes are simpler to implement and are used by currently available modems (e.g., WHOI Micromodem).

Our modem uses a simple FSK technique with an additional parameter which allows us to trade off data rate for energy per bit required. Trading off data rates for power makes sense in the context of sensor networks as they are typically deployed for long periods of time. In addition, given that the bulk of data download is done optically, the need for acoustic communication is limited to signaling, localization and tracking which require limited bandwidth.

To reduce the transmitter energy-per-bit requirement while keeping the error rate constant our modem uses a low transmitter duty cycle by spreading the symbols over time and reducing the self-generated noise (i.e. due to multi-path). In general the bit-error rate is a function of E_b/N_0 .




Figure 3-14: Saving power through lowering transmitter duty cycle in multi-path environments.

$$BER = f(E_b/N_0) \quad (3.1)$$

where

BER Bit-error rate

E_b Received energy per bit

N_0 Ambient spectral noise

f A descending monotonic function depending on the modulation type

In a modem characterized by a transmission duty cycle d_t the energy per bit can be expressed as:

$$E_b = G * d_t * P * T \quad (3.2)$$

where

d_t transmission duty cycle (between 0 and 1)

G transmitter to receiver channel gain

P transmitter power

T symbol size (time)

In a shallow water channel, a significant part of the noise at the receiver is caused by the transmitter itself through the the unquantified additional paths .

$$N_0 = N_0^{(a)} + G_n * P * d_t \quad (3.3)$$

where

$N_0^{(a)}$ ambient noise spectral density

G_n transmitter to receiver channel “noise gain”

(how much of the transmitter signal is travels through indirect paths)

For a constant bit-error rate we will plot the symbol rate $1/T$ as a function of E_b .

Constant *BER* will be achieved by varying d_t the transmitter duty cycle.

$$N_0 = E_b/f^{-1}(BER) \quad (3.4)$$

$$d_t = \frac{E_b/f^{-1}(BER) - N_0^{(a)}}{G_n * P} \quad (3.5)$$

$$1/T = \frac{G * d_t * P}{E_b} \quad (3.6)$$

$$1/T = \frac{G}{G_n} * \frac{E_b/f^{-1}(BER) - N_0}{E_b} \quad (3.7)$$

Equation 3.7 is plotted in Figure 3-15. For low E_b values the ambient noise, $N_0^{(a)}$, dominates and small increases in transmitter energy result in higher data rates. As the transmitter duty cycle increases the multi-path dominates through interference and the data rates do not increase significantly with the duty cycle.

Based on experimental measurements we operate our modem at 330 symbols/second and 30% duty cycle, which approximately corresponds to the corner of the plot in Figure 3-15.

3.5.3 Medium Access Control

Acoustic communication uses water as a broadcast medium. Each network node can receive the transmission of all other nodes in the network. Two nodes transmitting at the same time results in packet collisions and loss of transmitted data. Packet




Figure 3-15: Maximum theoretical data rate vs the transmitted energy per bit as the transmitter duty cycle is varied. The graph corresponds to a constant error rate. Only the data rates below the 100% duty cycle limit are achievable (i.e. below transmitter maximum power)

collisions waste bandwidth and power, and, therefore should be avoided (the sensor nodes have limited power available, and even without collisions the acoustic channel supports a very limited bandwidth). We devised a medium access control (MAC) protocol that operates efficiently underwater. The protocol was designed with the following assumptions:

- Collisions cannot be sensed by the transmitting nodes. There is a single transducer for both transmission and reception. Even with separate transducers collisions are impossible to detect due the wide dynamic range difference between the local node transmission and the remote node's signal.
- The sensor nodes operate as a distributed system, and there is no global knowledge of when the nodes have data to be transmitted or what the destination is. Any node may need to communicate with any other node. This assumption results from our typical applications of signaling, self-localizations, mobile nodes routing.

- The speed of sound in water (approx. 1500m/s) is very slow (e.g., compared to radio based networks). The difference between the time of transmission and the time of arrival of a packet is important. Even with the low data rates available for the acoustic channel, these latencies are significant relative to the packet duration.
- The sensor nodes are battery-powered and battery consumption is an important characteristic of the protocol.

We considered several established MAC protocols: CSMA/CD, CSMA/CA, FDMA, CDMA and TDMA.

- CSMA/CD (carrier sense multiple access with collision detection) is not feasible as the transmitting nodes cannot detect the collision.
- CSMA/CA (carrier sense multiple access with collision avoidance) is still difficult to implement due to the high channel latency. For example a distance of 200m between two nodes results in a latency of 0.2s, which is already half the specified duration (0.4s) of our typical 16 bytes packet. Not sensing a carrier is not a reliable indication that a packet is not already traveling. In addition colliding packets carry a particularly high price tag for self powered networks where energy consumption should be minimized.
- FDMA (frequency division multiple access) requires each node to operate on orthogonal frequencies bands, and therefore their transmission would not interfere with each other. In the case of underwater acoustics there are two limitations. The propagation characteristics of the acoustic channel are frequency dependent, limiting the usable bandwidth to a relatively narrow band of less than 10Khz for our application. This would put severe constraints on the scalability of the network. In addition the piezo-transducer used as an antenna is a high-Q device, and therefore operates efficiently in a limited band. Separating properly signals from several transmitters require that the signals arrive at the receiver with relatively similar strength. Due to the distributed nature of the

network, this arrival signal strength is impossible to insure as it will depend on the position of the transmitters and the receivers. For example if 2 sensor nodes A and B are transmitting at the same time while the sensor nodes C and D are receiving, and C is closer to A than to B while D is closer to B than to A, there is no power setting for A and B that will work for both C and D at the same time. Finally, the fact that a node cannot transmit and receive at the same time leaves the possibility for packets to be lost: if a node A transmits to node B while the node C transmits to the node A, the packet from C to A will be lost.

- CDMA (code division multiple access) does not work for the same reasons as FDMA: the input dynamic range at the receiver side and the possibility of collision between the a packet for node A and a packet transmitted by node A.

This leaves the TDMA (time division multiple access) option. TDMA works by dividing the time in slots, and each node would transmit in slots agreed upon globally (i.e. the rest of the nodes expect a transmission during those slots and do no transmit). The challenge for implementing a TDMA scheme underwater are long delays caused by the slow speed of sound, which, in the case of geographically distributed networks causes significant difference between a packet's transmission time and reception time. This problem is often ignored in current literature due to the use of TDMA mainly in radio networks where the delays are small. In the case where the delays are significant (e.g., satellite communication, GSM networks) all nodes communicate with a base station. In this case the problem is simpler as the base station time is the reference time, and the rest of the nodes can adjust their transmission time.

Our scheme divides the time into N equal size time slots labeled $0..N-1$. The slots are allocated in numerical progression; $N-1$ is followed by slot 0. Each slot can be owned by at most one network node. In each slot we allow exactly two packets: one is sent by the slot owner (called QUERY packet) and the second one is sent by the node which was interrogated by the slot owner (called RESPONSE packet)). We chose this particular scheme for several reasons: (1) due to the unreliable nature of

acoustic communication very often a confirmation of correct reception is needed, and having an immediate slot for the ack simplifies the design, (2) by having an immediate response the slot owner can measure its distance to the queried node by measuring the two-way time-of-flight (which is important for our application (self-localization, tracking) and slot time synchronization). Slot allocation to nodes is left to the higher protocol layers. In the simplest scenario there are as many slots as nodes in the network and each node owns exactly one slot.

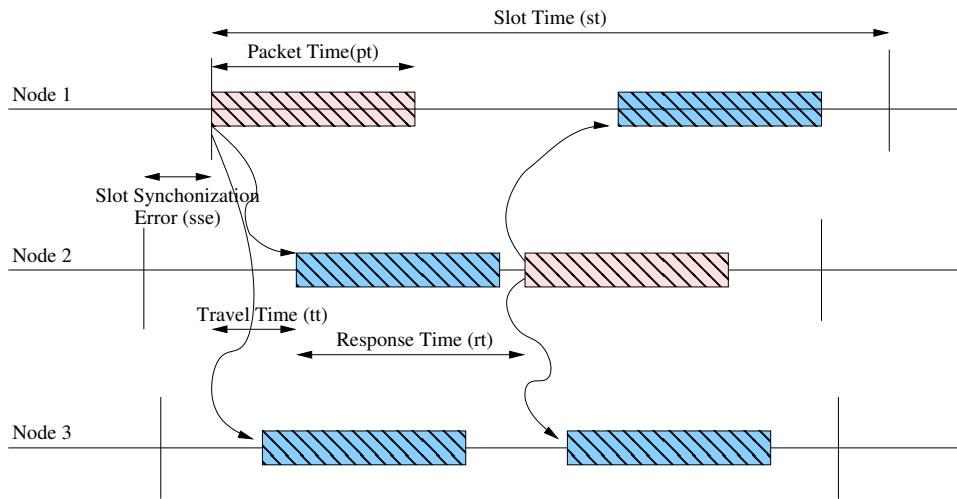


Figure 3-16: Time lines of Node 1, 2 and 3 while communicating with non-negligible desynchronization and packet travel time. Node 1 interrogates Node 2 and receives a response while node 3 is passively listening

Two inter-dependent problems are: (1) deciding the slot size, and (2) keeping the network in sync, as each node has its own time base. The slot size must be as small as possible (for low latency) but large enough to accommodate the packets' travel time and the time desynchronization between the nodes. Figure 3-16 presents a typical communication scenario: node 1 interrogates node 2 and receives a response from node 2 while node 3 passively listens. For successful communication (both the QUERY and RESPONSE packets remain in their originating) the following inequalities must hold true:

$$ST > TT + RT + TT + PT \quad \text{constrained by 1st node's time line} \quad (3.8)$$

$$ST > SSE + TT + RT + PT \quad \text{constrained by 2nd node's time line} \quad (3.9)$$

$$SD > SSE + TT + RT + TT + PT \quad \text{constrained by 3rd node's time line} \quad (3.10)$$

where

ST Slot duration

TT Maximum packet travel time

SSE Maximum slot desynchronization between the two nodes

PT Packet duration

RT Response time

Every time a node receives a packet, it can resynchronize its time base. That limits the slot desynchronization to the maximum travel time ($SSE < TT$). In addition the response time (RT) is the packet duration (PT) plus the computation time needed to create a response packet. We can approximate $RT = PT$. The three constraints can be reduced to a single one which encompasses them:

$$ST > 3TT + 2PT \quad (3.11)$$

The slot time should be greater than 2 times the packet size plus 3 times the maximum travel time. The maximum travel time is given by the network diameter (i.e the maximum distance between any two nodes) or the maximum communication range of the modem, whichever is smaller. If there are no slot synchronization errors between nodes $SSE = 0$, the constraint becomes:

$$ST > 2TT + 2PT \quad (3.12)$$

which is the natural minimum constraint for a network with 2 packets per slot.

Algorithm 3 summarizes our TDMA protocol. It is run on every node of the network. The slot-synchronization between nodes is insured by updating the local slot position with every packet received correctly (line 13). The procedures UPDATE-OFFSET and CURRENTSLOT are presented in Algorithm 4. The C implementation is included in Appendix F.

Algorithm 3 Self-Synchronizing TDMA

Require: *myID* the local node ID
mySlots the set of currently owned slots
CURRENTSLOT returns the current slot
 (based on the time and current slot offset)
UPDATEOFFET updates the slot 0 time offset

```

1: loop
2:   wait till beginning of a slot
3:   if currentSlot()  $\in$  mySlots then
4:     transmit packet(myID, destination, currentSlot(), QUERY)
5:     listen till a package is received, or end of slot
6:     if packet(dest, myID, data, RESPONSE) received at T then
7:       UPDATEDISTANCE(dest, T)
8:     end if
9:   else
10:    repeat
11:      if packet(sender, destination, slotID, type) received at T then
12:        if type = QUERY then
13:          UPDATEOFFSET(sender, slotID, T)
14:          if destination = myID then
15:            wait till time T + responseTime
16:            send packet(myID, sender, currentSlot(), RESPONSE)
17:          end if
18:        end if
19:      end if
20:    until end of slot & next slot  $\in$  mySlots
21:  end if
22: end loop

```

Algorithm 3 continuously updates the slot positions and insures that they remain aligned over time. Initializing the slots is not necessary provided that at least one node transmits on a unique time slot (not colliding with any other node). This will insure that its QUERY packet will not collide with other packets and will be received

Algorithm 4 Slot Synchronization

Require: ST slot duration
 $TIME$ the current time
 N number of slots

- 1: **global** $slotOffset$ ▷ Start time for slot zero
- 2: **procedure** UPDATEOFFSET($sender, slotID, T$)
- 3: $slotOffset \leftarrow (T - slotID * ST)$
- 4: **end procedure**
- 5: **function** CURRENTSLOT
- 6: **return** $\lfloor (TIME - slotOffset) / ST \rfloor \bmod N$
- 7: **end function**

by other nodes. The receiving nodes will then update their slot time offsets and the TDMA will synchronize. Given that there are N slots and each node owns one slot initially, the probability of at least one node transmitting in a time when no other node is transmitting initially is very high. However, to eliminate the small possibility of synchronization failure, we run Algorithm 5 before the regular TDMA. It essentially waits until the first packet is received correctly (thus insuring the necessary initial condition for self-synchronization), while transmitting at the regular time slots. The nodes transmits with a small probability P on other slots as well. This ensures that even if initially all the nodes have packet collisions, eventually their transmissions will be separated.

3.6 Experiments

We conducted several pool and ocean experiments to demonstrate the basic capabilities of our sensor nodes: sensing, data logging, optical communication and acoustic communication.

3.6.1 Optical Communication and Data Muling

We tested the optical communication range and quality by positioning two nodes in water, facing each other, at increasing distances. For each distance one of the nodes transmits a train of 200 numbered packets of 255 bytes each including a checksum.

Algorithm 5 Slots' positions initialization

Require: $myID$ the local node ID
 $mySlots$ the set of currently owned slots
 N number of slots
 P probability of using a slot not owned
UPDATEOFFSET updates the slot 0 time offset

```
1: if  $mySlots = \emptyset$  then  
2:   exit ▷ This node is listener, no need for init  
3: end if  
4: loop  
5:   for  $i = 0$  to  $N$  do  
6:     wait for the beginning of slot  $i$   
7:     if  $i \in mySlots$  or with probability  $P$  then  
8:       transmit  $packet(myID, all, i, QUERY)$   
9:     else  
10:      listen till the end of the slot  
11:      if  $packet(sender, destination, slotID, type)$  received at  $T$  then  
12:        UPDATEOFFSET( $sender, slotID, T$ )  
13:        Exit ▷ Node synchronized  
14:      end if  
15:    end if  
16:  end for  
17: end loop
```

The second node counts the packages received correctly. We carried out the experiment for 2 sensor nodes fitted with standard modems, and then by substituting one of the modems with our long range modem (Figure 3-8). The results are presented in Table 3.2 and Table 3.1.

Distance	Received Packets	Missed Packets	Packet Success Rate
0.5m	200	0	100.0%
1.5m	200	0	100.0%
2.0m	197	3	96.0%
2.4m	195	5	97.5%
2.5m	0	200	0.0%

Table 3.1: Reception rates for optical communication in clear water. The reception rate drops sharply at the modem’s range limit.

Distance	Received Packets	Missed Packets	Packet Success Rate
2.1m	200	0	100.0%
4.3m	200	0	100.0%
5.3m	198	2	99.0%
6.4m	200	0	100.0%
7.0m	200	8	96.0%

Table 3.2: Reception rates for optical communication in clear water for the long range modem. No reception was observed beyond 7m.

The optical modem has a very good packet success rate up to the range limit where it drops sharply. This is due to the simple demodulation method. We can envision more sophisticated methods of demodulation which would increase the range at the expense of the data rate. The transmission cone was measured to be 90° for the standard modem and 20° for the long range modem. The receiver has 120° reception angle.

We evaluated the optical modems in the context of data muling. 4 sensor nodes were placed on the bottom of the pool, to collect sensory data including temperature,




Figure 3-17: Starbug AUV [36, 119] acting as data mule. (a) The sensor nodes were deployed at the bottom of the pool for data collection (temperature, depth and images). Starbug fitted with an optical modem visits the nodes them in turn, downloading the data optically. (b) A picture of Starbug taken by one of the sensor nodes and downloaded optically by Starbug.

pressure and pictures. An AUV (Starbug, [36, 119]) fitted with an optical modem was sent to download the data (Figure 3-17). The AUV visited the nodes in turn, using its cameras for navigation and node acquisition, and downloaded the data collected by the sensor nodes optically. This experiment was repeated several times with various node configurations. In total over a megabyte of data was transmitted optically.

A third type of experiment done with the second version of the sensor nodes (Figure 3-3), used the optical modems as the primary mode of communication for upgrading the software in the sensor nodes. We estimate we transmitted over 10 megabytes optically while programming the nodes.

3.6.2 Acoustic Communication and Ranging

We performed several experiments to test the performance of our acoustic modem in real world conditions. We deployed 4 AQUANODES in lake Otsego, NY. The AQUANODES were deployed at arbitrary locations about 50 meters apart as represented in

Ground truth (m)	Measure average (m)	Standard Deviation (m)	Error (m)
4.75	5.01	0.04	0.26
9.67	9.72	0.08	0.05
16.15	16.19	0.07	0.04
24.7	24.35	0.11	-0.35
N/A	37.9	0.06	N/A
N/A	49.1	0.14	N/A

Table 3.3: Acoustically measured ranges between node A and node E compared to the ground truth.

Figure 3-21. The nodes are labeled A through D.

For the first set of experiments we tested the precision of the range measurements. We set a fifth node E to continuously measure and log the distance to node A. We moved node E to increase the distance to node A. For ground truth we used a measuring tape. We kept the E node stationary for 5 different distance as reported in Table 3.3. The ranges measured during the entire experiment are plotted in Figure 3-18. The standard deviation is very small (less than 10cm). Over more than 200 measurements were performed.

In a second experiment we set node E to continuously interrogate all static nodes (A-D). We attached node E to a boat that was driven around the sensor field. We plotted the measured range from node E to the nodes A-D in Figure 3-19. We noticed a very low number of outliers ($< 0.4\%$) over 1000 measurements, which were in the end tracked down to a bug in our software. We also plotted the round trip packet success rates in Figure 3-20. The success rate results are summarized in Table 3.4. We noticed that the success rate did not depend on distance.

3.6.3 Localization and Tracking

In this experiment we tested the ability of AQUANODES to self localize and track a mobile robot. We set the 4 static nodes to measure and exchange all the pairwise




Figure 3-18: Acoustically measured range between node A and node E. Node A of positioned in a fixed location while node E was moved at increasing distances to node A.

distances. Each node had one time slot allocated. During its time slot, each node attempted to measure the distance between itself and other nodes, while using the data portion of the packets to broadcast known ranges or interrogate for unknown ranges. The matrix of collected ranges is presented in [Table 3.5](#).

Based on the pairwise distances, the nodes constructed a system of coordinates using a least-square error algorithm. Node A's position was chosen to be the origin of the coordinate system while Node B's position was chosen along the X axis. The positions of the nodes is summarized in [Table 3.6](#).

Finally, we attached node E to the underwater robot AMOUR (described in [Chapter 4](#)) so that the robot becomes a mobile node in the sensor network and set it to measure continuously its ranges to all the static nodes. During this part of the experiment we allocated all the slots to node E for faster update rate. We drove AMOUR under remote control in the sensor field, while keeping it close to surface. We logged the ranges between AMOUR and the static sensor nodes. We also logged AMOUR's

Node	Round trip success rate (measured)	One way success rate (computed)
A	58%	76%
B	64%	80%
C	64%	80%
D	50%	71%

Table 3.4: Acoustic packet success rate between node E and nodes A-D. The one way packet rates success rates were computed considering equal that the two packets in a round trip communication have an equal probability of success.

(m)	A	B	C	D
A	0	68.80	99.42	44.67
B	68.80	0	59.15	78.84
C	99.42	59.15	0	77.44
D	44.67	78.84	77.44	0

Table 3.5: The ranges between the static sensor nodes A-D as measured and exchanged by the nodes. Each node collected a copy of this matrix.

(m)	X	Y
A	0	0
B	68.8	0
C	80.8	57.9
D	3.72	44.52

Table 3.6: The position of the sensor node in the locally constructed coordinate system, based on pairwise ranges.




Figure 3-19: The acoustic tracking experiment

GPS position for ground truth. We reconstructed AMOUR’s path from the acoustic range measurements using the algorithm described in [29]. The reconstructed path and the GPS logged paths are plotted in Figure 3-21. The GPS path was rotated and translated to fit the local coordinate system of the sensor nodes. The robot traveled over 250m. The mean error between the GPS path and the acoustically computed path was 2.05m. This is comparable to GPS error. It is hard to separate the contribution of the GPS error from the error of our acoustic tracking system.

3.7 Summary

This chapter presented the static underwater sensor network component of our system. This component is responsible for long term data collection, storage and networking. We presented: (1) the optical modem capable of short range, low power, high speed communication, (2) the acoustic modem capable of broadcast communication and inter-node range measurements, (3) communication protocols optimized for low power. The next chapter presents the robotic component of our system which




Figure 3-20: The acoustic packet success rate experiment.

enables optical data muling, high resolution surveys and automatic sensor network deployment, relocation and recovery.




Figure 3-21: Tracking AMOUR

Chapter 4

Underwater Autonomous Vehicle: AMOUR

This section presents the design and operation of AMOUR¹ (Autonomous Modular Optical Underwater Robot), an AUV tailored for operations in coastal waters and reef environments. We set the following requirements for AMOUR: (1) ability to hover; (2) ability to operate in an underwater sensor network; (3) ability to pick-up, travel with and drop off payloads, (4) ability to carry standard marine instruments over medium range missions (10km); (5) a small, light-weight package that can be carried by a single person and fit in a suitcase; (6) low-cost modular configuration. After considering the option of using and adapting an existing underwater robot design we decided that this path would not easily satisfy our design requirements. We opted to design a new robot system from scratch for several reasons. First, we were interested in exploring and understanding the limits of current underwater technology. Second, we were interested in incorporating the latest technological advances in electronics, motors, sensors, batteries and fabrication techniques. Third, we wanted to approach our system design requirements in an organic integrated fashion, as opposed to an add-on approach.

The key capability of AMOUR is the ability to operate in symbiosis with a wireless

¹The results in this chapter have been presented in [30, 37, 121]



Figure 4-1: AMOUR and a few sensor nodes. AMOUR has attached the underwater camera and the adaptive illumination flash.

underwater sensor network. AMOUR and the sensor network nodes are fitted with custom developed acoustic and optical modems for low data rate broadcast and high data rate point-to-point communication respectively. AMOUR can visit the sensor nodes and download their data using its high speed, low power optical modem, enabling real time, in-situ data muling, which is not practical with acoustic telemetry and without using mobility. The sensor network also acts as a GPS-like constellation of satellites, providing the robot precise position information. This eliminates the need for the robot to include large expensive underwater position information sensors for navigation, such as Doppler velocity loggers or laser ring gyros.

The synergy between the network and the robot is more extensive. AMOUR is fitted with a docking mechanism, which enables the robot to autonomously pick up sensor nodes for deployment, relocation and recovery. In order to transport efficiently the dynamic payload of the sensor nodes, the robot was fitted with internal buoyancy and balance control. AMOUR relies on a water piston mechanism to compensate for changes in buoyancy of up to 1kg due of payload. In addition an internal battery moving mechanism enables the robot to change its center of mass position and ac-

commodate the same 1kg payload at its tail end without help from the thrusters for balance. The buoyancy and balance control systems provide an efficient way for controlling the robot’s buoyancy, center of mass and overall orientation. Both of these systems save significant energy during the robot’s mission as they require energy during adjustment only. In the absence of buoyancy and balance control mechanisms the robot would have to use the main thrusters (and hence a large amount of additional energy) to compensate for the payload. Additional applications for our underwater robot are complex autonomous and collaborative underwater pick and place operations – for example assembling underwater structures, collecting samples or using tools for repairs.

