The Dissertation Committee for Karl Frazier Nieman
certifies that this is the approved version of the following dissertation:

Space-Time-Frequency Methods for Interference-Limited Communication Systems

Committee:

Brian L. Evans, Supervisor

Ross Baldick

Robert W. Heath, Jr.

Russell Pinkston

Preston S. Wilson

Space-Time-Frequency Methods for Interference-Limited Communication Systems

by

Karl Frazier Nieman, B.S., M.S.E.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

December 2014

To my parents.

Acknowledgments

Foremost, I would