4.1 Related Work

Our work builds on a large body of previous research in underwater robotics. Underwater autonomous robots began to be fielded in early 1990s with robots such as Woods Hole’s ABE [128] and MIT’s Odyssey [11]. Since then the field has undergone continuous technological improvement. For a survey of AUVs and ROVs evolution see [129, 6].

Today’s AUV landscape includes several companies such as Hydroid and Bluefin Robotics that produce AUVs as their main product. These AUVs are torpedo shaped robots geared toward commercial and military applications (offshore oil, security, environmental monitoring). These are optimized for long surveys using various sensors. The robots are unable to hover (as they are underactuated and need constant motion for control). These AUVs include the Remus series from [112] and the Bluefin series from Bluefin [4]. These robots can perform missions of mine detection, mapping, marine raster surveys. The research community has also developed a number of non-hovering AUVs. SAUV [27] has solar cells which allow it to recharge at the surface for long endurance missions. The Starfish AUV [96] was designed to be modular and easily upgradeable.

The ability of an underwater vehicle to hover is commonly found in remotely oper-

ated vehicles (ROVs). Autonomous vehicles with hovering abilities include AQUA [34] which is an amphibious robot with six flippers, Finnegan the RoboTurtle [66] which is a biologically inspired four flipper robot and The ODIN-III [24] has eight thrusters and a one degree of freedom manipulator. Seabed [105] AUV is designed for high resolution underwater imaging. It has a dual torpedo shape with enough degrees of freedom to support hovering. It is designed to be a statically stable platform for imaging. WHOI's Sentry [68, 127] AUV (the successor of ABE) is designed for efficient cruising but is also able to rotate its thrusters to hover for near-bottom operations. The WHOI Jaguar and Puma AUVs [63] (successors of Seabed) were developed to study the ocean under the Arctic ice cover. A smaller scale robot, adapted for coastal water is CSIRO's Starbug [36]. It can hover and at the same time it is very efficient while moving forward, due to its flat later thrusters. The HAUV [118] is fully actuated vehicle designed to perform autonomous operations such as ship hull inspection. AMOUR introduces a modular composable thruster design that enables navigation and hovering in both vertical configuration (for high maneuverability) as well as in a more streamlined horizontal configuration (for long duration missions.)

Most underwater robot systems are neutrally buoyant and do not execute tasks that require control of their buoyancy. A few robots that do change their buoyancy include Divebot [69] which changes its density by heating oil, SubjuGator [64, 77], which uses two solenoids which regulate the amount of ballast in a buoyancy compensatory, and gliders such as the Spray gliders by Bluefin Robotics [102] and University of Washington's Sea Gliders [39] which change their buoyancy by battery-powered hydraulic pumps in order to glide forward. Our work provides simultaneous control of balance and buoyancy to enable the pickup of a payload. Gliders are the only class of robots that also control both balance and buoyancy. However, gliders do so for different reasons. They use buoyancy and balance control as their main propulsion mechanism.

Navigation and position estimation for AUVs is typically achieved by using a combination of a Doppler velocity logger (for dead reckoning) and laser ring gyros (for precise orientation). These sensors are far too large and expensive for the scale

of AMOUR. Our robot instead relies on an external localization system provided by a statically deployed sensor network. The network self localizes and provides localization information to our robot [29, 31, 122]. This differs from the traditional long base line (LBL) systems (which is how many AUVs localize [118]) in that our system does not require a priori localization of the beacons (sensor nodes).

4.2 Hardware Description




Figure 4-2: AMOUR, three generations. (a) AMOUR I, used for studying the feasibility of optical guided docking. (b) AMOUR II, capable of optical docking and data muling, acoustic communication, tracking and sensor nodes transport. (c) AMOUR V has all the features of AMOUR III plus the ability to adjust its buoyancy and balance, and higher payload capability (in the image its carrying the camera and the adaptive illumination flash).

AMOUR's body is built from a 63cm long, 16.5cm diameter acrylic tube with a wall thickness of 0.6cm. The cylindrical body shape was chosen for hydrodynamic reasons. The body is an efficient pressure vessel and can be easily built from low cost materials. In order to dock and carry a sensor node efficiently, the robot was designed to be operational with the body both in horizontal and vertical configurations. Horizontal configuration is used for long distance, streamlined travel. The vertical configuration

is used for docking and data muling. Vertical docking allows the sensor node to be attached at the end of the body, thus maintaining the streamlined shape during travel. The robot's configuration is modular. The system was designed so that it is easy to add or remove modules such as sensor network nodes, sensing modules, as well as the buoyancy and balance control modules.

Attached to the top cap of the tube is the *head* of the robot. It contains the electronics for planning and control (Section 4.2.5) and the Inertial Measurement Unit IMU (Section 4.2.5). The head is contained in a splash proof container for easy removal in-situ or immediately after the mission. The mid section of the body contains the battery module supported by 4 stainless steel rails. The balance control mechanism sets the battery module's position along the length of the robot (Section 4.2.3). The battery module is composed of 72 rechargeable Lithium Ion cells with a nominal voltage of 3.7V and a charge of 2Ah each. We used readily available 18650 laptop cells. The battery is organized as a 4-series cell with a nominal voltage of 14.8V and a charge of 36Ah, thus a total energy capacity of over 500Wh.

Below the battery module there is the power PCB, which contains the power distribution and voltage conversion circuits. The space between the power PCB and the lower cap can be used for additional payloads (e.g., a video camera). The lower cap of the robot is replaceable with the buoyancy and docking module (Section 4.2.3) or other modules. The two end caps of the main tube are held in place by the lowering the inside pressure of AMOUR's body below atmospheric pressure. We typically operate the robot with inside absolute pressure of 0.9Bar, which creates a holding force on the end caps of about 180N at sea level, and increasingly higher as the robot dives. The end caps are sealed using standard o-rings. An internal pressure sensor can easily identify potential leaks, before deployment of the robot in water.

AMOUR has five thrusters (Section 4.2.1) attached to the main body using easily removable stainless steel clamps. The thrusters are electrically connected to the main body through flexible PVC tubes, an effective and significantly less expensive solution than underwater connectors. Three thrusters are collinear with and evenly distributed around the main body tube. The remaining two thrusters are orthogonal to the main

body. This thruster configuration allows the robot to be oriented and controlled in both horizontal and vertical configurations. Additional thrusters can be added to the robot. Previous revisions of the robot included four thrusters only.

The robot is more streamlined horizontally than vertically for energy efficient motion. In the vertical configuration the robot is more maneuverable. This configuration is used when docking with sensor nodes or downloading data via the optical modem. In both configurations the robot has independent control over its pitch, roll and yaw angles as well as its depth and forward/reverse motion (details in [Section 4.3.2](#)). Side motions are not possible in either configuration. In the vertical configuration the robot is able to spin along the yaw axis very rapidly mitigating the lack of direct lateral control. This allows the robot to respond to currents and unexpected water motions. In the horizontal configuration the rotation is slower, however, this is less important as this configuration is used for long distance travel when maneuverability is less important.

In addition to the thrusters, the robot can carry additional sensors and devices attached to its body: video camera, still camera, scanning sonar, range finder, and/or one of our sensor nodes. We typically run the robot with one of our network sensor node attached to its body. The node's CPU runs the high level mission control. In addition the sensor node contains a radio modem and GPS (for surface link with the robot), acoustic and optical modem for underwater communication, an SD card for data logging, temperature sensor and 24bit ADCs for additional survey sensors (e.g., salinity). When in horizontal configuration and traveling close to surface the GPS and radio antenna can stay above water for localization and mission control.

4.2.1 Propulsion

We opted to design and fabricate AMOUR's thrusters as the commercially available units were too big, did not have the required thrust, and were generally very expensive. The thrusters are built around Himax HB3630-1000 sensorless brushless DC motors fitted with a 4.3:1 planetary gear box. The motors are rated for 600W and are intended for the RC market. They have a very low winding resistance and thus

are able to generate high power efficiently in a small package (with proper cooling). Brushless motors were chosen for their higher power and efficiency compared to the mechanically commutated DC motors. The placement of the windings on the stator (the body) is also advantageous, as efficient cooling can be easily insured by thermally connecting the motor's body to the thruster case.

The motors are driven by in-house developed power brushless controllers build on two PCBs attached to the motor's back. The controllers receive commands from the main CPU over a RS485 bus and regulate the motor speed accordingly. The complete schematics are included in [Appendix G](#).

The motors, gearbox and electronics are housed in a watertight aluminum case. The 6mm stainless steel shaft of the propeller penetrates the top cap and it is sealed by a double o-ring. Between the two o-rings there is an oil cavity that assure lubrication and prevents overheating. The propeller used is Prop Shop K3435/4 with a 90mm diameter and 100mm pitch. The propeller size was determined experimentally to maximize the static thrust and efficiency. The propellers were statically balanced in house (by removing material from their blades) to reduce vibrations and improve efficiency.

4.2.2 Docking

The underside of AMOUR's buoyancy module is a conical shaped cavity ([Figure 4-4](#)) that allows docking with any object that holds a docking element — a 15cm long rod of 1cm diameter with a modulated light source at its base as beacon. The mechanism is illustrated in [Figure 4-3](#). The conical shape of the docking cavity provides the mechanical compliance during the final stage of docking. This alignment mechanism is similar to the probe and drogue mechanism used by NASA's Apollo program. A latch plate with a variable width hole located at the apex of the cone is actuated by a servo-motor in the horizontal plane in a tub-like compartment. To dock, the plate is actuated such that the smaller side of its variable width hole is latched around a corresponding groove on the docking rod. This creates a strong link between the




Figure 4-3: The docking mechanism.

robot and the docked object (e.g sensor node) capable of withstanding forces of up to 200N.

AMOUR uses 2 pairs of optical sensors distributed around the cavity circumference for optical alignment during docking procedure (one for the X axis and one for the Y axes). Each pair determines the position of the light beacon based on the absolute value and the difference in reading of its two sensors. The absolute value is an indication for the distance to the optical beacon and the difference is an indication for the angle to the beacon on that the pair's axis. A recent docking sensor is presented in [Figure 4-5](#). This sensor is based on 2 elongated photodiodes separated by a opaque wall. The wall casts a shadow on one of the photodiodes depending on the position of the beacon light. The sensor uses a bandpass filter tuned on the beacon's modulation to eliminate the ambient light influence. This sensor has not been evaluated yet in practice.

4.2.3 Buoyancy and Balance

We extended the robot with buoyancy and balance control mechanisms, packaged as auxiliary modules (see [Figure 4-6](#)). The main design requirement for the mechanisms




Figure 4-4: The docking cone. View from of the underside of AMOUR.

was to achieve adaptation to an additional payload of up to 1 kg within 30 seconds. The robot uses its thrusters to achieve a desired depth and orientation. Extra weight would make the thrusters work harder. The role of the buoyancy control module is to bring the robot back to a neutrally buoyant state (and thus save energy). The role of the balance control module is to change the center of mass of the robot when additional weight is added or in response to the change in the buoyancy control module.

The buoyancy mechanism (see [Figure 4-6](#)) controls the buoyancy of the robot by moving a piston inside a cylinder. The cylinder has a 16.5 cm diameter. The piston can travel 5.6 cm over 24 sec and has 1.2 kg lift capacity. The effect on buoyancy is $\Delta m = A\rho\Delta h$. The buoyancy system coordinates with the balance control mechanism ([Figure 4-8](#)), which alters the center of mass of the robot by moving the battery pack up and down in the robot. The buoyancy system has an integrated docking mechanism ([Figure 4-3](#), [121, 37]) which enables the automatic pickup of payloads that are compatible with the docking mechanism (for example, our underwater sensor network nodes can be picked up by this robot). The buoyancy control module is




Figure 4-5: A new docking sensor design. (a) view of the signal processing board. (b) The sensor is based on two long photodiodes, covered by green filters, separate by a opaque wall. The wall will cast a shadow on one of the photodiodes depending on the direction of the incoming light.

contained in a water-tight cylindrical tube that can be attached to the robot with an underwater connector. It includes a piston moved by a set of three ball screws with ball nuts. A custom-designed gear box ensures that the ball screws are turned simultaneously and provides a gear ratio that can compensate for forces arising at up to 40m depth². The output power P_{out} of the motor is $P_{out} = \frac{(V - k_r \frac{T_m}{k_m}) k_s T_m 2\pi}{60}$, where V is voltage, T_m is the motor torque, k_r is the terminal resistance in Ω , k_m the torque constant in $\frac{mNm}{A}$ and k_s as the speed constant in $\frac{RPM}{V}$. We found that the best performance was for gears of 64 teeth for the ball screws and 32 for the motor gear.

4.2.4 Inertial Measurement Unit

The IMU (Figure 4-7) is used to determine the absolute orientation in pitch, roll and yaw of AMOUR. The robot is neutrally balanced and consequently can float in a random orientation. The role of the IMU is to measure precisely its orientation. The robot can be controlled to hold a desired orientation using the angles derived by the IMU.

²The force on the piston is about 10500N at 40m depth.




Figure 4-6: AMOUR with the buoyancy and balance mechanism installed. The right side module of AMOUR is the buoyancy control mechanism. It is an actuated piston which can expand or reduce AMOUR's volume and thus change its buoyancy. AMOUR's balance control module is constituted by its battery pack (in the center of its body). The balance is changed by moving the battery pack's position with a lead screw.

The IMU uses 10 sensors to measure AMOUR's attitude and depth: three orthogonally mounted acceleration sensors, three solid state gyroscopes, three magnetometers and one pressure sensor.

The accelerometers are used to determine the direction of the gravitation vector, and thus the absolute orientation of the robot to vertical direction. The accelerometers provide an absolute measurement but they are noisy and prone to biased values due to the robot's acceleration during motion.

The gyros are used to measure the rate of turn on all axes. Their readings are integrated to provide orientation and are thus prone to drift over time. However their readings are robust relative to accelerations and have good signal to noise ratio.

The three magnetic sensors provide the direction of the Earth magnetic field. The IMU uses the magnetic vector and the orientation of the robot to compute the robot's tilt compensated heading.

The pressure sensor measures the depth, ascent and descent rates. All the sensor measurements are converted to digital values and filtered by the IMU's CPU, (which is a LPC2148 ARM7 processor). The CPU fuses the raw sensory data and computes




Figure 4-7: The Inertial Navigation Unit.

accurate orientation angles and depth of the robot ([Section 4.3.1](#)).

4.2.5 Central Controller Board

AMOUR's motions are coordinated by the Central Controller Board (CCB). The CCB receives commands from the sensor node high level mission control and sensory information from the IMU. The CCB sends commands to the 5 thrusters, buoyancy module and balance module. The CCB is responsible for running the low level control loops, maintaining the commanded configuration, attitude, depth and speed.

The CCB is built around an 32bit LPC2148 ARM7 processor running at 60Mhz for computation. A Cyclone II FPGA is used as communication co-processor. All the serial links between the LPC and the rest of robot's subsystems (thrusters, buoyancy/balance, IMU, sensor node) are implemented in the FPGA. All the serial links are buffered inside the FPGA. This greatly reduces the CPU time necessary for communication, and frees cycles for control loops calculations.

4.2.6 Sensor Node

AMOUR's high level mission control and sensing capabilities are coordinated by an AQUANODE unit attached to its body. We opted for this solution in order to maintain homogeneity in our system and reduce software complexity. The AQUANODE stores and executes the mission profile. It plans the robot motions and executes them by sending commands to the CCB. The robot's sensors are sampled by the AQUANODE and the data stored on the SD card. In addition, the robot can travel close to water surface, such that the AQUANODE's radio and GPS antennas are above the surface, and therefore operational (see [Figure 4-15](#)). The robot can receive commands and missions updates over the radio.

4.3 Algorithms and Control

In this section we describe the algorithms used by the robot to navigate and to carry payloads. The robot uses its actuators to actively maintain the desired pose based on the sensory input from the IMU, pressure sensor, localization network and GPS (when on surface).

4.3.1 Pose estimation

The robot uses the IMU to measure the absolute orientation of its body in Roll, Pitch and Yaw relative to the Earth's coordinate system (defined by the gravitation and the magnetic field). The three types of sensors used are: acceleration sensors, that measure static as well as dynamic accelerations, gyroscopes that measure rotations around the IMU's internal X, Y and Z-axis, and the magnetic field sensors which can be used to measure the earth's magnetic field.

The IMU uses three orthogonally placed acceleration sensors to measure the robot's acceleration. If the robot is static or in motion with constant velocity, the robot's acceleration corresponds to the gravitation vector and can be used to determine the roll and pitch angles. If the robot is accelerating or decelerating, the

acceleration sensors' readings do not accurately reflect the roll and pitch angles, as the gravitational acceleration is composed with robot's acceleration. The IMU also uses three orthogonally placed gyroscopes to measure the roll, pitch and yaw angles by integrating their outputs. The results are very accurate over short term and insensitive to translational accelerations. However, due to accumulated errors during integration the computer angles drift over time. The angles can drift as much as several degrees per minute. The IMU uses the magnetic field sensors to determine the robot's heading (yaw), by measuring the direction of the Earth's magnetic field. During motion the robot's body is not always horizontal. To get heading angle, the magnetic field vector is first projected on the horizontal plane, based on the roll and pitch angles. The direction of the projected vector is the used as robot's heading.

The acceleration and gyroscope measurements are fused in order to get accurate measurement of the roll and pitch angles. The Kalman Filter is the optimal linear estimator for this type of problems (cf. [44]), however, the matrix-operations required for a Kalman-Filter require significant CPU time. To achieve higher update rates for reduced latency, the IMU uses a simple and efficient algorithm presented in Algorithm 6.

Algorithm 6 Pose Estimation

- 1: read the accelerometers on all three axes X, Y, Z
 - 2: read the gyroscopes on all three axes X, Y, Z
 - 3: compute $(Roll_t^{(g)}, Pitch_t^{(g)}) = (Roll_{t-1}, Pitch_{t-1}) + \int gyroscopes$ the new gyro pose estimation
 - 4: compute $(Roll_t^{(a)}, Pitch_t^{(a)})$ based on accelerometer readings
 - 5: compute $(Roll_t, Pitch_t) = \alpha(Roll_t^{(g)}, Pitch_t^{(g)}) + (1 - \alpha)(Roll_t^{(a)}, Pitch_t^{(a)})$, the new pose estimation by reconciling the gyro and accelerometer readings
-

The algorithm computes the pose estimation as a weighted average between the gyroscopes estimation (integration of the rate of turn) and accelerometers estimation (assumes the acceleration vertical, causes just by gravitation). The weighting factor α was chosen close to one (i.e. heavy weight on gyro estimation) due to the immunity of gyro readings to robot's accelerations. The accelerometers estimation is used only to prevent the estimation drift due the gyro bias and integration.

Using the reliable estimate of the current pitch and roll angles stored in *angle_pitch* and *angle_roll* the heading is computed using the magnetic vector projection described above.

Simple functions using only fixed-point-numbers, lookup-tables and linear interpolation were programmed to do the computation of all trigonometric functions in a reasonable amount of time. These functions can be executed in about 600 cycles per second in the IMU’s microcontroller.

4.3.2 Hovering and Motion

AMOUR can hover and navigate in two orientations — vertical and horizontal — denoted by the robot’s body orientation. In both orientations the pose is actively maintained using the thrusters. In the horizontal configuration, the two thrusters perpendicular to the body are used for controlling depth and roll while the three thrusters parallel to the body are used to control pitch, yaw and forward/backward motion. In vertical configuration, the two thrusters perpendicular to the robot’s body are used for controlling yaw and forward/backward motion, while the three thrusters parallel to body are used for controlling pitch, roll and depth. In both configurations the robot is fully actuated, therefore pitch, roll, yaw, depth and forward/backward motion can be controlled independently.

Control Output	T0	T1	T2	T3	T4
Pitch	0	0	-1	$+\frac{3}{4}$	$+\frac{3}{4}$
Roll	0	0	0	+1	-1
Yaw	+1	-1	0	0	0
Depth	0	0	-1	-1	-1
Forward Speed	+1	+1	0	0	0

Table 4.1: Control output mapping for vertical orientation of the robot

The CCB uses the robot pose estimates computed by the IMU to control the pose. The pitch, roll, yaw and depth controllers are very similar. They are PD loops operating on the difference between the desired pose and the current pose estimated

Control Output	T0	T1	T2	T3	T4
Pitch	0	0	-1	$+\frac{3}{4}$	$+\frac{3}{4}$
Roll	-1	+1	0	0	0
Yaw	0	0	0	+1	-1
Depth	+1	+1	0	0	0
Forward Speed	0	0	+1	$+\frac{3}{4}$	$+\frac{3}{4}$

Table 4.2: Control output mapping for horizontal orientation of the robot

by the IMU. Figure 4-10 presents the depth and pitch controllers. The outputs of the PD controllers are used both for thruster commands and as an input for the buoyancy and balance mechanisms (Section 4.3.4). The roll and yaw controllers are similar, but their output is used only for thruster commands. The four controllers are run independently, and their output, depending on the the robot’s orientation is mixed for the correct thrusters. For example, in horizontal configuration the output of the roll controller is added to the output of the depth controller and sent to the two thrusters perpendicular to the robot’s body. The output of the depth controller is added equally to the two thrusters (i.e. they have to generate equal thrust for depth adjustment) and the output of the roll controller is added differentially to the two thrusters (i.e. they have to exert opposite thrust to generate roll motion). Table 4.1 and Table 4.2 present the complete controllers’ output mapping to the robot thrusters.

The values for the PD loop gains were determined experimentally. Compensations were added for the asymmetric thrust generated by the propellers (i.e. for the same rotation speed the propellers generate twice the forward thrust compared to back thrust).

Switching between the horizontal and vertical modes of operation is described by Algorithm 7. The robot starts with a request for a change in pitch which is followed by the thruster mapping that corresponds to vertical or horizontal configuration as appropriate.

For navigation between waypoints, the robot maintains roll and pitch angles of 0° and the specified depth profile. The heading is computed as a direct vector between the current position estimate and the desired waypoint. The speed is set as specified

Algorithm 7 Switch between horizontal and vertical orientation

- 1: request Pitch angle of 45°
 - 2: wait till the AMOUR's Pitch angle is 45°
 - 3: switch to the mapping table corresponding to the desired orientation (Table 4.1 or Table 4.2)
-

by the mission and mapped into thruster output based on the current orientation (Table 4.1 or Table 4.2).

4.3.3 Docking

We implemented a docking behavior which relies on the information available from the docking sensor (Section 4.2.2) to dock AMOUR with a sensor node. The behavior can be invoked if AMOUR is within the effective range of the optical beacon (approx. 90° cone 2.5m high). We envision using acoustic communication and ranging to guide AMOUR to within the optical docking range. The docking behavior is designed using the subsumption architecture [15, 18, 17, 16]. It is composed of several interacting state machines controlling the robot depth, orientation and XY position. We chose this architecture due to the non-linearity and non-monotony of our docking sensor, which made a linear controller impossible to use. For example, when the docking rod is close to the docking cone, the docking sensor readings become unreliable due to light reflexions and non-uniformity in the LED radiation pattern. In such circumstances the robot descends unconditionally 10cm and checks if the dock was successful. If not it moves up 25cm and attempts a new alignment and docking.

In summary the docking behavior attempts to align the robot with the sensor node in the XY plane. The alignment is performed by rotating the robot and moving forward or backward, as the robot is non-holonomic. If the alignment is within a specified threshold the robot descends slowly while maintaining the alignment. If the light is lost or the alignment is out of some specified bounds the robot ascends. If the light intensity reaches a specified threshold (the robot is closed to the rod), the robot ignores the reading of the sensors and attempts the open loop final docking step, described above. The C implementation is included in Appendix H.

4.3.4 Buoyancy and Balance




Figure 4-8: Adjusting the center of mass by moving the battery pack. From left to right: no payload, battery down (the robot is evenly balanced), no payload, battery up (the robot is top heavy), payload, battery down (the robot is bottom heavy) and payload, battery up (the robot is evenly balanced).

Suppose a payload of mass m (in water) is attached to the robot at depth d and carried for time t . If the robot uses thrusters only, the energy consumption is: $E = K_p \times m^2 \times t$, where K_p is a constant related to the thrusters efficiency (which is about $20W/kg^2$ in our system). The longer the payload is carried the more energy it uses. If instead the buoyancy engine is used the energy consumption is $E = K_b \times m \times d$, where K_b is a constant related to our buoyancy engine efficiency (which is about $35J/(kg \cdot m)$ in our system). Figure 4-9 shows the trade-offs between using thrusters only and using thrusters and a buoyancy engine. The energy does not depend on the time the payload will be carried, but does depend on the depth at which it is picked up. For example, for a 1kg payload collected at 10m depth the thrusters will use $20J/sec$ while the buoyancy engine will use $350J$ so the break even point is 17.5 secs. For the first 17.5 seconds the thrusters are more efficient beyond

- Thrusters

$$E = K_p \times m^2 \times t$$

- Buoyancy

$$E = K_b \times m \times d$$




Figure 4-9: Energy Analysis for a system that uses (a) thrusters (top curve) only and (b) thrusters and buoyancy engine (bottom curve). The x-axis shows the system’s energy consumption in Wh. The y-axis shows time in sec. The simulations were done for a payload of 1kg at 5m depth. The constants used were $K_p = 40W/(kg^2)$ and $K_b = 160J/(Kg * m)$.

that the buoyancy engine. But if the object is to be carried for 5min, the thrusters will use 20 times more energy than the buoyancy engine.




Figure 4-10: The control loops for the buoyancy and balance systems.

Figure 4-10a shows the control loop for the depth and buoyancy engine. The thrusters are controlled by a PD feedback loop that corrects for the desired depth. The buoyancy engine receives a low pass filtered version of the thruster command as input for a PID feedback loop that controls the piston’s position. The buoyancy engine moves the piston in the direction that brings the robot to neutral buoyancy (close to 0%) thruster output.

Figure 4-10b shows the control loop for the battery engine. The thrusters are

controlled by a PD feedback loop that corrects for the desired pitch. The battery motor receives a low pass filtered version of the thruster command. The battery is moved in the direction that brings the robot to neutral balance by a PID feedback loop that controls position.

4.4 Experiments

We conducted two types of experiments: in the pool and in the ocean. For pool tests we were able to mark the bottom of the pool and examine video data in order to extract the ground truth used to evaluate the results. During the ocean experiments the robot was driven close to the water surface so that a GPS antenna would gather information about the location of the robot. The relevant navigation data was logged and each experiment was also videotaped.

4.4.1 Control

The first series of experiments were done in MIT's Alumni Pool. The Pool is 4 m deep and 25 m long. The following experiments and tasks were evaluated in the pool:

- Controller performance tests
- Maximum speed and power tests
- Remote controlled travel
- Autonomous drive test

The controllers were tested by applying external forces to the robot while it was floating at a depth of 1.2 meters. The robot was manually tilted in all axes, one after another. Furthermore, it was forced to different depths. All input- and output-data was recorded using the control software's logging function. The whole experiment was also recorded by an underwater camera, whose time was synchronized to the timestamps in the log file. This was necessary to be able to assign the recorded

data-sets to the respective movement. With the data gathered in these experiments impulse-responses of all controllers could be plotted.

For the following experiments the robot was floating in horizontal orientation, with the thrusters on. Manual perturbations were applied on the three axes. Even though the control algorithms are the same in pitch, roll and yaw, there are differences in their behavior. This is due to a different thruster configuration in all axes. As seen in the following figures, all controllers are stable and precise.

- With the current parameter set there is no overshoot at all in the pitch control (Figure 4-11). The robot was pushed more than 30 degrees in this axis and it came back to the desired orientation within one second.
- The roll controller (Figure 4-12) performed well, but not as good. The robot was moved about 22 degrees in this axis and it took almost two seconds until it came back to neutral after an overshoot of about 7 to 8 degrees. The performance was considered acceptable.
- The yaw controller (Figure 4-13) works well. There is almost no overshoot and it takes about one and a half seconds to compensate for a perturbation of more than 45 degrees.
- The resolution of the pressure sensor is obvious in the impulse-response plot of the depth controller in Figure 4-14. The robot was supposed to stay at a depth of 1.2 meters. It did come back to its original depth after it was moved up about 60 cm. With the current pressure sensor's resolution it was difficult to get better controller performance, as the derivative signal used for the differential input is very noisy. The performance was considered acceptable.

The 25m pool was used to test the maximum speed of the robot underwater. The robot was commanded to stay at 1000mm depth while driving full speed from one end to the other in horizontal and vertical configuration. A video camera was used to




Figure 4-11: Impulse response of the pitch axis.

record the robot together with yard marks on the floor next to the pool to calculate the maximum speed.

The robot needed 16s to travel a distance of about 20m resulting in an average speed of $v_{avg} = 2.6knots$. The tests were repeated three times. In vertical configuration the speed was about 2/3 of the horizontal speed. Considering the size and maneuverability of AMOUR, this is a very high speed for an AUVs.

For the autonomous drive test a small square trajectory was programmed. For the square the robot was commanded to a depth of 500mm. 10 waypoints were used to describe the square and a turn of 90 at each corner. The entire path is shown in the following table.

The robot was commanded to follow the programmed path. AMOUR performed as expected and followed the waypoints very precisely.

4.4.2 Navigation

In this section we describe a suite of experiments designed to test AMOUR's ability to perform autonomous missions in natural ocean settings. We gave AMOUR the goal




Figure 4-12: Impulse response of the roll axis.

to drive a total of 5 autonomous missions during the same experiment. Figure 4-16 shows the path traveled by the robot during one of these autonomous missions. During this particular mission the robot was instructed to drive to mission markers distributed at the corners of a 50m x 50m square. Each drive was represented by a mission leg. During the first of these mission legs the robot was instructed to dive to a depth of 5m and travel using dead reckoning. The rest of the mission was driven at surface level using GPS for localization. Because in this mission the robot was too deep to log its GPS location and the underwater sensor network was not deployed along with the robot, there were no means of localization under water and the robot had to use dead reckoning. This is the reason why the robot did not estimate correctly the first waypoint, surfaced before reaching the marker for mission leg 1 (bottom right in Figure 4-16), and started to execute mission leg 2 (top right). The remaining three markers were reached to a desired error tolerance of 10m as planned. After reaching the last marker the mission controller switched off AMOUR's motors and we were able to lift the robot into the boat.

During the mission we were able to monitor execution within the mission user interface when we had a radio link with the robot. The robot drove a total of 200m for this particular mission (and each of the other four missions.) All five missions




Figure 4-13: Impulse response of the yaw axis.

were completed within two hours including programming of the mission and the robot deployment. During one of the missions the robot was sent to a depth of 10m and resurface after a total of 50m traveled at that depth.

In a second suite of ocean experiments we tested the robot’s endurance (e.g., battery life) and ability to travel in realistic environments. We commanded the robot to travel in the open ocean and measured its performance over a 1.65km trajectory. The robot traveled at surface level to facilitate feedback and trajectory corrections transmitted by radio. The robot was observed and followed from the boat. The water current was 0.5 m/sec in the opposite direction of travel and the wave surf was 0.5m from lowest to highest surface point. The robot traveled at the surface in order to collect ground truth for its trajectory in the form of GPS waypoints. As a result it was exposed to the waves and water currents because of windy weather. The current was moving in the direction of the open ocean and out the bay (Figure 4-17) and was thus working strongly against the robots intended direction of travel. The robot took 45 minutes to travel 1650m, which is equivalent to an absolute speed of 0.61 m/sec. Theoretically the robot should have traveled at 1 m/sec. The difference in speed coincides well with our estimate of water current speed. The battery was charged




Figure 4-14: Impulse response of the depth controller.

fully and used only for a few short system test runs before the actual experiment. After the endurance experiment the battery charge level was 50%. Thus with a fully charged battery the robot could have traveled 3.3 km against a water current of 0.5 m/sec and with a surf of 0.5m.

This experiment demonstrated the ability of the robot to travel long distances. The robot was able to travel on the surface in averse conditions. The need to travel more than 1000m can occur when the robot is autonomously deployed from the dock and has to travel to the measurement site and back. It can also occur when sensor readings are needed over a long transect. The robot can traverse the entire transect that needs to be measured and collect the necessary sensor data.

4.4.3 Docking

We performed two types of experiments to evaluate the optical docking performance of AMOUR.

For the first set of experiments we placed a sensor node on the bottom of the pool. We deployed AMOUR in the vicinity of the sensor. The robot performed a spiral search pattern based on dead reckoning, in order the acquire the optical beacon. Once the optical beacon was located AMOUR switched the docking procedure presented in

Waypoint No.	Roll	Pitch	Yaw	Depth	Forward speed	Time [1/100s]
0	0	0	0	500mm	0%	400
1	0	0	0	500mm	20%	200
2	0	0	90	500mm	0%	400
3	0	0	90	500mm	20%	200
4	0	0	180	500mm	0%	400
5	0	0	180	500mm	20%	200
6	0	0	270	500mm	0%	400
7	0	0	270	500mm	20%	200
8	0	0	360	500mm	0%	400
9	0	0	360	500mm	20%	200

Table 4.3: Waypoints used for the autonomous drive test

[Section 4.3.3](#). AMOUR docked the sensor node, picked it up, and moved it to a new location. We repeated the experiment more than 30 times. Images of the robot during the docking experiment are presented in [Figure 4-18](#). We increased the difficulty of the task by manually moving the sensor node or the robot during various stages of the docking, simulating unexpected waves. The robot successfully docked the sensor node in most of experiments (approx. 80%). We observed two failure modes: (1) the docking latch failed to grip properly onto the docking rod — this was a mechanical design problem and was fixed during later revisions of AMOUR, (2) if the sensor node was significantly tilted the light direction caused the docking state machine to move the robot around the sensor node without ever approaching it fully — this was not a problem for the later revisions of the sensor nodes which are moored and balanced to always stay vertical. We measured the docking optical beacon range to 2.5m above the node and approximately 90 degrees angle (in clear water).

For the second set of experiments [\[37\]](#) placed a sensor node on a second AUV, Starbug [\[36\]](#). AMOUR was programmed to dock the sensor node while Starbug was hovering maintaining depth and location. This experiment simulates a mission in which the two AUVs (AMOUR and Starbug) would cooperate to deploy and recover sensor nodes. Starbug’s task is to carry the sensor nodes to the destination area. Once

there, AMOUR's task is to pick up the nodes, one by one, and deploy them to their specified locations. For recovery, AMOUR picked up the nodes up from their sensing locations and placed them on Starbug's top side. Images from the experiments are presented in [Figure 4-19](#). The experiment was repeated over 20 times. Only 4 failures were noted due to the failed latch or the spiral search failing to locate the beacon. The latch problems were fixed in the next robot's revision, while we envision to replace the spiral search with an acoustic based approach to within optical range.

4.4.4 Buoyancy and Balance

Several experiments were done to evaluate the buoyancy and balance control. The buoyancy experiment examines the aggregate thruster output as the robot's weight is changed by adding or removing a weight of 950g. [Figure 4-20](#) shows the results. At time $t = 0$ the robot hovers at constant depth in vertical configuration. The thrusters oscillate around 0% output. The robot's depth is maintained at 0.5m (+/- 2cm) below the water surface. The depth controller output oscillations are due to the latency of the sensorless motor controllers while changing direction (we note that the larger oscillations are due to human touch in preparation for adding a weight). At time $t = 100$ s a weight of 950g is added to the robot. The depth controller oscillates till time $t = 110$ s in response to this event. Between times $t = 110$ s and $t = 170$ s the depth controller is stable at 40% output to compensate for the additional weight. We note that the oscillations are smaller in this phase since the controller does not need to change the direction of rotation as in the neutrally buoyant case. At time $t = 170$ s the buoyancy engine is activated. Between $t = 170$ s and $t = 185$ s the buoyancy engine effects work to compensate for the additional weight of the robot. At time $t = 185$ s the robot is again neutrally buoyant and the depth controller operates around 0%. At time $t = 230$ s the robot releases the attached weight and the depth controller starts compensating for the positive buoyancy of the robot. At time $t = 250$ s the buoyancy engine is activated again to bring the robot back to neutral buoyancy which happens at time $t = 265$ s. We have repeated the buoyancy control experiment 10


times with similar results. The average time required to compensate for the robot's weight change is 15s.

The next group of experiments illustrates the effect of moving the battery inside robot on the pitch controller. [Figure 4-21](#) shows the data. The robot starts in horizontal configuration with a 950g weight attached to its bottom. The battery is at its bottom position in the robot's cylindrical enclosure. The thrusters work at -40% to compensate for the weight. At time $t = 10\text{s}$ the battery begins moving toward its top-most position. It takes 25s to complete the move. We observe the controller output at 10% (in the other direction). The movement is repeated back and forth.

A third suite of experiments concerns balance control. [Figure 4-22](#) shows the data. In these experiments the robot starts by hovering in horizontal configuration with the battery placed at the bottom of the robot's cylindrical body. The thrusters keep the robot hovering in horizontal stance. At time $t = 20\text{s}$ the battery is moved to its mid-point in the body of the robot and subsequently we observe the effect on the pitch controller's output. At time $t = 45\text{s}$ a weight of 950g is added manually at the rear end of the robot and again we observe the effect on the pitch controller's output. At time $t = 67\text{s}$ the robot starts moving the battery autonomously to adjust the robot's balance. At time $t = 85\text{s}$ we observe that the robot has successfully achieved neutral balance and subsequently the pitch controller oscillates around 0. The balance control experiment was repeated 10 times with similar performance. The average time to achieve neutral balance was 20s.

4.5 Summary

This chapter presented the robotic component of our system which enables data muling, sensor network deployment, relocation and recovery. We presented the hardware and the control algorithms that enable navigation, optical guided docking, and energy efficient transport of dynamic payloads. These tasks are performed to support the activity of the underwater sensor network.




(a)




(b)

Figure 4-15: AMOUR during a (a) daylight mission and (b) night mission. At night the advantage was the higher visibility of the illuminated robot (for easy recovery).



(a)



(b)

Figure 4-16: Plot of an autonomous mission path (a) and depth profile (b). The crosses denote the sides of a 50m by 50m square that the robot was told to drive. The dashed path was a dive at 5m depth while the other three sides were driven on the surface. Because the dive was executed with dead reckoning the robot surfaced early and started the next mission leg.




Figure 4-17: GPS path during endurance mission. The robot was driven by a human operator through a wireless link. The robot traveled at 75% of its maximum speed for a total of 45 minutes. The distance traveled was 1650 meters from the test site to the research station dock.





Figure 4-18: Autonomous docking experiment (frames from the experiment video). (a) AMOUR is executing a spiral search pattern, looking for sensor node's optical beacon. (b) AMOUR is descending to the sensor node while using the optical beacon to maintain XY alignment. (c) AMOUR is docked and is locking the attachment latch. (d) AMOUR is carrying the sensor node to a new location.



(a)

(b)

Figure 4-19: AMOUR docking StarBug




Figure 4-20: Buoyancy experiment: thruster controller output (in percentage) over time as we vary the robot's weight and adjust the buoyancy




Figure 4-21: Battery experiment: thruster controller output (in percentage) over time as we change the robot's balance by moving the battery's position.




Figure 4-22: Balance experiment: thruster controller output (in percentage) over time as we change the robot's balance by adding weight.

Chapter 5

Conclusions

In this thesis we proposed a new approach for marine data collection and retrieval which provides new types of data and simplifies the task of the researcher, through automation. Our core innovation is the extensive use of light and mobility for information transfer underwater. We use light for sensing, communication and manipulation.

It is hard not to notice the similarity between our approach and nature's approach to sensing underwater. Nature may have chose to evolve "eyes" for the 380nm-750nm region of the electromagnetic spectrum, precisely because of its propagation properties in water (i.e., low attenuation). As humans, we gather most of our information through our eyes — and we believe automated systems can do the same.

5.1 Contributions and Lessons Learned

We built and integrated devices to test and demonstrate our concepts. The system was build ground up with very few pre-made components, in an attempt (1) to avoid proprietary, hard to modify parts and (2) to better understand the technology limits, costs and constraints of building underwater equipment. In doing so, we noticed the complexity of a system growing exponentially (and not linearly) with the number of features and components. This is due to increasing the number of interactions and also to the fact that the probability of the system working correctly being the product of all its parts' probabilities of working correctly. The problem is aggravated in the

ocean which is a harsh and unforgiving environment. A single failure can result to equipment loss and important setbacks.

Because of these experimental challenges, we focused on extensively test every part of the system separately. Even simple basic mechanical components such as watertight cases and robot thrusters went through many revisions till they became reliable and tolerant to operator errors. We understood the why the industry appears to proceed cautiously. Even if a new component has higher potential, an older component that is tried and tested is easily preferred. The high penalty for failure is a key factor in the apparent slow development of the field. Advances in the underwater monitoring can be realized only when there will be enough researchers who will work to perfect every individual component of the system.

We built a new device (i.e, an underwater flash) that enables true color underwater imaging. Our device illuminates the underwater objects with an adjustable spectrum light that compensates for the color distortions caused by water. We demonstrated an algorithm for optimizing the light composition necessary to take accurate color images with our device. The algorithm uses the existing knowledge about light propagation underwater and the distance between the subject and the camera. We extensively tested our device in pool and ocean experiments and compared it with existing techniques. One aspect that we have not addressed is measuring the distance from the camera to the subject. In our experiments we measured the distance manually or we tried to use acoustic range sensors. We found that the acoustic range sensors are inadequate: they do not have the precision required and are prone to errors as not all underwater objects reflect sound (e.g., soft coral). A better system would have been to use the distance information optically, from the camera focus system. This was not possible as all the SLR cameras we investigated were proprietary and do not make this information available. A laser range sensor can be envisioned as well.

We built and analyzed underwater optical communication modems. On short distances they provide orders of magnitude higher data rates and lower power consumption than traditional acoustic systems. We chose to develop optical communication as it was the only practical information carrier capable of high data rates and low

power consumption (due to relatively low attenuation in water). We demonstrated the optical communication system as a component of an underwater robotic data muling process. This optical communication approach is scalable with the network size and information density, especially when compared to acoustic networks.

Without doubt the optical systems' performance underwater depends on the water clarity. For some applications (e.g., reef studies in South Pacific) this is not a critical issue due to year-round high water clarity. For other applications (e.g., port security) we envision the problem being mitigated by the robotic and acoustic part of the system. For example, to download data the robot needs to be come within the visibility distance of the sensor node. Even very polluted rivers (e.g., Charles River in Boston) have at least 90cm of visibility. The robots can be guided acoustically to within that distance.

We built and experimented with an underwater acoustical modem. The goal of our acoustic system was to provide a broadcast environment for signaling, status, data, localization and tracking. We demonstrated the use of our acoustic network for self-localization, tracking of a mobile robot — tasks which simplify the researcher's job and the robot's navigation. We also designed a low power modulation (sparse symbols) and a low power medium access protocol (self synchronizing TDMA) which trade off data rates for lower energy per transmitted bit. Lower energy per bit translates into longer deployments for the sensor network. Lower data rates were acceptable as direct consequence of using optical communication for the bulk of the data transfer.

We integrated sensors, computation, data logging and communication into our underwater sensor network prototype which we tested in the pool and in ocean experiments. Our platform is an excellent test bed for exploring the possibilities of underwater sensor networks. Our sensor nodes are small, flexible, easy to transport, program and deploy. Their endurance is on the order of weeks (in continuous acoustic transmission mode) to months (in low power modes).

Finally we designed, fabricated and tested an underwater robot AMOUR with several innovative features. AMOUR is a hovering robot capable of both precise maneuvering and efficient long distance travel. We demonstrated docking and sensor

node pick up behavior which are key for data mulling and autonomous deployment, relocation and retrieval of underwater sensor networks. In addition our robot's internal buoyancy and balance mechanism enable it to carry dynamic payloads efficiently, up to a significant fraction of its size. We demonstrated experimentally the robot's interaction with the sensor network: docking, localization and tracking. We also demonstrated the robot's ability to carry the underwater camera and adaptive flash system and autonomously take pictures of coral reef areas.

5.2 Near Future Goals

Our vision for persistent maritime surveillance encompasses a fully integrated system, which is yet to be done. We demonstrated so far all the pieces of the systems, as well as some integrated experiments. A viable integrated system requires significant additional software infrastructure to support the integration and coordination between the system components, the high-level plans, and the mission interface. This software can only be reliable and usable in field tests with a systematic approach to testing, debugging and fault tolerance. An experiment typically takes hours. A very useful feature for streamlining such experiments is the ability to freeze the experiment to its current status, correct software issues and restart it from where it was left. A simulation platform is another alternative. Having such tools is the only way to make significant progress toward a functional complex system capable of persistent surveillance.

In addition there are many improvements we can make to the individual parts of our system. For the underwater flash we hope to better integrate it with the camera and get the distance information. For the sensor nodes we envision a winch system which will enable them to adjust their depth for better acoustic communication, water column sampling, radio and GPS on the surface, solar recharging. We envision higher data rates and higher sensitivity for the optical modem in order to cope with turbid waters. Finally our robot needs better control loops to deal with waves and currents, acoustic servoing for closing in to the sensor nodes, and simpler mechanical design

for easy transport, assembling and disassembling.

5.3 Future Research Directions

The field of underwater sensing and data retrieval is very young and much remains to be explored. Several long-term results would make a difference.

For the underwater illumination system we envision a LED-based continuous illumination version for videos. We would also like to see it integrated with methods coping with water turbidity. In addition, our correction method can be integrated with 3D reconstruction techniques for color correcting highly three-dimensional objects.

For the sensor nodes we would like to implement methods of harvesting power from the environment (e.g., tides, waves) and extending their autonomy. For the acoustic network we would like to implement distributed error correcting codes, which would improve the network throughput while keeping the latency low. Finally we envision longer range optical modems based on frequency modulation which will reduce the constraints on the AUV maneuverability capabilities.

Finally, AUV navigation underwater remains an open topic. We require our AUV to perform significant autonomous navigation tasks with the aid of acoustic and optical beacons. The potential of existing techniques localization and mapping techniques (e.g., SLAM) can be explored in this unique environment.

Appendix A

Adaptive Light Optimization

```
% computes the relative power of the strobes for a given distance
% between camera and subject and a given white power
function [CRI coefs times color_dist white_CRI white_color_dist ] ...
    = optimize_CRI( depth, white_power )
% loads saved flash output function and output spectra
[ref_times wavelengths spect_pval] = get_radiance_function;
max_time      = ref_times(size(ref_times,2));
min_time      = ref_times(1);
lights        = size(spect_pval.spectrums, 1);
%white flash position in the device
white_position = 4;
%salt water absorption for the specified depth
absorption    = get_saltwater_absorption(depth, wavelengths);
% load CRI test colors and reference illuminant
[ref_light, XYZBar, TestColors] = init_cri(wavelengths);
ref_light = get_total_radiance([0 0 0 1000 0 0], spect_pval);
white_time  = get_time_from_power(white_power, spect_pval.pval);
initial     = ones(1, lights) * max_time * 0.01;
low_bound   = ones(1, lights) * min_time;
upper_bound = ones(1, lights) * max_time;
low_bound(white_position) = white_time;
upper_bound(white_position) = white_time;
```

```

initial    (white_position) = white_time;
options = optimset('MaxFunEvals',4000);
% the optimization
[times CRI] = fmincon( ...
    @(x) -compute_CRI( get_total_radiance(x, spect_pval) ...
    .* absorption, ref_light, wavelengths), ...
    initial, [], [], [], [], low_bound, upper_bound, [], options);

CRI = -CRI;

%CRI calculation
function [CRI color_dist] = compute_CRI(light, ref_light, wavelengths)
    [CRI color_dist] = ...
        compute_CRI_with_gain(light, ref_light, wavelengths, [1 0 0 ; 0 1 0 ; 0 0 1]);

function [CRI color_dist] = ...
    compute_CRI_with_gain(light, ref_light, wavelengths, gain)
unity_vector = ones(size(wavelengths));
[ref_light_old XYZ_bar test_colors] = init_cri(wavelengths);
[Y_ref u_ref v_ref c_ref d_ref] = ...
    compute_Yuvcd(ref_light, test_colors, XYZ_bar);
[Y_src_ref u_src_ref v_src_ref c_src_ref d_src_ref] = ...
    compute_Yuvcd(ref_light, unity_vector, XYZ_bar);
[W_ref U_ref V_ref] = ...
    compute_WUV(Y_ref, u_ref, v_ref, Y_src_ref, u_src_ref, v_src_ref);

[Y u v c d] = ...
    compute_Yuvcd(light, test_colors, gain * XYZ_bar);
[Y_src u_src v_src c_src d_src] = ...
    compute_Yuvcd(light, unity_vector, gain * XYZ_bar);
[W U V] = ...
    compute_scaled_WUV(Y, c, d, Y_src, c_src, d_src, ...
        u_src_ref, v_src_ref, c_src_ref, d_src_ref);
dE = sqrt(sum(( [W U V] - [W_ref U_ref V_ref] ) .^ 2, 2));
Ri = 100 - 4.6 * dE;
color_dist = [W U V] - [W_ref U_ref V_ref];

```

```
CRI = mean(Ri);

%flash output spectrum calculation
function radiance = get_total_radiance(times, spect_pvals)
    radiance = zeros(1, size(spect_pvals.spectrums,2));
    for i=1:size(times,2)
        radiance = radiance + spect_pvals.spectrums(i,:)*ppval(spect_pvals.pval, times(i));
    end
end
```


Appendix B

Flash Optimization Results

```
1 // flash durations in micro-seconds for distances
2 // between 0 and 5m, in 10cm increments
3 // each line contains the 6 durations, the the 6 flashes
4 { 0, 0, 1, 0, 1, 0},
5 { 0, 0, 0, 6, 3, 0},
6 { 4, 0, 0, 14, 0, 0},
7 { 5, 0, 10, 21, 9, 0},
8 { 7, 0, 12, 27, 13, 0},
9 { 10, 0, 16, 35, 17, 0},
10 { 15, 0, 0, 44, 0, 13},
11 { 17, 0, 23, 52, 27, 0},
12 { 21, 0, 27, 57, 32, 4},
13 { 26, 0, 31, 61, 39, 5},
14 { 25, 26, 36, 63, 28, 29},
15 { 36, 0, 40, 67, 53, 6},
16 { 41, 0, 45, 70, 58, 7},
17 { 47, 0, 49, 73, 62, 8},
18 { 54, 0, 54, 77, 66, 9},
19 { 57, 0, 57, 80, 70, 15},
20 { 61, 0, 60, 84, 74, 20},
21 { 64, 0, 63, 88, 80, 18},
22 { 68, 0, 66, 92, 85, 19},
```

23 { 71, 0, 69, 97, 92, 17},
 24 { 75, 0, 71, 101, 99, 15},
 25 { 79, 0, 74, 106, 109, 15},
 26 { 83, 0, 77, 111, 118, 12},
 27 { 88, 0, 80, 116, 128, 12},
 28 { 92, 1, 83, 121, 138, 7},
 29 { 98, 3, 86, 127, 151, 5},
 30 { 104, 4, 89, 132, 163, 4},
 31 { 111, 4, 93, 138, 177, 6},
 32 { 117, 1, 97, 143, 191, 21},
 33 { 104, 128, 103, 148, 100, 108},
 34 { 132, 2, 105, 154, 226, 10},
 35 { 140, 2, 110, 160, 246, 12},
 36 { 148, 4, 115, 166, 268, 15},
 37 { 157, 2, 119, 172, 293, 15},
 38 { 166, 1, 124, 178, 320, 16},
 39 { 143, 190, 133, 184, 148, 157},
 40 { 151, 204, 139, 190, 158, 168},
 41 { 158, 217, 144, 196, 168, 179},
 42 { 166, 233, 150, 203, 179, 192},
 43 { 175, 250, 156, 210, 191, 206},
 44 { 183, 268, 162, 216, 199, 220},
 45 { 191, 290, 168, 223, 207, 238},
 46 { 202, 314, 175, 230, 221, 256},
 47 { 211, 339, 182, 237, 230, 276},
 48 { 214, 339, 189, 244, 188, 370},
 49 { 229, 352, 196, 252, 232, 383},
 50 { 3895, 1277, 2254, 259, 571, 980},
 51 { 3946, 1319, 2052, 267, 610, 1015},
 52 { 270, 480, 219, 274, 298, 392},
 53 { 284, 511, 227, 282, 316, 417},
 54 { 299, 541, 237, 291, 333, 444},
 55 { 314, 570, 246, 299, 348, 470},
 56 { 330, 599, 256, 307, 363, 498},
 57 { 347, 633, 266, 315, 382, 522},
 58 { 364, 666, 277, 324, 399, 552},

59 { 382, 699, 288, 332, 422, 581},
60 { 401, 734, 300, 341, 448, 616},
61 { 3904, 1266, 1146, 350, 589, 986},
62 { 440, 814, 325, 358, 501, 697},
63 { 461, 865, 338, 367, 530, 739},
64 { 483, 914, 352, 376, 566, 784},
65 { 506, 968, 365, 385, 610, 830},
66 { 527, 1027, 379, 394, 659, 881},
67 { 3973, 1197, 852, 402, 559, 959},
68 { 566, 1190, 407, 410, 761, 1006},
69 { 585, 1298, 421, 419, 822, 1086},
70 { 605, 1429, 436, 428, 890, 1169},
71 { 625, 1578, 451, 437, 951, 1264},
72 { 645, 1742, 467, 445, 1001, 1367},
73 { 667, 1947, 482, 454, 1059, 1494},
74 { 692, 2527, 497, 463, 1148, 1581},
75 { 723, 2713, 511, 471, 1329, 1781},
76 { 4000, 1199, 599, 479, 515, 905},
77 { 4000, 1127, 577, 487, 479, 850},
78 { 855, 3753, 547, 496, 2731, 4000},
79 { 956, 3664, 543, 504, 2653, 4000},
80 { 1158, 3680, 556, 513, 2230, 4000},
81 { 4000, 3880, 664, 520, 1998, 4000},
82 { 4000, 979, 517, 527, 420, 765},
83 { 1792, 3799, 581, 534, 1678, 4000},
84 { 2307, 3809, 589, 542, 1388, 4000},
85 { 3219, 3932, 599, 549, 1612, 4000},
86 { 3965, 4000, 590, 556, 1583, 3998},
87 { 4000, 3879, 568, 563, 1290, 4000},
88 { 3979, 4000, 548, 569, 1463, 3970},
89 { 4000, 3983, 522, 576, 1237, 4000},
90 { 3955, 4000, 508, 583, 2573, 3996},
91 { 3964, 3952, 478, 590, 2037, 4000},
92 { 3985, 3938, 444, 597, 1805, 4000},
93 { 4000, 3937, 403, 604, 1584, 4000},
94 { 3969, 4000, 354, 610, 1228, 3608},

95 { 4000, 3899, 315, 617, 1263, 4000},
96 { 3984, 3922, 0, 624, 1239, 4000},
97 { 4000, 3924, 246, 631, 1276, 4000},
98 { 4000, 3947, 214, 639, 1209, 4000},
99 { 3978, 3712, 195, 646, 2917, 4000},
100 { 3979, 3728, 166, 653, 2930, 4000},
101 { 3999, 4000, 0, 661, 1188, 3846},
102 { 3995, 3837, 0, 669, 1585, 4000},
103 { 3979, 3784, 0, 677, 2952, 4000},
104 { 4000, 4000, 2, 685, 1061, 3300}

Appendix C

Adaptive Flash Schematics

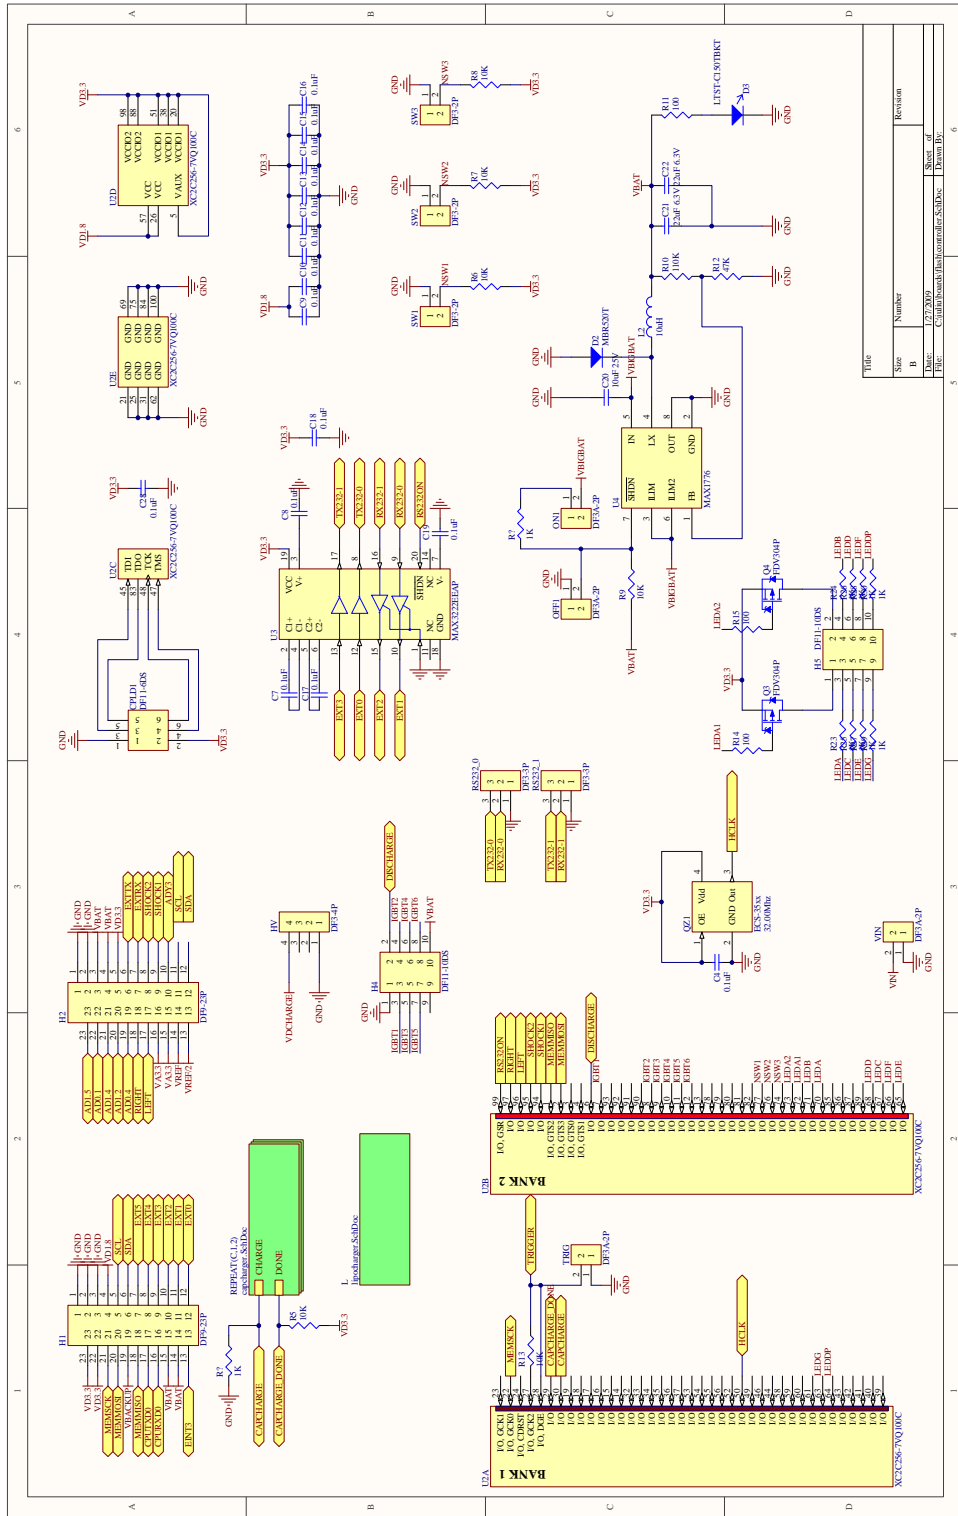




Figure C-1: Timing Controller 1



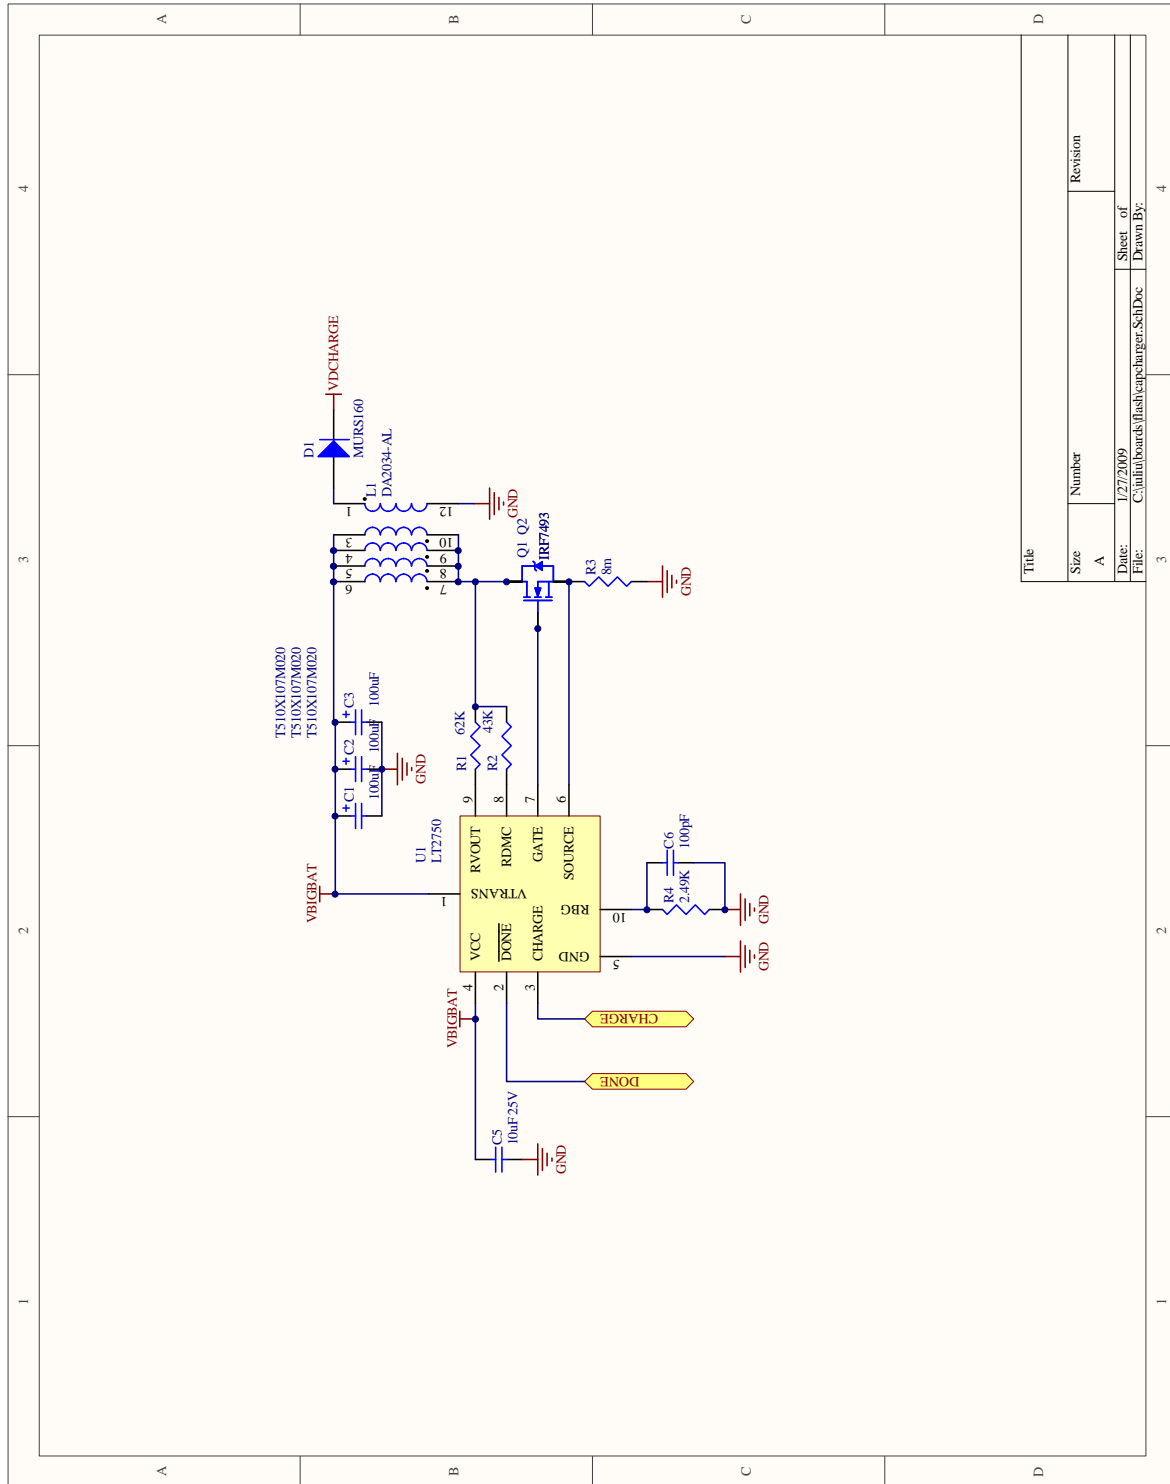
Title	
Size	Revision
A	
Date:	1/27/2009
File:	C:\initrigger board.SchDoc
Sheet of	4
Drawn By:	

Figure C-2: Timing Controller 2



Title	
Size	Revision
A	
Date: 1/27/2009	
File: C:\info\boards\flash\minger.SchDoc	
Sheet of 3	
Drawn By: 4	

Figure C-3: Timing Controller 3



Title	
Size	Revision
A	
Date:	1/27/2009
File:	C:\info\boards\flash\capcharger.SchDoc
Sheet of	3
Drawn By:	

Figure C-4: Voltage Inverter

Appendix D

Flash Timing Control

```
`timescale 1ns / 1ps

module main(CS, CLK, DI, TRIGGER, RCLK, IGBT);
    parameter FLASHES = 6;
    parameter COUNTER_SIZE = 10;
    localparam DATA_SIZE = FLASHES * COUNTER_SIZE;
    input CS, CLK, DI, TRIGGER, RCLK;
    output [FLASHES-1:0] IGBT;
    reg [FLASHES-1:0] IGBT;
    reg [COUNTER_SIZE-1:0] counter;
    reg state;

    wire [COUNTER_SIZE-1:0] stop_reg [FLASHES-1:0];
    wire [DATA_SIZE-1:0] spi_data;
    genvar i;
    generate
        for(i=0;i<FLASHES;i=i+1) begin: a
            assign stop_reg[i] = spi_data[(i+1)*COUNTER_SIZE-1:i*COUNTER_SIZE];
        end
    endgenerate
endmodule
```

```

spi #(.DATA_SIZE(DATA_SIZE)) sp(.CS(CS), .CLK(CLK), .DI(DI),
    .OUT(spi_data));

integer j;
always @(posedge RCLK) begin
    if(TRIGGER) begin
        counter <= counter + 1;
        if(state == 0) begin
            for(j=0;j<FLASHES;j=j+1)    IGBT[j] <= 1;
            state <= 1;
        end else begin
            for(j=0;j<FLASHES;j=j+1)
                if(counter == stop_reg[j]) IGBT[j] <= 0;
        end

        end else begin
            state <= 0;
            OUT <= 0;
            counter <=0;
        end
    end
endmodule

// spi module, for timing data upload
module spi(CS, CLK, DI, OUT);
    parameter DATA_SIZE = 10;
    input  CS, CLK, DI;
    output [DATA_SIZE-1:0]OUT;

    reg    [DATA_SIZE-1:0]OUT;

    always @(negedge CLK) begin
        if(~CS) begin
            OUT <= { OUT[DATA_SIZE-2:0] , DI};
        end
    end
end

```

```
end  
endmodule
```


Appendix E

AQUANET Schematics

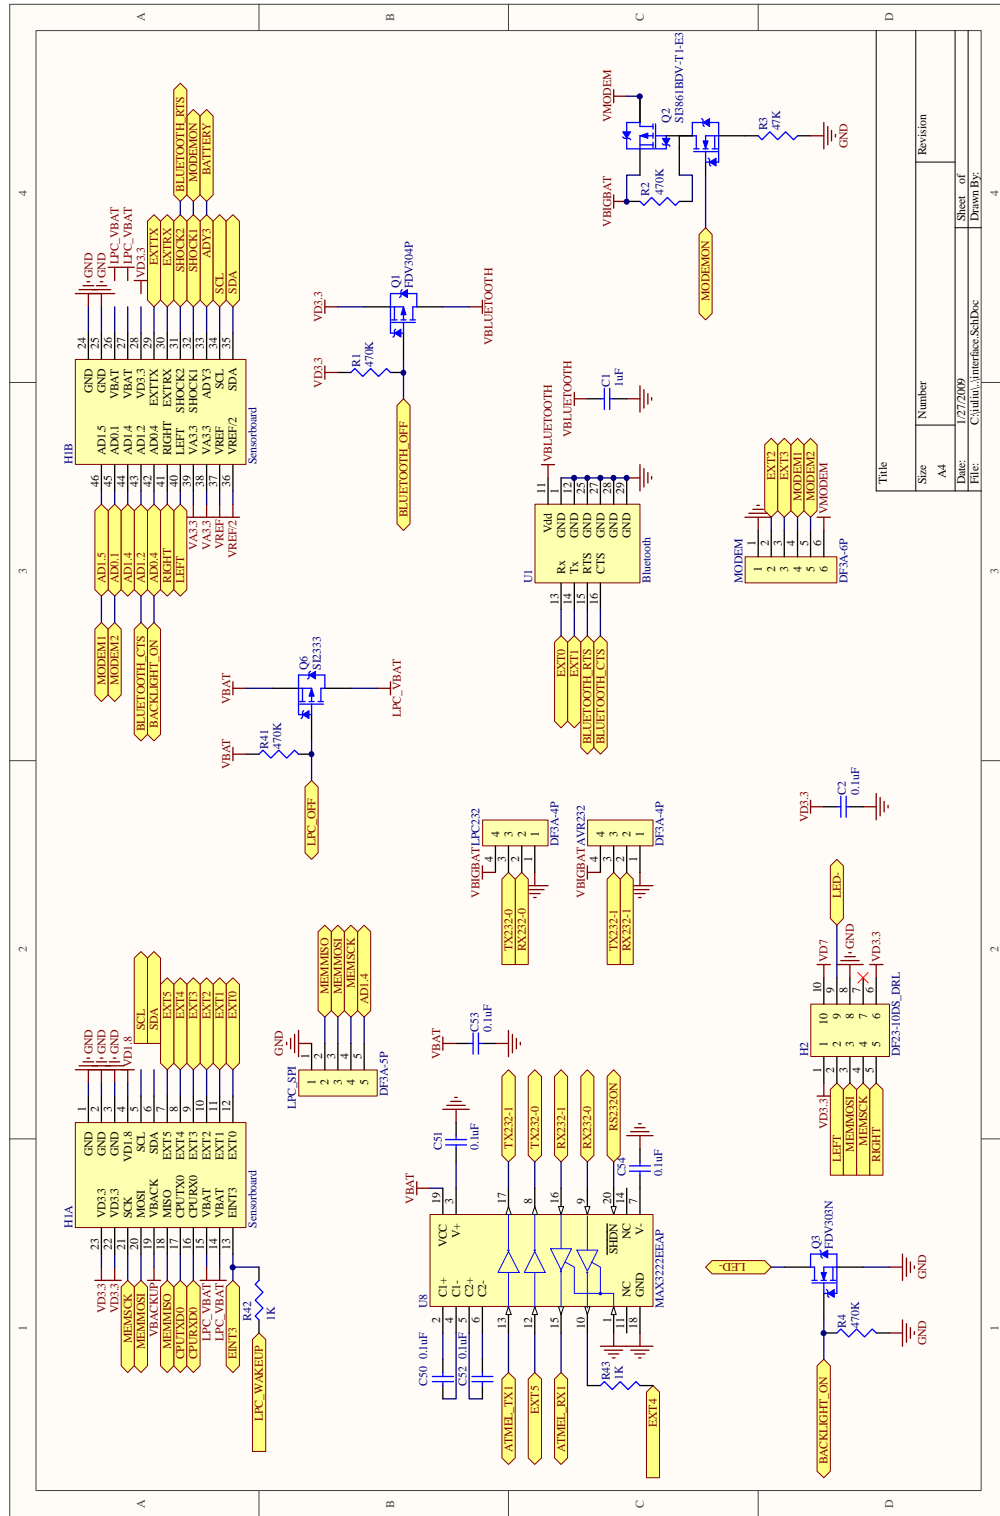


Figure E-1: AQUANODE Interface

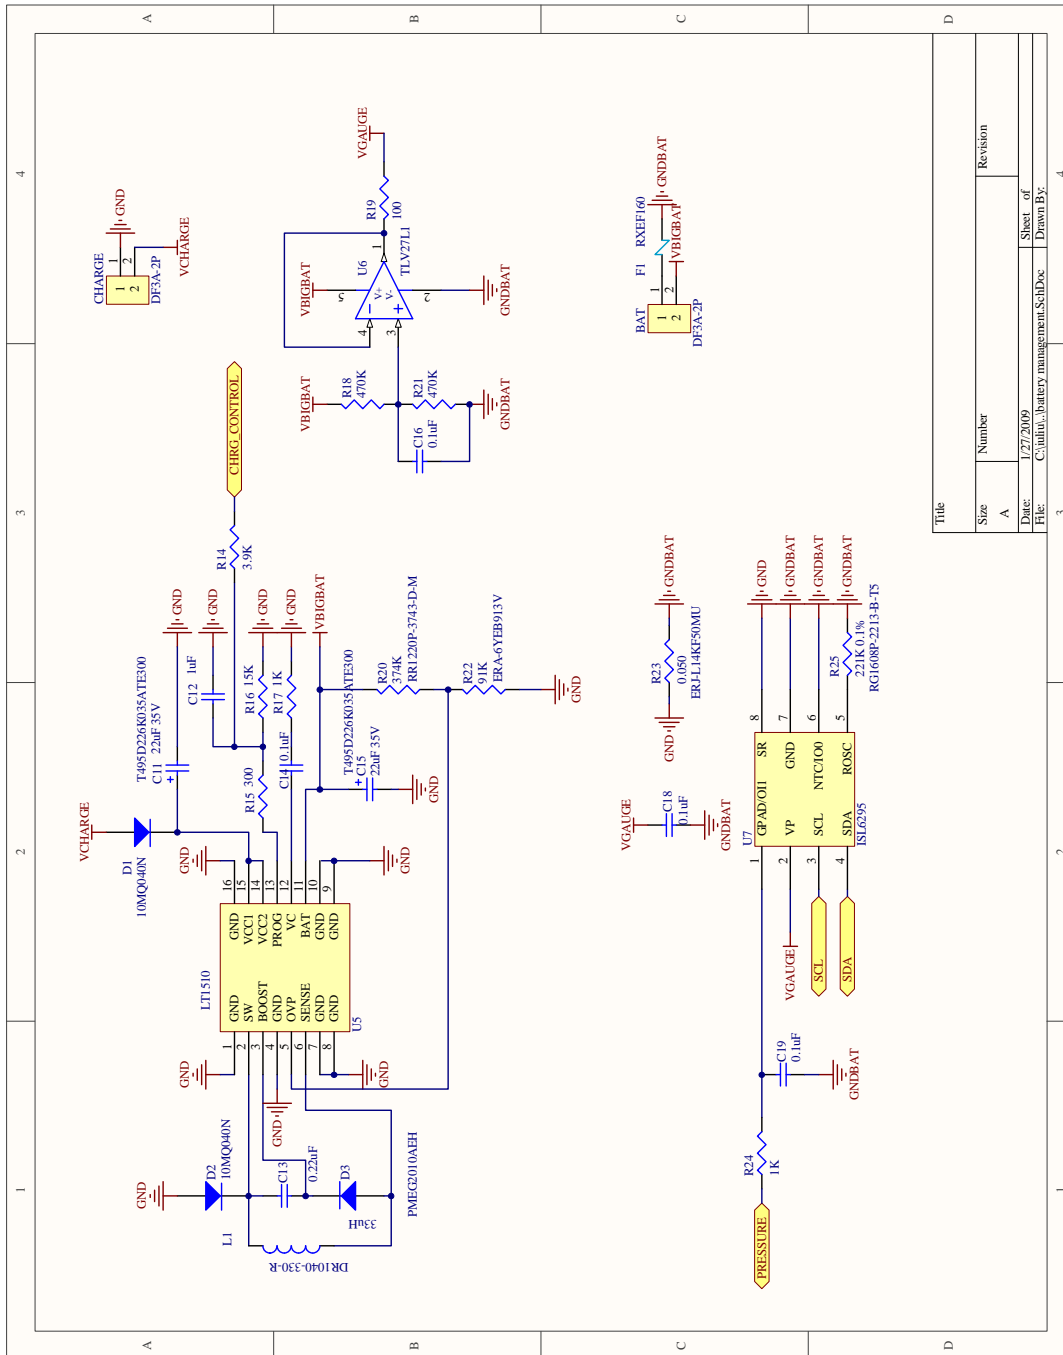


Figure E-2: AQUANODE Battery Management




Figure E-3: AQUANODE Pressure Sensors

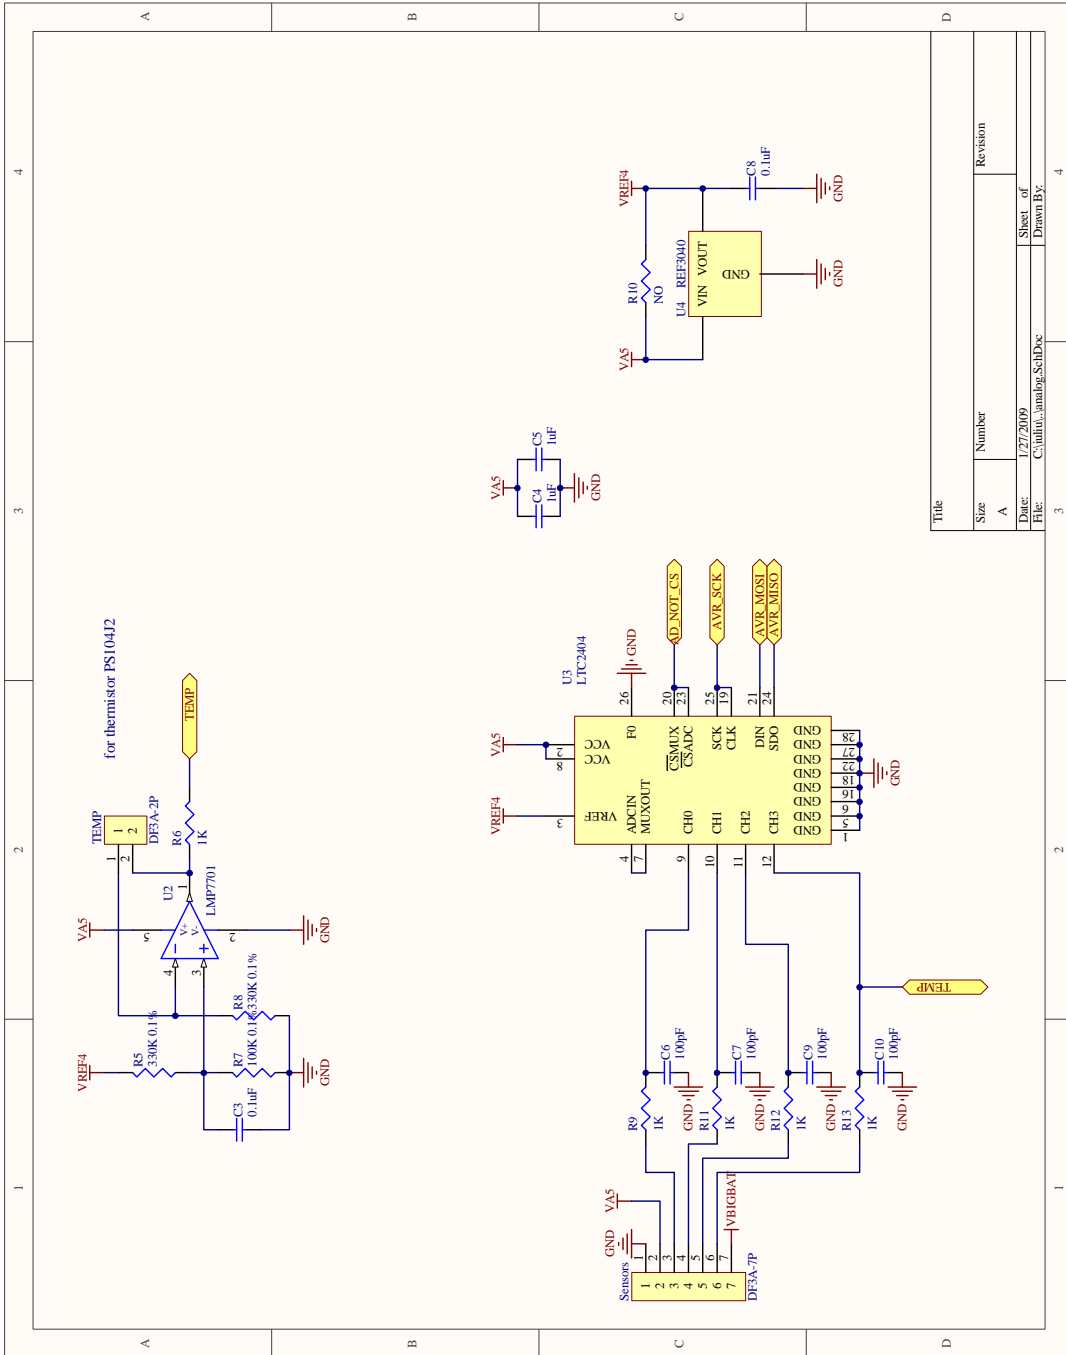


Figure E-4: AQUANODE A/D Converter

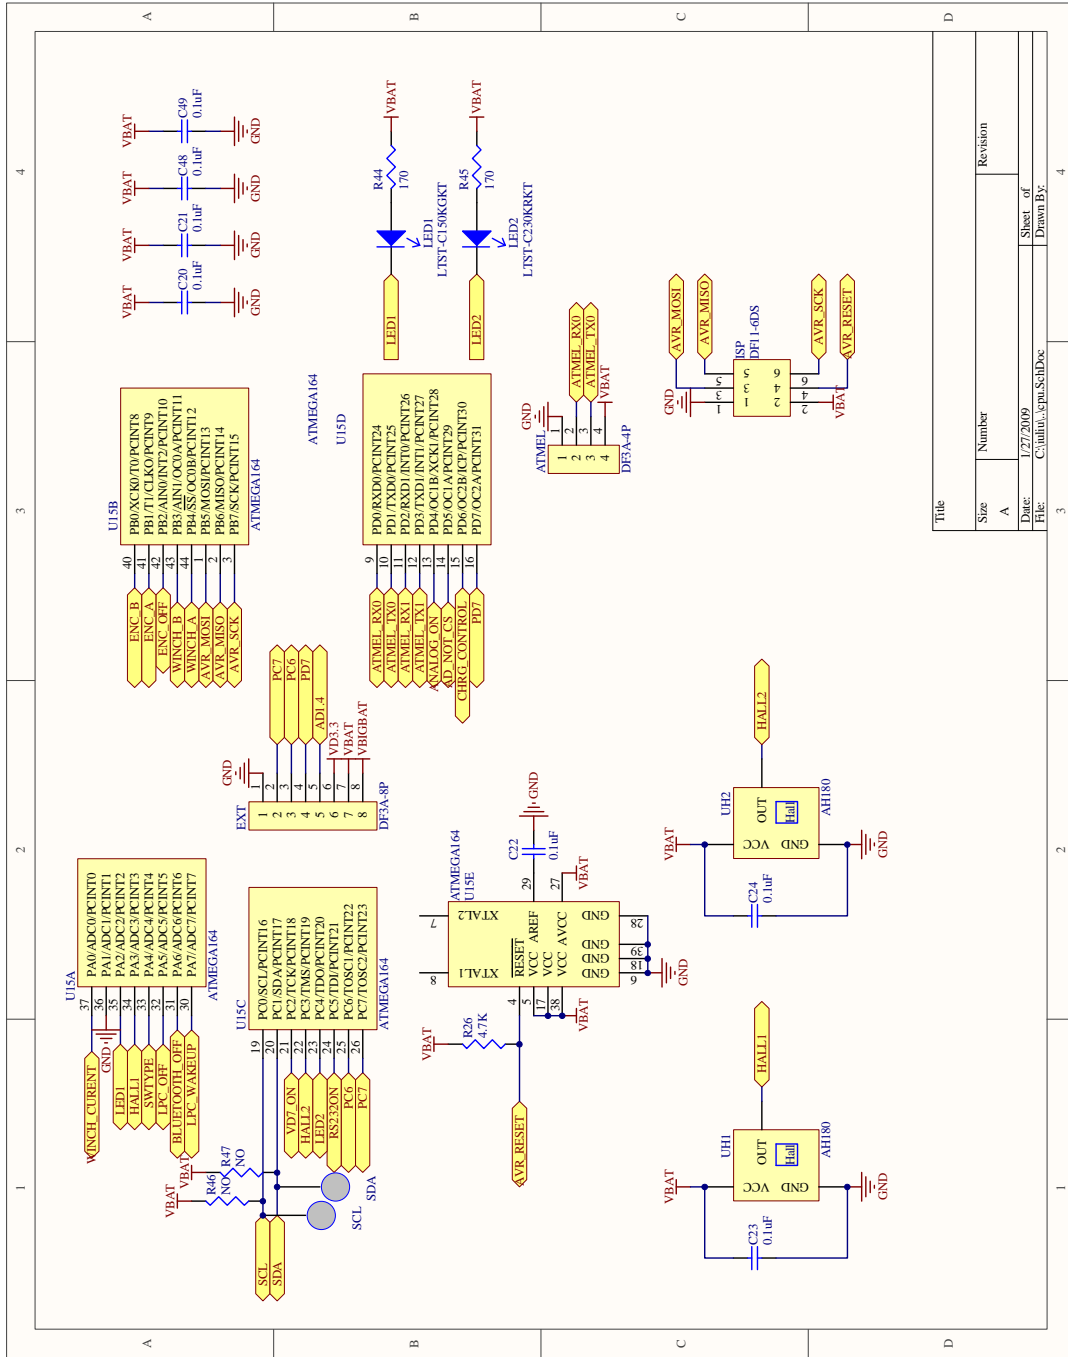


Figure E-5: AQUANODE Power Management CPU

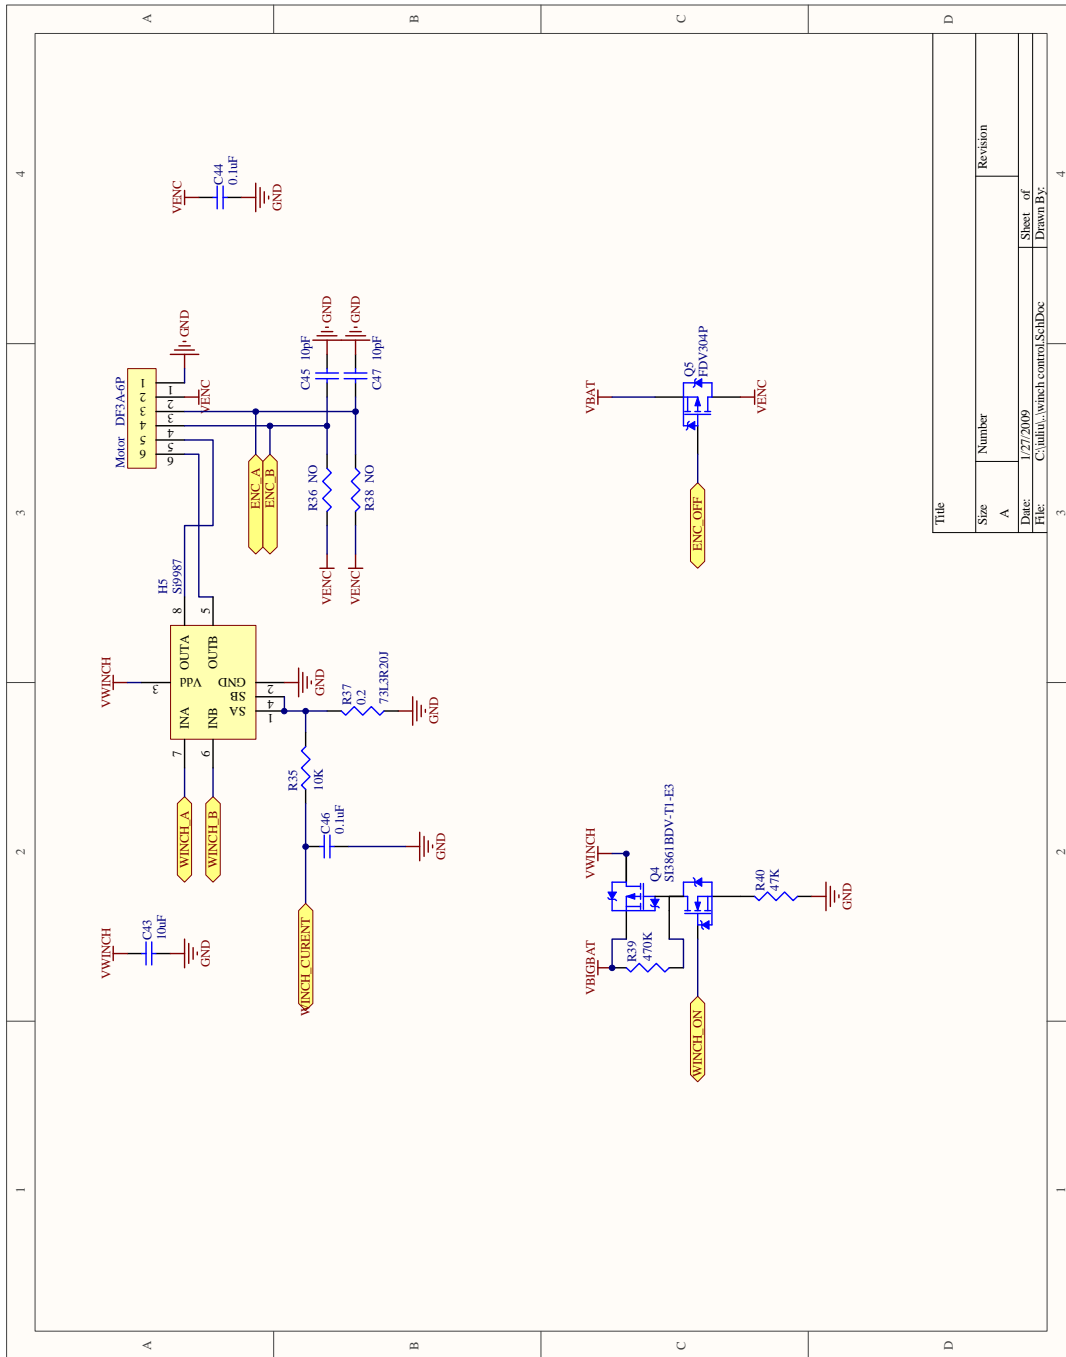


Figure E-6: AQUANODE Winch Control

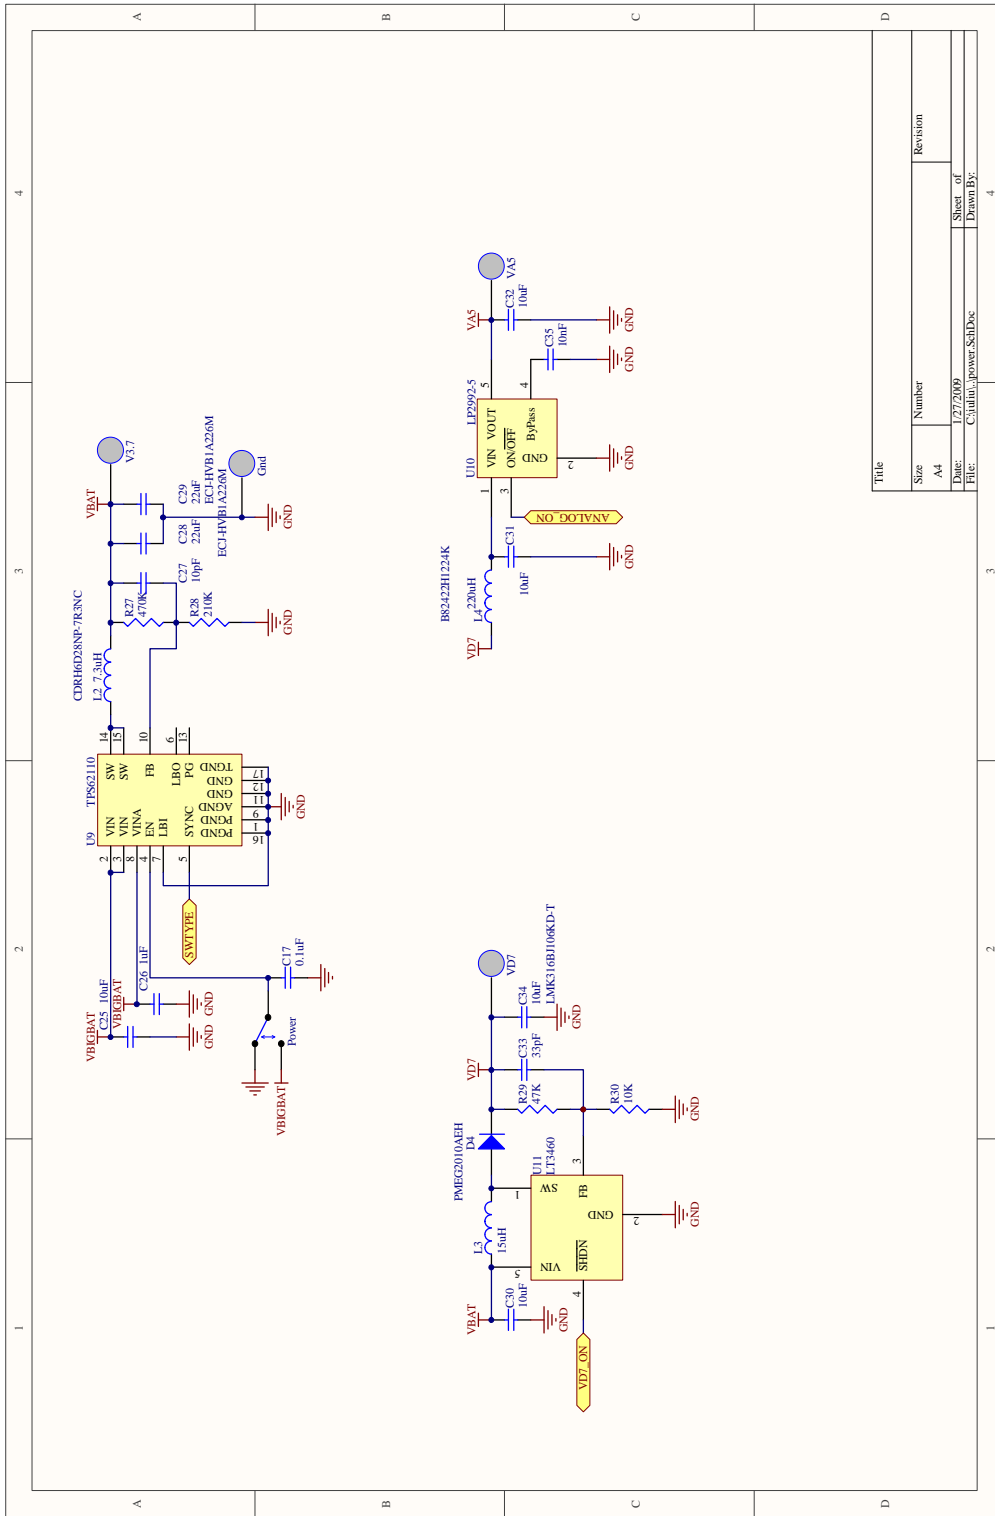


Figure E-7: AQUANODE Power Management

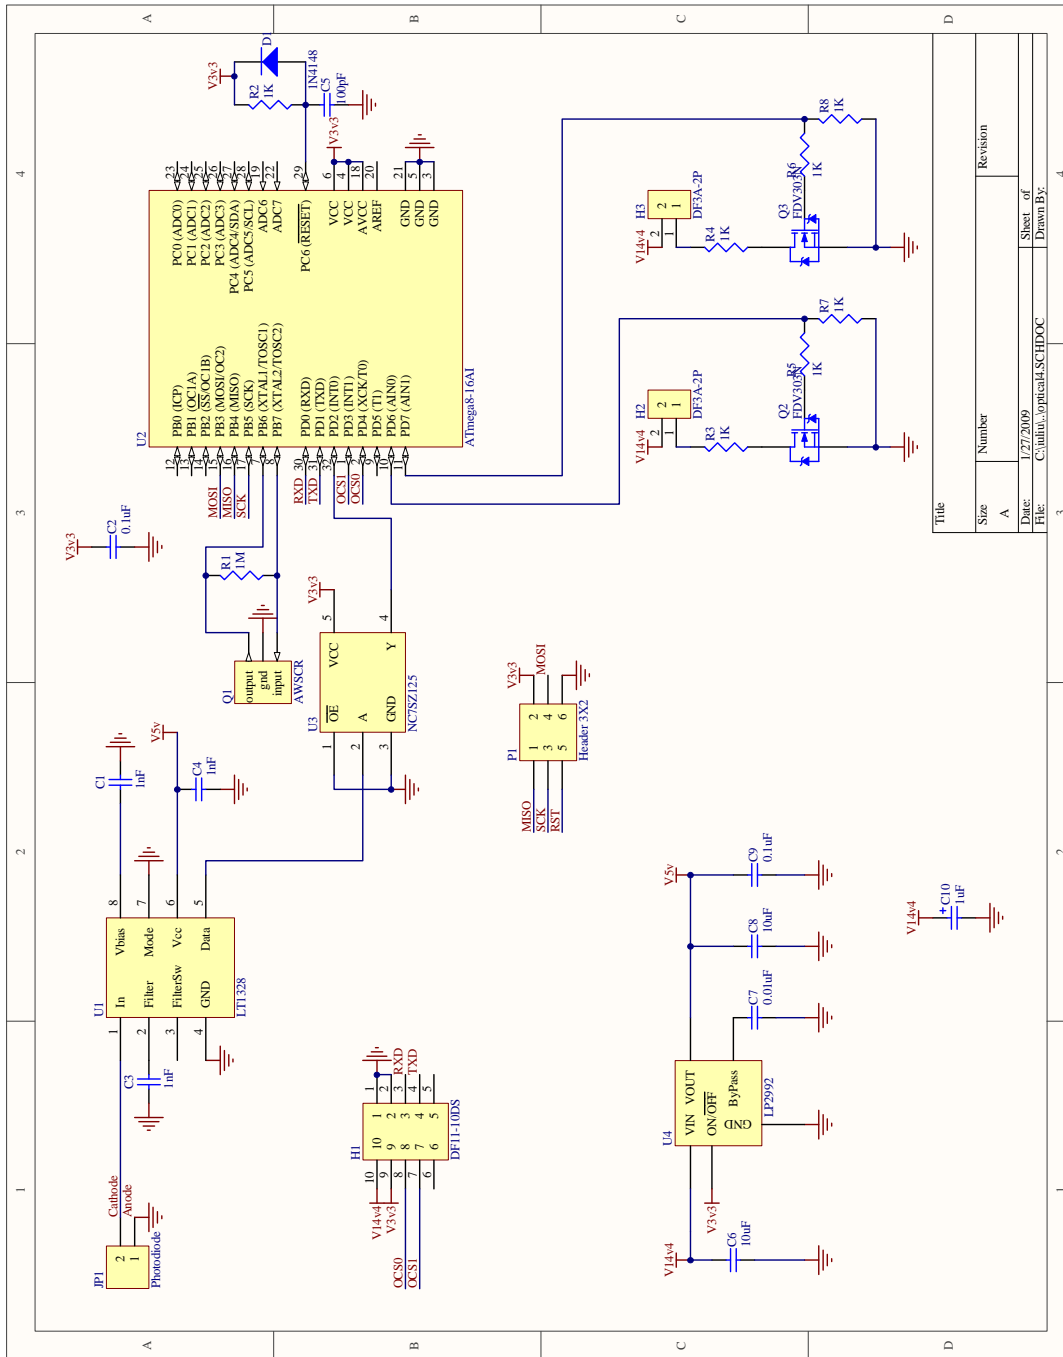


Figure E-8: Optical Modem Schematic

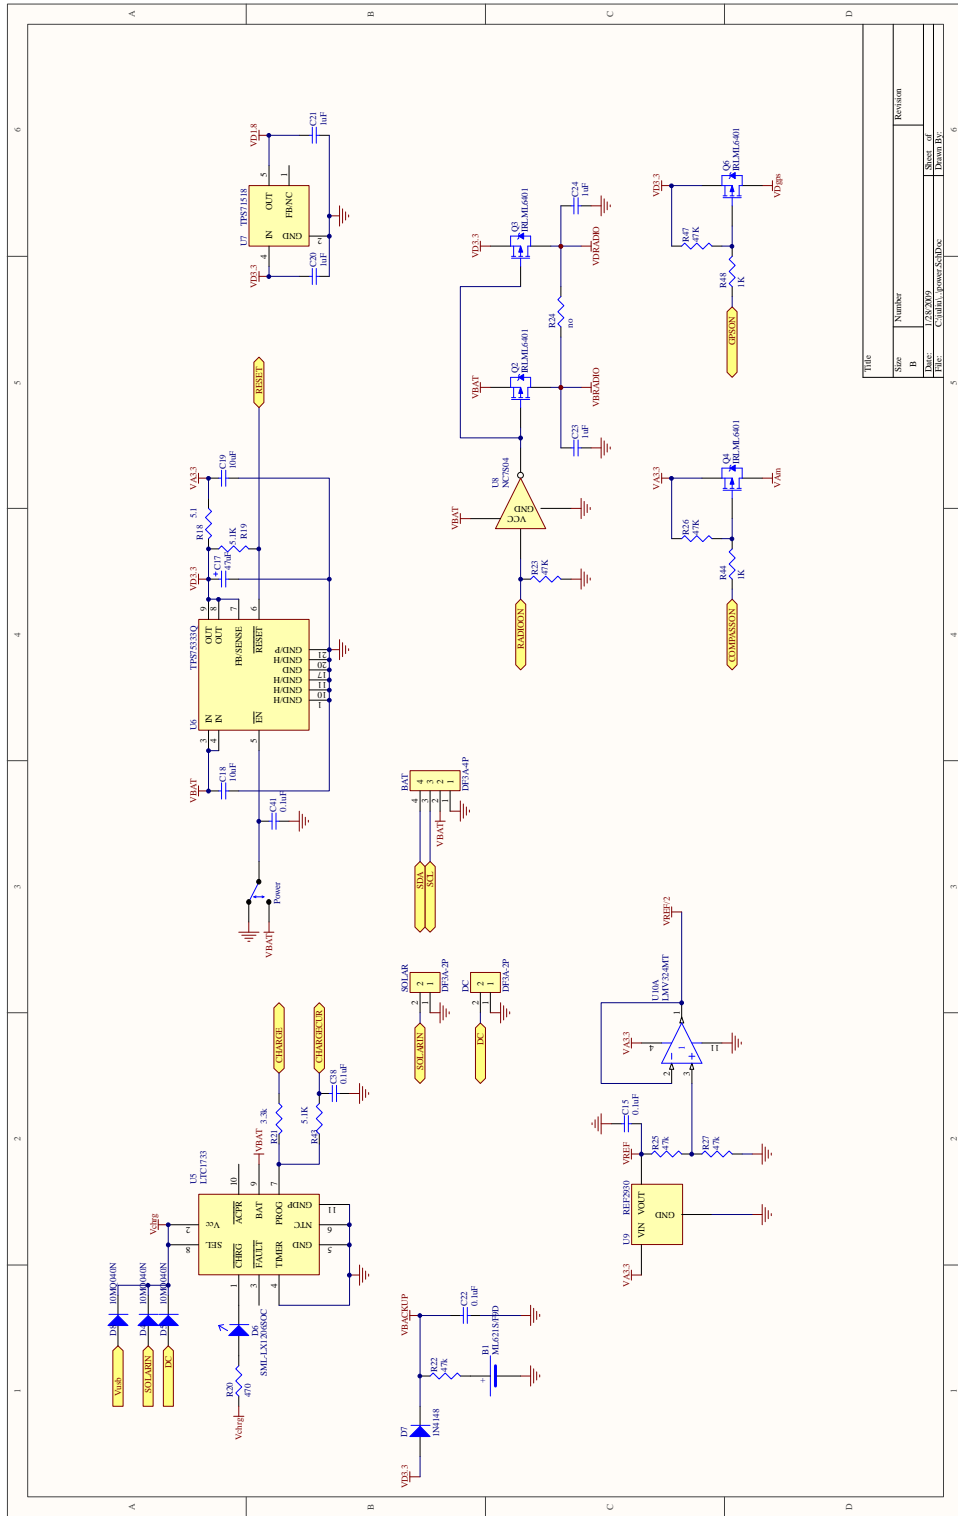


Figure E-9: AQUANODE CPU Power

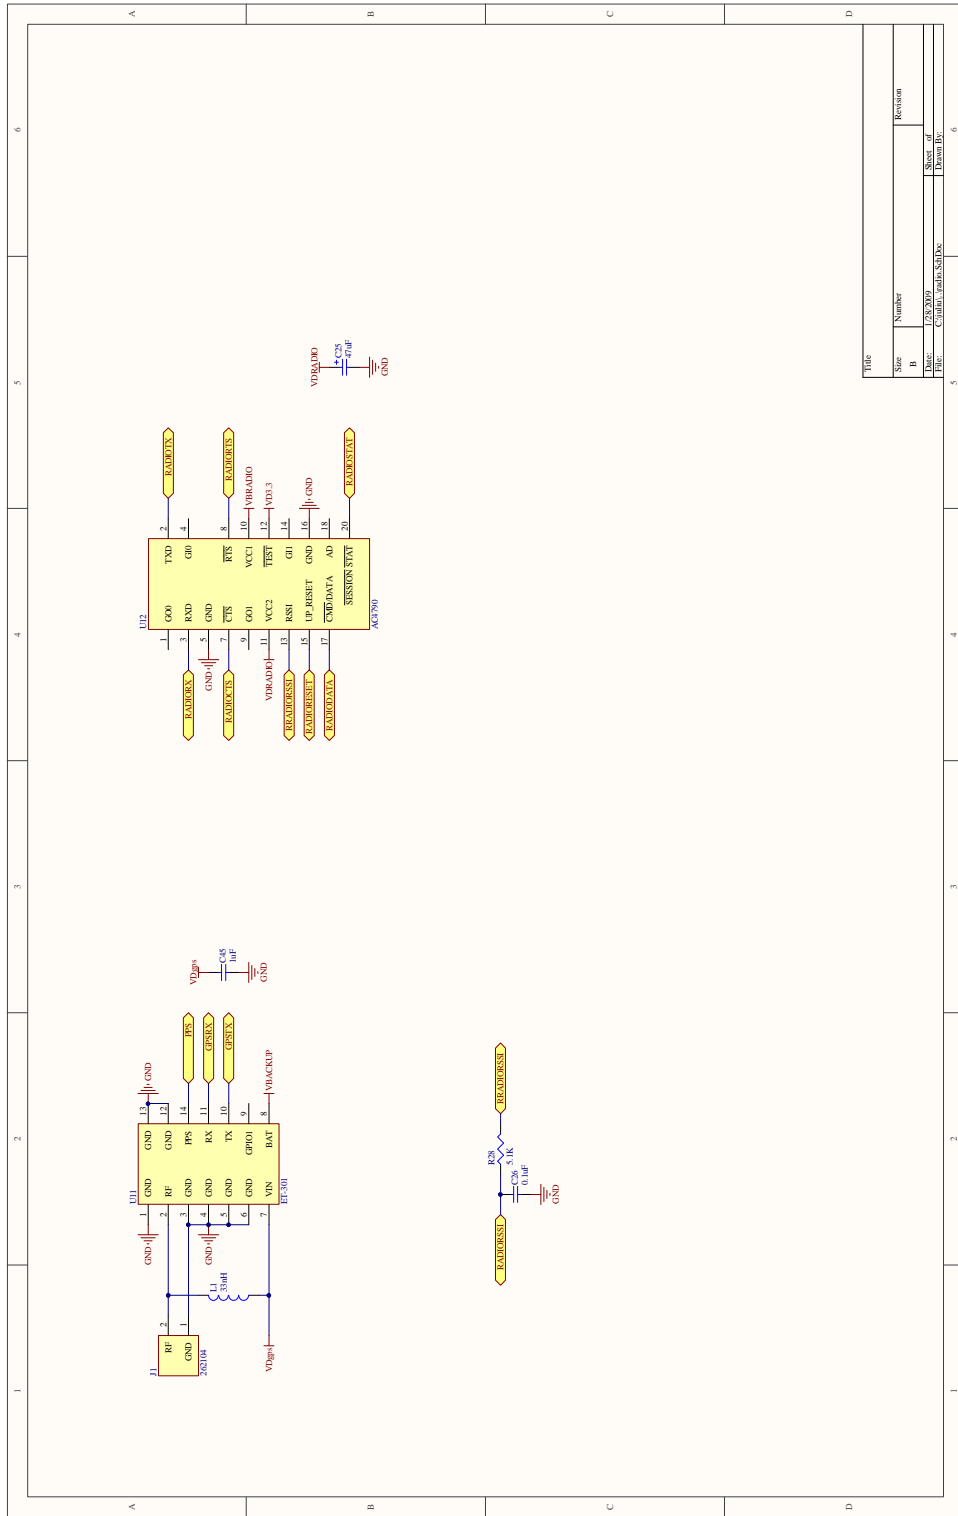


Figure E-10: AQUANODE Radio and GPS

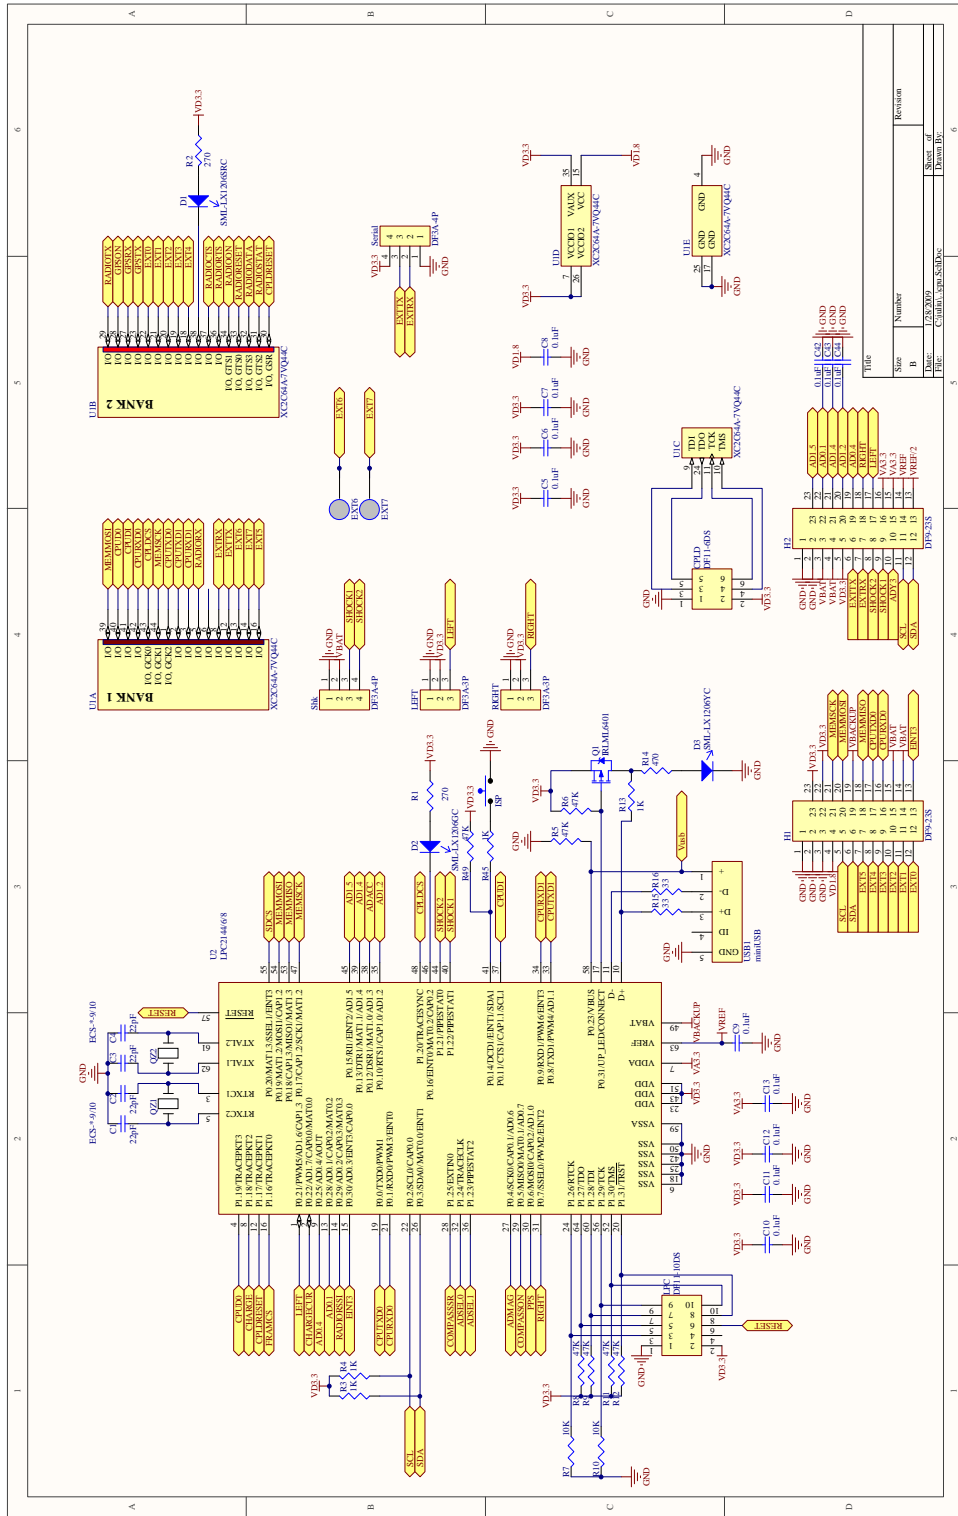
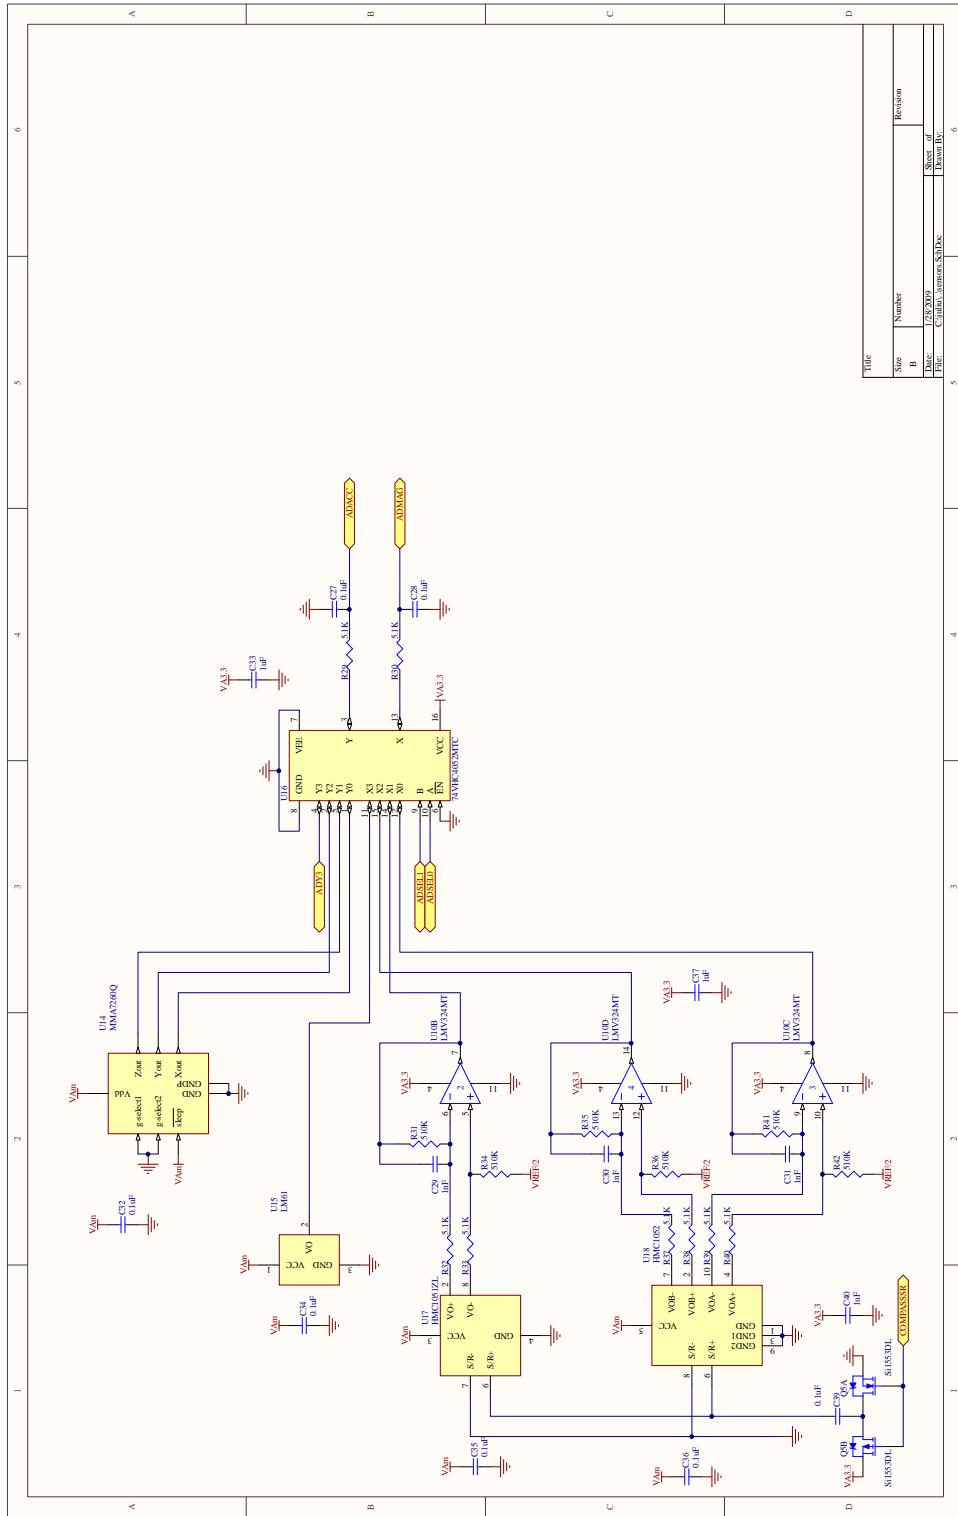


Figure E-11: AQUANODE Accelerometer and Compass



Title		Revision	
Size	Number	Sheet of	Revision
B			
Date:	1/25/2009	Sheet of	
File:	C:\path\aquanode\CPU.Dx	Drawn By:	

Figure E-12: AQUANODE CPU Flash Storage

Appendix F

TDMA Implementation

```
int slots = DEFAULT_SLOTS;
char slots_type[MAX_SLOTS_COUNT];
long long slot_offset = 0;

int local_id = DEFAULT_ID;

void init_slots(void) {
    int i;
    for(i=0; i<slots; i++) slots_type[i] = (i==local_id);
};

unsigned char send_buffer[MAX_PACKET_LENGTH] ;

packet received;

void tdma_dispatch(void) {
    start_time();
    receive_init();

    init_slots();

    while(1) {
```

```

while(get_time_in_slot >= SLOT_GUARD_TICKS);
int current_slot = get_slot;

if(slots_type[get_slot]){ //send slot
    packet to_send=*getMasterPacket();
    if(get_time_in_slot >= SLOT_GUARD_TICKS || get_slot != current_slot)
        continue;
    to_send.type |= PACKET_MASTER;
    to_send.slot = get_slot;
    prepare_raw_packet(&to_send, send_buffer);

    while(get_time_in_slot < SLOT_GUARD_TICKS);
    to_send.time=get_time();

    send_raw_packet(send_buffer, MAX_PACKET_LENGTH);

    if(to_send.type & PACKET_ALL_RANGES){
        int last_id = -1;
        while(get_time_in_slot < SLOT_TICKS - ALL_RANGES_SLOT_TICKS &&
            get_slot == current_slot){
            receive(&received, get_time()+SLOT_TICKS-get_time_in_slot, RECEIVE_SYNC);
            if(received.packetValid){
                int id = (received.time - to_send.time - ALL_RANGES_DELAY_TICKS) /
                    ALL_RANGES_SLOT_TICKS;
                if(id==last_id) break; // eliminate trivial multipath
                last_id = id;
                double distance =
                    (double)((received.time - to_send.time - ALL_RANGES_DELAY_TICKS)
                received.type = PACKET_ALL_RANGES;
                received.origin = id;
                received.destination = get_id();
                received.distance = distance;
                received.distanceValid = 1;
                int i;
                for(i=0;i<PACKET_DATA_LENGTH;i++) received.data[i]=0;

```

```

        packetReceived(&received);
    }
}
//Now that we got all of our ranges, do the tracking (500ms to do this)
doTracking();
}else{ // waiting response from one node
receive(&received, get_time()+SLOT_TICKS-get_time_in_slot, RECEIVE_PACKET);
if(received.packetValid){
if(!(received.type&PACKET_MASTER) &&
    received.destination == to_send.origin &&
    received.origin == to_send.destination){
received.distance=(float) (
    ((double) (received.time-to_send.time - RESPONSE_DELAY_TICKS))
    * SPEED_OF_SOUND / TIME_CLK / 2);
received.distanceValid = 1;
}else{
    received.distanceValid = 0;
}
    packetReceived(&received);
}
}
}else{ // receive slot
while((get_time_in_slot < SLOT_TICKS - TIME_CLK / 2)
    && get_slot == current_slot){
receive(&received, get_time()+SLOT_TICKS-get_time_in_slot, RECEIVE_PACKET);

if(received.packetValid){
    received.distanceValid = 0;
if(received.type&PACKET_MASTER)
    adjust_offset_to_guard(received.time, received.slot);
if(received.type & PACKET_ALL_RANGES) {
    while(get_time()<received.time
        + ALL_RANGES_DELAY_TICKS + get_id() * ALL_RANGES_SLOT_TICKS);
    send_sync();
    packetReceived(&received);
    break;
}
}
}
}
}

```

```

    }else if(received.destination == get_id() && (received.type&PACKET_MASTER)){
        packet *to_send=getResponsePacket(&received);
        to_send -> type &= ~PACKET_MASTER;
        prepare_raw_packet(to_send, send_buffer);
        while(get_time()<received.time + RESPONSE_DELAY_TICKS);
        send_raw_packet(send_buffer, MAX_PACKET_LENGTH);
    }else{
        packetReceived(&received);
    }
}
}
}
}
}

void adjust_offset_to_guard(long long time, int slot){
    long long slot_guard_time = (long long)slot * SLOT_TICKS + SLOT_GUARD_TICKS;
    slot_offset = slot_guard_time - time;
}

```

Appendix G

AMOUR Motor Controllers




Figure G-1: AMOUR Motor controller power board


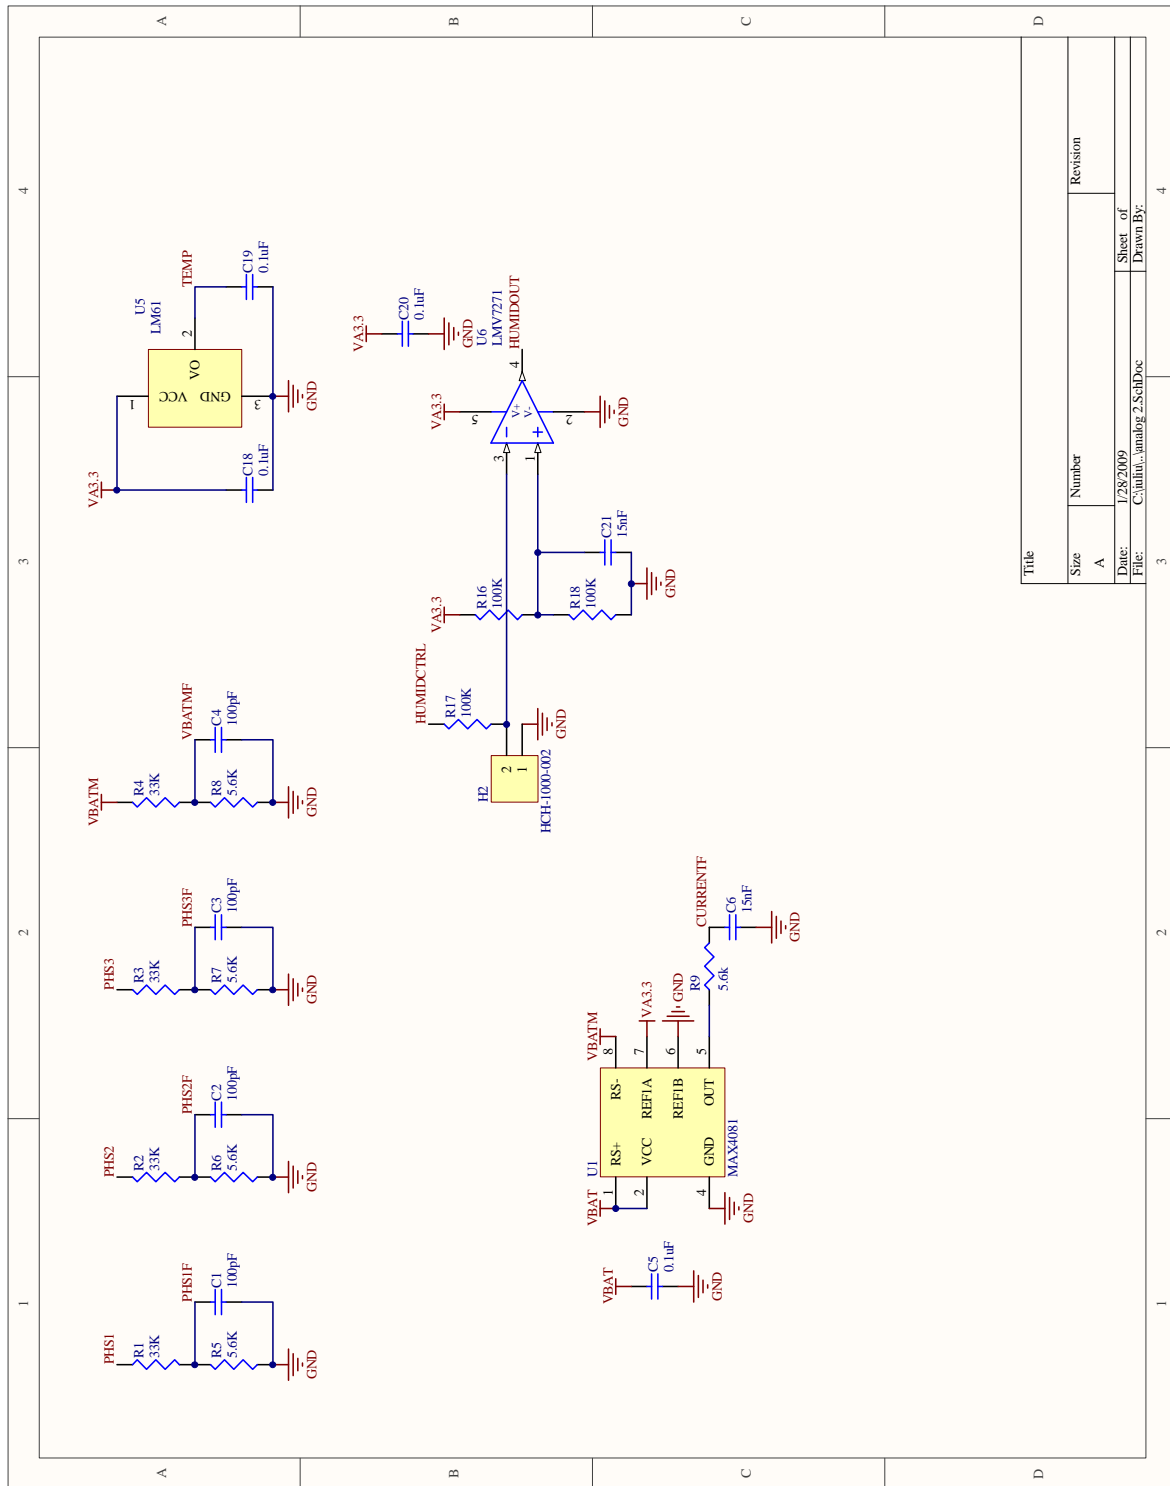


Figure G-2: AMOUR Motor controller regulator



Title	
Size	Number
A	
Date:	Revision
File:	
Sheet of	
Drawn By:	

Figure G-3: AMOUR Motor controller feedback

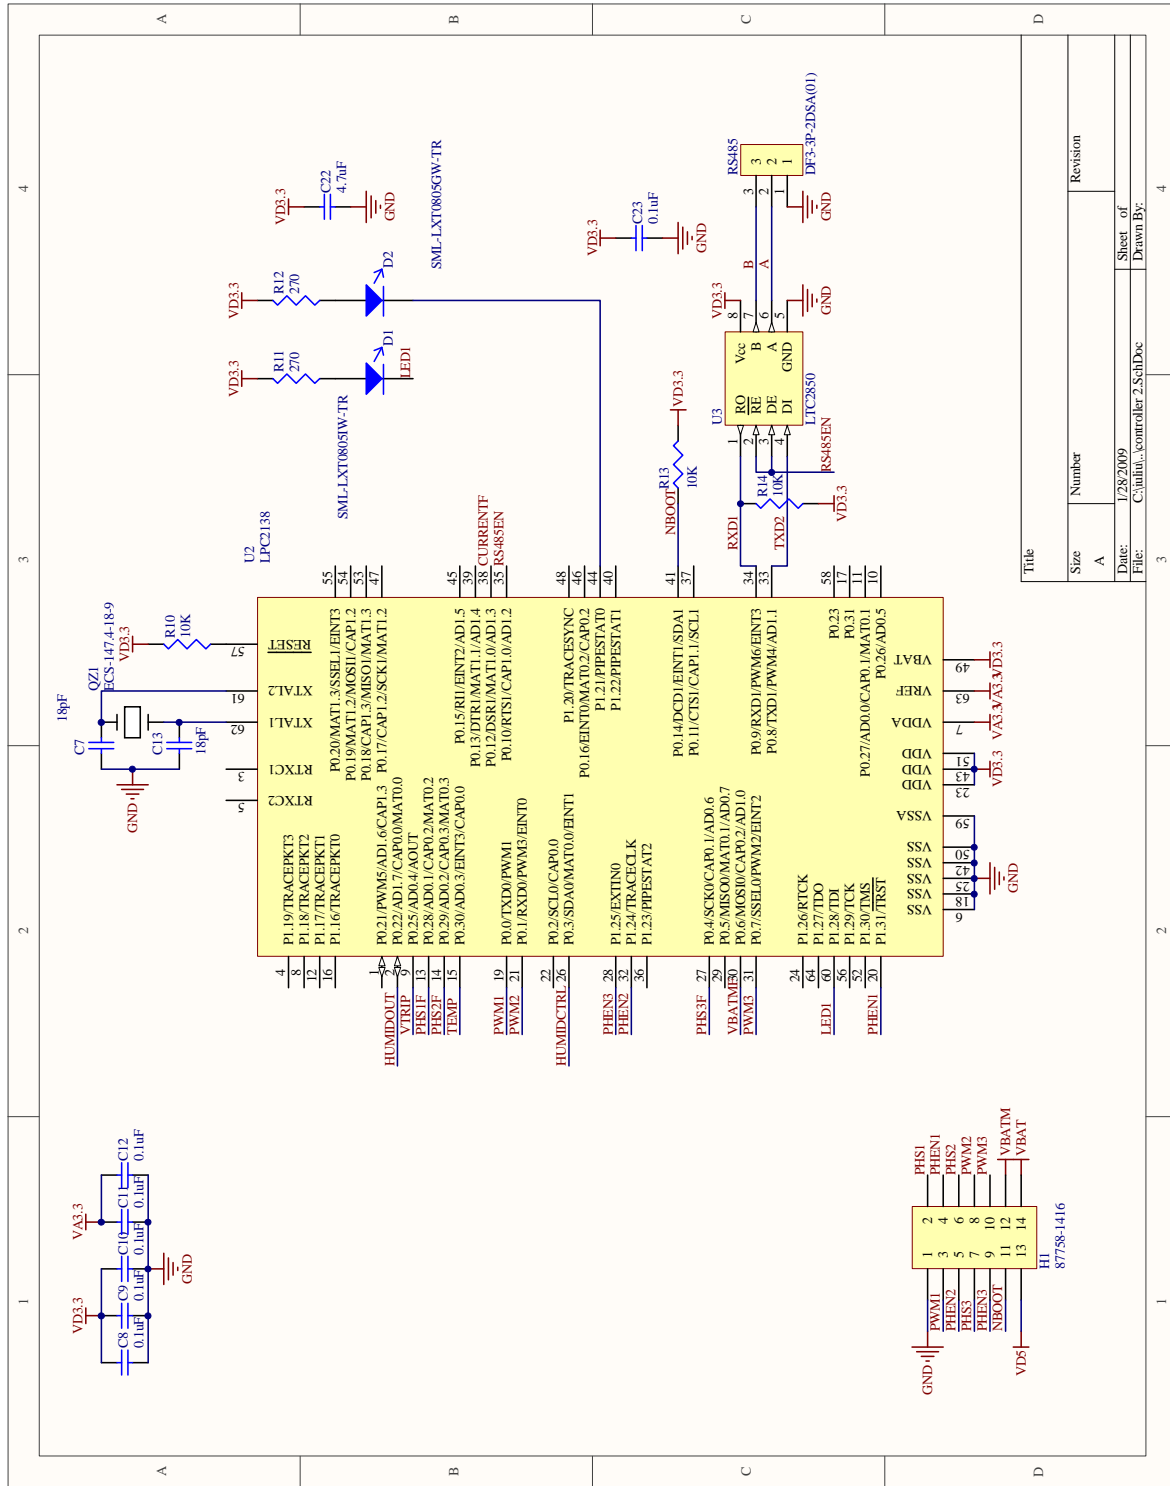


Figure G-4: AMOUR Motor controller CPU

Appendix H

AMOUR Docking Controller

```
#define TURN 24
#define MOVE 25
#define DOWN 15
#define MIN_DIFF 8

uint8_t SEE_LIGHT=30;
#define ON_TOP 220

#define ADJB(x) (x>0?x:x*3/2)

struct PID {
    int32_t comm, e1, e2, val;
    int32_t p, d, i, t;
    int32_t rez;
};

struct PID phot, dept, dire ;
```

```

void updatePID(struct PID *p, int32_t sensor){

    int32_t err=p->val-sensor;

    p->comm = p->p * err + ((err-p->e1)*p->d) + (p->i * (err+2*p->e1+p->e2))/2;
    //p->comm+= - (sensor - p->s1)*p->p + p->i * p->t * (p->val-sensor) - (p->d * (sensor
    if(p->comm>500) p->rez=500;
    if(p->comm<-500) p->rez=-500;
    p->rez = p->comm / 10;
    p->e2=p->e1; p->e1=err;
}

void initPID(struct PID *p, int pp, int i, int d, int t){
    p->comm=0;p->e1=0;p->e2=0;
    p->p=pp; p->i=i; p->d=d; p->t=t;
}

int8_t adjComm(int16_t comm){
    if(comm<0) comm=(comm*3)/2;
    if(comm<-50) comm=-50;
    if(comm>50) comm=50;
    return comm;
}

void initControllers(){
    initPID(&phot, 40, 0, 40, 10);
    initPID(&dept, 40, 0, 40, 10);
    initPID(&dire, 50, 0, 50, 10);
}

int32_t dl=422, depth=774, spee=0;

void goStraight2(int8_t comp, int32_t val){

```

```

    dire.val=val;
    updatePID(&dire, compass.data[comp]);
    BTLFT=adjComm(dire.rez+spee);
    BTRGT=adjComm(dire.rez-spee);
}

void keepDepth(int32_t depth){
    dept.val=depth;
    updatePID(&dept, getChannel(PRESSCHANNEL));
    TLFT=TRGT=adjComm(dept.rez);
}

void setPID(){
    struct PID *p;
    switch(si()){
        case 'p': p=&phot;break;
        case 'c': p=&dire;break;
        case 'd': p=&dept;break;
        default: return;
    }

    p->p=si();
    p->i=si();
    p->d=si();
}

void readPID(){
    struct PID *p;
    switch(si()){
        case 'p': p=&phot;break;
        case 'c': p=&dire;break;
        case 'd': p=&dept;break;
        default: return;
    }
}

```

```

    so(p->p);
    so(p->i);
    so(p->d);
}

void controlDock(void) {
    int32_t diffflr=((int32_t)LFT-(int32_t)RGT)*20/((int32_t) LFT+(int32_t)RGT);
    //int32_t diffffb=((int32_t)FWD-(int32_t)BCK)*20/((int32_t) FWD+(int32_t)BCK);
    phot.val=0;
    int8_t left, right;

    updatePID(&phot, diffflr);

    if(FWD>BCK){
        left  =-phot.rez;
        right =-phot.rez;
    }else{
        left  =phot.rez;
        right =phot.rez;
    }

    if(diffflr>-3 && diffflr<3){
        if(FWD-8>BCK){ left-=22; right+=22;}
        if(BCK-8>FWD){ left+=22; right-=22;}
    }

    BTLFT=adjComm(left);
    BTRGT=adjComm(right);

}

uint16_t dock_count=0, dock_state=0, optimal_depth;

```



```
uint8_t latch, latch_reset;
```

```
uint8_t min,max;
```

```
void computeMinMax(void) {  
    min=max=msg[1];int8_t i;  
    for(i=2;i<5;i++){  
        if(msg[i]>max)max=msg[i];  
        if(msg[i]<min)min=msg[i];  
    }  
}
```

```
void dock(void) {  
    computeMinMax();
```

```
    switch(dock_state){  
        case 0:  
            if(min<SEE_LIGHT && depth>optimal_depth) depth--;  
            if( (max-min)<max/8 || LIGHT) depth++;  
            controlDock();  
            if( min>ON_TOP){  
                BTLFT=BTRGT=0;  
                dock_state=1;  
                dock_count=0;  
            }  
            break;  
        case 1: // fast decent  
            if(!LIGHT && min<100) {dock_state=0; break;}  
            dock_count++;  
            if(dock_count==2) {depth++;dock_count=0;}  
            if(min<60) {dock_count=0; dock_state=2; depth+=4;latch=2;latch_reset=0; break;}  
            break;  
        case 2: // latch open  
            dock_count++;  
            if(dock_count>=8){ latch=4; dock_count=0; dock_state=3;latch=3;latch_reset=0;}  
            break;
```

```

    case 3: //latch close
        dock_count++;
        if(dock_count>=10){ latch=4; dock_count=0; dock_state=4; }
        break;
    case 4:
        if(!LIGHT) {dock_state=0; depth=optimal_depth; break;}
        dock_count++;
        if(depth>optimal_depth) depth--;
        if(dock_count>=20 && depth<=optimal_depth){ dock_state=10;}
        break;
}

}

void sendPhotoSerial(){
    so('a');so(LIGHT);so(msg[1]);so(msg[2]);so(msg[3]);so(msg[4]);

    so16(getChannel(PRESSCHANNEL));so16(depth);

    so16(compass.data[1]);so16(compass.data[2]);

    so(state);so(dock_state);

    so(13);so(10);
}

int main (void){
    int8_t i;

    initBoard();
    stopMotors();

```

```

initControllers();
delay(200);
latch=1;optimal_depth=770;

while(1){
  stopMotors();
  delay(1);

  switch(state){
    case 0: continue;
    case 1: // test just go straight
goStraight2( 1, d1);
break;
    case 2: // test just depth
    depth=optimal_depth;
keepDepth(depth);
break;

    case 3: // test go straight and depth
    depth=optimal_depth;
goStraight2( 1, d1);
keepDepth(depth);
break;
    case 4: // lock the dock
depth=optimal_depth;
keepDepth(depth);
controlDock();
break;

    case 5: //dock and land
depth=optimal_depth;
dock_state=0;
dock_count=0;

```

```

keepDepth(depth);
state=51;
break;
    case 51: //5-1
        dock();
keepDepth(depth);
if(dock_state==10) state=0;
    break;

    case 9:
        depth=optimal_depth;
goStraight2(1, d1);
keepDepth(depth);
computeMinMax();
if(max>SEE_LIGHT){
    state=91;
}
break;
    case 91:
        keepDepth(depth);
        if(max>SEE_LIGHT){
            state=92;
            depth=optimal_depth;
            dock_state=0;
            dock_count=0;
        }else
            state=9;
break;
    case 92:
        dock();
keepDepth(depth);
if(dock_state==10) state=0;
    break;
}

```

```
    delay(12);  
  }  
  stopMotors();  
  while(1);  
}
```

Bibliography

- [1] Color correction filters. <http://www.urprofilters.com/>.
- [2] X-rite colorchecker chart. http://www.xrite.com/product_overview.aspx?ID=820.
- [3] Ian F. Akyildiz, Dario Pompili, and Tommaso Melodia. Underwater acoustic sensor networks: research challenges. *Ad Hoc Networks (Elsevier)*, 3:257–279, 2005.
- [4] G.I. Allen, G. Sulzberger, J.T. Bono, J.S. Pray, and T.R. Clem. Initial evaluation of the new real-time tracking gradiometer designed for small unmanned underwater vehicles. In *OCEANS, 2005. Proceedings of MTS/IEEE*, pages 1956–1962 Vol. 3, 2005.
- [5] K. M. Aly and E. Esmail. Refractive index of salt water: effect of temperature. *Opt. Mater.*, 2:195–199, 1993.
- [6] Gianluca Antonelli, Thor I. Fossen, and Dana R. Yoerger. *Underwater Robotics*, pages 987–1008. 2008.
- [7] R. Austin and G. Halikas. The index of refraction of seawater. Technical report, Univ. Calif. SIO, La Jolla, California, 1976.
- [8] A. Baggeroer. Acoustic telemetry—an overview. *Oceanic Engineering, IEEE Journal of*, 9(4):229–235, Oct 1984.
- [9] K. Baker and R. C. Smith. Bio-optical classification and model of natural waters. *Limnol. Oceanogr.*, 27:500–509, 1982.

- [10] S. Bazeille, I. Quindu, L. Jaulin, and J.P. Malkasse. Automatic underwater image pre-processing. *Sea Tech Week*, 2006.
- [11] J.G. Bellingham, C.A. Goudey, T.R. Consi, J.W. Bales, D.K. Atwood, J.J. Leonard, and C. Chryssostomidis. A second generation survey auv. *Autonomous Underwater Vehicle Technology, 1994. AUV '94., Proceedings of the 1994 Symposium on*, pages 148–155, Jul 1994.
- [12] D. Bellwood. The 'lost' language of fishes. http://www.coralcoe.org.au/news_stories/fishcolour.html.
- [13] L. P. Boivan, W. F. Davidson, R. S. Storey, D. Sinclair, and E. D. Earle. Determination of the attenuation coefficients of visible and ultraviolet radiation in heavy water. *Appl. Opt.*, 25:877–882, 1986.
- [14] A. Bricaud, M. Babin, A. Morel, and H. Claustre. Variability in the chlorophyll-specific absorption coefficients of natural phytoplankton: analysis and parameterization. *Journal of Geophysical Research*, 100:13321–13332, 1995.
- [15] R. Brooks. A robust layered control system for a mobile robot. *Robotics and Automation, IEEE Journal of [legacy, pre - 1988]*, 2(1):14–23, Mar 1986.
- [16] Rodney A. Brooks. A robust layered control system for a mobile robot. Technical report, CSAIL, MIT, Cambridge, MA, USA, 1985.
- [17] Rodney A. Brooks. Integrated systems based on behaviors. *SIGART Bull.*, 2(4):46–50, 1991.
- [18] Rodney A. Brooks. *Cambrian intelligence: the early history of the new AI*. MIT Press, Cambridge, MA, USA, 1999.
- [19] H. Buiteveld, J. M. H. Hakvoort, and M. Donze. The optical properties of pure water. In J. S. Jaffe, editor, *SPIE Proceedings on Ocean Optics XII*, volume 2258, pages 174–183, 1994.

- [20] INC Campbell Scientific. Light absorption and scattering in water samples. http://www.campbellsci.ca/Download/LitNote_lightabsorption.pdf, 2008.
- [21] M. Chambah, D. Semani, A. Renouf, P. Courtellemont, and A. Rizzi. Underwater color constancy: Enhancement of automatic live fish recognition. In *Color Imaging Conference*, 2004.
- [22] MARK ALAN CHANCEY. Short range underwater optical communication links. Master’s thesis, North Carolina State University, 2005.
- [23] M. Chen, Z. He, and F. Ao. Study and implementation for range-gated underwater laser imaging system. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6625 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, March 2008.
- [24] H.T. Choi, A. Hanai, S.K. Choi, and J. Yuh. Development of an underwater robot, odin-iii. In *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 1, pages 836–841 vol.1, 2003.
- [25] Cathy Church. Underwater photography. http://www.cathychurch.com/OnlineGallery/Gallery1_New_DESIGN.html.
- [26] P. Corke, C. Detweiler, M. Dunbabin, M. Hamilton, D. Rus, and I. Vasilescu. Experiments with underwater robot localization and tracking. *Robotics and Automation, 2007 IEEE International Conference on*, pages 4556–4561, April 2007.
- [27] D.M. Crimmins, C.T. Patty, M.A. Beliard, J. Baker, J.C. Jalbert, R.J. Komer-ska, S.G. Chappell, and D.R. Blidberg. Long-endurance test results of the solar-powered auv system. In *OCEANS 2006*, pages 1–5, 2006.
- [28] A.C. Crook. Colour patterns in a coral reef fish: Is background complexity important? *J. of eperimental marine biology and ecology*, 217(2), 1997.

- [29] C. Detweiler, J. Leonard, D. Rus, and S. Teller. Passive mobile robot localization within a fixed beacon field. In *Proceedings of the 2006 International Workshop on Algorithmic Foundations of Robotics*, New York, aug 2006.
- [30] C. Detweiler, S. Sosnowski, I. Vasilescu, and D. Rus. Saving energy with buoyancy and balance control for underwater robots with dynamic payloads. In *Proceedings of the 11th International Symposium on Experimental Robotics*, Athens, Greece, July 2008.
- [31] C. Detweiler, I. Vasilescu, and D. Rus. An underwater sensor network with dual communications, sensing, and mobility. *OCEANS 2007 - Europe*, pages 1–6, June 2007.
- [32] N. E. Dorsey. *Properties of the Ordinary Water Substance in all Its Phases*. Reinhold, New York, 1940.
- [33] David Doubilet. Undersea images. <http://www.daviddoubilet.com/>.
- [34] Gregory Dudek, Philippe Giguere, Chris Prahacs, Shane Saunderson, Junaed Sattar, Luz-Abril Torres-Mendez, Michael Jenkin, Andrew German, Andrew Hogue, Arlene Ripsman, Jim Zacher, Evangelos Miliotis, Hui Liu, Pifu Zhang, Marti Buehler, and Christina Georgiades. Aqua: An amphibious autonomous robot. *Computer*, 40(1):46–53, 2007.
- [35] M. Dunbabin, P. Corke, I. Vasilescu, and D. Rus. Data muling over underwater wireless sensor networks using an autonomous underwater vehicle. *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 2091–2098, 15-19, 2006.
- [36] M. Dunbabin, J. Roberts, K. Usher, G. Winstanley, and P. Corke. A hybrid auv design for shallow water reef navigation. *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 2105–2110, April 2005.

- [37] M. Dunbabin, I. Vasilescu, P. Corke, and D. Rus. Cooperative navigation experiments with two autonomous underwater robots. In *ISER06 - International Symposium on Experimental Robotics*, 2006.
- [38] Martin Edge. The underwater photographer, third edition: Digital and traditional techniques (paperback), 2006.
- [39] C.C. Eriksen, T.J. Osse, R.D. Light, T. Wen, T.W. Lehman, P.L. Sabin, J.W. Ballard, and A.M. Chiodi. Seaglider: a long-range autonomous underwater vehicle for oceanographic research. *Oceanic Engineering, IEEE Journal of*, 26(4):424–436, Oct 2001.
- [40] D. Estrin, L. Girod, G. Pottie, and M. Srivastava. Instrumenting the world with wireless sensor networks. In *International Conference on Acoustics, Speech, and Signal Processing*, 2001.
- [41] R.M. Eustice. *Large-area visually augmented navigation for autonomous underwater vehicles*. PhD thesis, MIT, 2005.
- [42] Ryan M. Eustice, Hanumant Singh, John J. Leonard, and Matthew R. Walter. Visually mapping the rms titanic: Conservative covariance estimates for slam information filters. *Int. J. Rob. Res.*, 25(12):1223–1242, 2006.
- [43] E. S. Fry, G. W. Kattawar, and R. M. Pope. Integrating cavity absorption meter. *Appl. Opt.*, 31:2055–2065, 1992.
- [44] G. Bishop G. Welch. An introduction to the kalman filter, university of north carolina, 2006.
- [45] National Geographic. Milestones in underwater photography. <http://photography.nationalgeographic.com/photography/photos/milestones-underwater-photography.html>.
- [46] J.W. Giles and I.N. Bankman. Underwater optical communications systems. part 2: basic design considerations. *Military Communications Conference, 2005. MILCOM 2005. IEEE*, pages 1700–1705 Vol. 3, Oct. 2005.

- [47] Kyle Gilpin, Keith Kotay, Daniela Rus, and Iuliu Vasilescu. Miche: Modular shape formation by self-disassembly. *Int. J. Rob. Res.*, 27(3-4):345–372, 2008.
- [48] L. Giosan, R. D. Flood, J. Grtzner, and P. Mudie. Paleoceanographic significance of sediment color on western North Atlantic drifts: II. Plio-Pleistocene sedimentation. *International Journal of Marine Geology, Geochemistry and Geophysics*, 189, 2002.
- [49] G. M. Hale and M. R. Querry. Optical constants of water in the 200 nm to 200 μm wavelength region. *Appl. Opt.*, 12:555–563, 1973.
- [50] Frank Hanson and Stojan Radic. High bandwidth underwater optical communication. *Appl. Opt.*, 47(2):277–283, 2008.
- [51] Julia Åhlén. *Colour Correction of Underwater Images Using Spectral Data*. PhD thesis, Uppsala Universitet, 2005.
- [52] Commission internationale de l’éclairage. Cie free publications. <http://www.cie.co.at/main/freepubs.html>.
- [53] Commission internationale de l’éclairage. Cie 1931 2° standard observer, 1931.
- [54] W. M. Irvine and J. B. Pollack. Infrared optical properties of water and ice spheres. *Icarus*, 8:324–360, 1968.
- [55] Jules Jaffe, Kad Moore, John McLean, and Michael Strand. Underwater optical imaging: Status and prospects. *Oceanography*, 14, 2007.
- [56] Azam Osman Kashif Iqbal, Rosalina Abdul Salam and Abdullah Zawawi Talib. Underwater image enhancement using an integrated colour model. *IAENG International Journal of Computer Science*, 2007.
- [57] Kebes. Water absorption. http://en.wikipedia.org/wiki/Water_absorption.
- [58] D.B. Kilfoyle and A.B. Baggeroer. The state of the art in underwater acoustic telemetry. *Oceanic Engineering, IEEE Journal of*, 25(1):4–27, Jan 2000.

- [59] O. V. Kopelevich. Optical properties of pure water in the 250–600 nm range. *Opt. Spectrosc.*, 41:391–392, 1976.
- [60] L. Kou. *Refractive indices of water and ice in the 0.65–2.5 μm spectral range*. PhD thesis, Dalhousie University (Canada), 1993.
- [61] L. Kou, D. Labrie, and P. Chylek. Refractive indices of water and ice in the 0.65–2.5 μm spectral range. *Appl. Opt.*, 32:3531–3540, 1993.
- [62] A.P.Kerswell K.R.N. Anthony. Coral mortality following extreme low tides and high solar radiation. *Marine Biology*, 151(4), 5 2007.
- [63] Clayton Kunz, Chris Murphy, Richard Camilli, Hanumant Singh, John Bailey, Ryan Eustice, Michael Jakuba, Ko ichi Nakamura, Chris Roman, Taichi Sato, Robert A. Sohn, and Claire Willis. 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*, pages 3654–3660, 2008.
- [64] J. L. Laine, S. A. Nichols, D. K. Novick, P. D. O Malley, D. Copeland, and M. C. Nechyba. Subjugator: a highly maneuverable, intelligent underwater vehicle. Technical report, Machine Intelligence Laboratory, University of Florida, Gainesville, FL, 1999.
- [65] E. H. Land and J. J. McCann. Lightness and Retinex Theory. *Journal of the Optical Society of America (1917-1983)*, 61:1–+, January 1971.
- [66] S. Licht, V. Polidoro, M. Flores, F.S. Hover, and M.S. Triantafyllou. Design and projected performance of a flapping foil auv. *Oceanic Engineering, IEEE Journal of*, 29(3):786–794, 2004.
- [67] Xavier Lurton. *An Introduction to Underwater Acoustics*. Springer, 2002.
- [68] S.C. Martin, L.L. Whitcomb, D. Yoerger, and H. Singh. A mission controller for high level control of autonomous and semi-autonomous underwater vehicles. In *OCEANS 2006*, pages 1–6, 2006.

- [69] D. McFarland, I. Gilhespy, and Honary E. Divebot: A diving robot with a whale-like buoyancy mechanism. *Robotica (2003)*, 21(04), aug 2003.
- [70] E. A. McLean, H. R. Burris Jr., and M. P. Strand. Short-pulse range-gated optical imaging in turbid water. *Applied Optics*, 1995.
- [71] V. Messmer, G.P. Jones, L. van Herwerden, and P.L. Munday. Genetic and ecological characterisation of colour dimorphism in a coral reef fish. *Environmental Biology of Fishes*, 74, 2005.
- [72] V. Messmer, L. van Herwerden, P.L. Munday, and G.P. Jones. Phylogeography of colour polymorphism in the coral reef fish pseudochromis fuscus, from papua new guinea and the great barrier reef. *Coral Reefs*, 24(3), 2005.
- [73] A. Morel and L. Prieur. Analysis of variations in ocean color. *Limnol. Oceanogr.*, 22:709–722, 1977.
- [74] S.G. Narasimhan and S.K. Nayar. Structured light methods for underwater imaging: Light stripe scanning and photometric stereo. In *OCEANS*, 2005.
- [75] S.G. Narasimhan, S.K. Nayar, B. Sun, and S.J. Koppal. Structured light in scattering media. In *OCEANS*, 2005.
- [76] Shahriar Negahdaripour and Pezhman Firoozfam. An rov stereovision system for ship hull inspection. In *OCEANS*, 2005.
- [77] D. K. Novick, R. Pitzer, B. Wilkers, C. D. Crane, E. de la Iglesia, and K. L. Doty. The development of a highly maneuverable underwater vehicle. In *Robotics 98: The 3rd International Conference and Exposition/Demonstration on Robotics for Challenging Environments*, pages 168–173, Albuquerque, NM, apr 1998.
- [78] K. F. Palmer and D. Williams. Optical properties of water in the near infrared. *J. Opt. Soc. Am.*, 64:1107–1110, 1974.
- [79] G. Pang, T. Kwan, H. Liu, and Chi-Ho Chan. Optical wireless based on high brightness visible leds. *Industry Applications Conference, 1999. Thirty-Fourth*

- IAS Annual Meeting. Conference Record of the 1999 IEEE*, 3:1693–1699 vol.3, 1999.
- [80] Jim Partan, Jim Kurose, and Brian Neil Levine. A survey of practical issues in underwater networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 11(4):23–33, 2007.
- [81] D. Perovich and J. Govoni. Absorption coefficients of ice from 250 to 400 nm. *Geophys. Res. Lett.*, 18:1233–1235, 1991.
- [82] Dario Pompili. Efficient communication protocols for underwater acoustic sensor networks, 2007.
- [83] R. M. Pope. *Optical absorption of pure water and sea water using the integrating cavity absorption meter*. PhD thesis, Texas A&M University, 1993.
- [84] R. M. Pope and E. S. Fry. Absorption spectrum (380–700 nm) of pure water. II. integrating cavity measurements. *Appl. Opt.*, 36:8710–8723, 1997.
- [85] G. J. Pottie and W. J. Kaiser. Wireless integrated network sensors. *Commun. ACM*, 43(5):51–58, 2000.
- [86] G.J. Pottie. Wireless sensor networks. *Information Theory Workshop, 1998*, pages 139–140, Jun 1998.
- [87] James Preisig. Acoustic propagation considerations for underwater acoustic communications network development. In *WUWNet '06: Proceedings of the 1st ACM international workshop on Underwater networks*, pages 1–5, New York, NY, USA, 2006. ACM.
- [88] L. Prieur and S. Sathyendranath. An optical classification of coastal and oceanic waters based on the specific spectral absorption curves of phytoplankton pigments dissolved organic matter, and other particulate materials. *Limnol. Oceanogr.*, 26:671–689, 1981.

- [89] J.G. Proakis, E.M. Sozer, J.A. Rice, and M. Stojanovic. Shallow water acoustic networks. *Communications Magazine, IEEE*, 39(11):114–119, Nov 2001.
- [90] Xiaohong Quan and E. S. Fry. Empirical equation for the index of refraction of seawater. *Appl. Opt.*, 34:3477–3480, 1995.
- [91] M. R. Querry, P. G. Cary, and R. C. Waring. Split-pulse laser method for measuring attenuation coefficients of transparent liquids: application to deionized filtered water in the visible region. *Appl. Opt.*, 17:3587–3592, 1978.
- [92] M. R. Querry, D. M. Wieliczka, and D. J. Segelstein. Water (H₂O). In *Handbook of Optical Constants of Solids II*, pages 1059–1077. Academic Press, 1991.
- [93] T. I. Quickenden and J. A. Irvin. The ultraviolet absorption spectrum of liquid water. *J. Chem Phys.*, 72:4416–4428, 1980.
- [94] S. Reed, A. Cormack, K. Hamilton, I. Tena Ruiz, and D. Lane. Automatic ship hull inspection using unmanned underwater vehicles. http://www.seebyte.com/docs/ShipHull_Arredondo.pdf.
- [95] F. Salva-Garau and M. Stojanovic. Multi-cluster protocol for ad hoc mobile underwater acoustic networks. *OCEANS 2003. Proceedings*, 1:91–98 Vol.1, 2003.
- [96] Mehul Sangekar, Mandar Chitre, and Koay Teong Beng. Hardware architecture for a modular autonomous underwater vehicle starfish. In *OCEANS*, Quebec City, Canada, 2008.
- [97] Shubha Sathyendranath and Trevor Platt. Analytic model of ocean color. *Applied Optics*, 36(12):2620–2629, 1997.
- [98] P. Schiebener, J. Straub, J. M. H. L. Sengers, and J. S. Gallagher. Refractive index of water and steam as function of wavelength, temperature and density. *J. Phys. Ch. R.*, 19:677–717, 1990.

- [99] Felix Schill, Uwe R. Zimmer, and Jochen Trumpf. Visible spectrum optical communication and distance sensing for underwater applications. In *AGRA*, 2004.
- [100] Mac Schwager, Carrick Detweiler, Iuliu Vasilescu, Dean M. Anderson, and Daniela Rus. Data-driven identification of group dynamics for motion prediction and control. *J. Field Robot.*, 25(6-7):305–324, 2008.
- [101] D. J. Segelstein. *The complex refractive index of water*. PhD thesis, University of Missouri-Kansas City, 1981.
- [102] J. Sherman, R.E. Davis, W.B. Owens, and J. Valdes. The autonomous underwater glider "spray". *Oceanic Engineering, IEEE Journal of*, 26(4):437–446, Oct 2001.
- [103] K. S. Shifrin. *Physical Optics of Ocean Water*. American Institute of Physics, New York, 1988.
- [104] H. Singh, J. Howland, and O. Pizarro. Advances in large-area photomosaicking underwater. *Oceanic Engineering, IEEE Journal of*, 29(3):872–886, July 2004.
- [105] Hanumant Singh, Roy Armstrong, Fernando Gilbes, Ryan Eustice, Chris Roman, Oscar Pizarro, and Juan Torres. Imaging coral i: Imaging coral habitats with the seabed auv. *Subsurface Sensing Technologies and Applications*, 5, 2004.
- [106] R. C. Smith and K. S. Baker. Optical properties of the clearest natural waters (200–800 nm). *Appl. Opt.*, 20:177–184, 1981.
- [107] F. M. Sogandares. *The spectral absorption of pure water*. PhD thesis, Texas A&M University, August 1991.
- [108] F. M. Sogandares and E. S. Fry. Absorption spectrum (340–640 nm) of pure water. I. Photothermal measurements. *Appl. Opt.*, 36:8699–8709, 1997.
- [109] E.M. Sozer, M. Stojanovic, and J.G. Proakis. Underwater acoustic networks. *Oceanic Engineering, IEEE Journal of*, 25(1):72–83, Jan 2000.

- [110] A. Stockman, D.I.A. MacLeod, and N.E. Johnson. Spectral sensitivities of human cones. *Journal of the Optical Society of America*, 10:2491–2521, 1993.
- [111] M. Stojanovic, L. Freitag, J. Leonard, and P. Newman. A network protocol for multiple auv localization. *Oceans '02 MTS/IEEE*, 1:604–611 vol.1, Oct. 2002.
- [112] R.P. Stokey, A. Roup, C. von Alt, B. Allen, N. Forrester, T. Austin, R. Goldsborough, M. Purcell, F. Jaffre, G. Packard, and A. Kukulya. Development of the remus 600 autonomous underwater vehicle. In *OCEANS, 2005. Proceedings of MTS/IEEE*, pages 1301–1304 Vol. 2, 2005.
- [113] S. A. Sullivan. Experimental study of the absorption in distilled water, artificial sea water, and heavy water in the visible region of the spectrum. *Opt. Soc. Am. J.*, 53:962–968, 1963.
- [114] A. C. Tam and C. K. N. Patel. Optical absorption of light and heavy water by laser optoacoustic spectroscopy. *Appl. Opt.*, 18:3348–3358, 1979.
- [115] L.A. Torres-Mndez and G. Dudek. A statistical learning-based method for color correction of underwater images. *Research on Computer Science*, 17, 2005.
- [116] Y. Tsuchida, Hama N., and Takahata M. An optical telemetry system for underwater recording of electromyogram and neuronal activity from non-tethered crayfish. *J of Neuroscience Methods*, 2004.
- [117] M. Tubaishat and S. Madria. Sensor networks: an overview. *Potentials, IEEE*, 22(2):20–23, April-May 2003.
- [118] J. Vaganay, M. Elkins, D. Esposito, W. OHalloran, F. Hover, and M. Kokko. Ship hull inspection with the hauv: Us navy and nato demonstration results. In *OCEANS*, 2006.
- [119] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke. Data collection, storage, and retrieval with an underwater sensor network. In *SenSys '05*:

Proceedings of the 3rd international conference on Embedded networked sensor systems, pages 154–165, New York, NY, USA, 2005. ACM.

- [120] I. Vasilescu, K. Kotay, D. Rus, L. Overs, P. Sikka, M. Dunbabin, P. Chen, and P. Corke. Krill: An exploration in underwater sensor networks. *Embedded Networked Sensors, 2005. EmNetS-II. The Second IEEE Workshop on*, pages 151–152, May 2005.
- [121] I. Vasilescu, P. Varshavskaya, K. Kotay, and D. Rus. Autonomous modular optical underwater robot (amour) design, prototype and feasibility study. *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 1603–1609, April 2005.
- [122] Iuliu Vasilescu, Carrick Detweiler, and Daniela Rus. Aquanodes: an underwater sensor network. In *WuWNet '07: Proceedings of the second workshop on Underwater networks*, pages 85–88, New York, NY, USA, 2007. ACM.
- [123] S. G. Warren. Optical constants of ice from the ultraviolet to the microwave. *Appl. Opt.*, 23:1026–1225, 1984.
- [124] D. M. Wieliczka, Shengshan Weng, and M. R. Querry. Wedge shaped cell for highly absorbent liquids: infrared optical constants of water. *Appl. Opt.*, 28:1714–1719, 1989.
- [125] G. Wyszecki. Proposal for a new color-difference formula. *JOSA*, 53:1318–1319, 1963.
- [126] A. Yamashita, M. Fujii, and T. Kaneko. Color registration of underwater images for underwater sensing with consideration of light attenuation. *Robotics and Automation, 2007 IEEE International Conference on*, pages 4570–4575, April 2007.
- [127] D. R. Yoerger, A. M. Bradley, S. C. Martin, and L. L. Whitcomb. The sentry autonomous underwater vehicle: Field trial results and future capabilities. *AGU Fall Meeting Abstracts*, 33:1674, December 2006.

- [128] Dana R. Yoerger, Albert M. Bradley, Barrie B. Walden, Hanumant Singh, and Ralf Bachmayer. Surveying a subsea lava flow using the autonomous benthic explorer (abe). *International Journal of Systems Science*, 29:1031–1044, 1998.
- [129] J. Yuh and M. West. Underwater robotics. *International Journal of Advanced Robotics*, 15(5):609–639, 2001.
- [130] H. R. Zelsmann. Temperature dependence of the optical constants for liquid H₂O and D₂O in the far IR region. *J. Mol. Struct.*, 350:95–114, 1995.
- [131] V. M. Zolotarev, B. A. Mikhilov, L. L. Alperovich, and S. I. Popov. Dispersion and absorption of liquid water in the infrared and radio regions of the spectrum. *Optics and Spectroscopy*, 27:430–432, 1969.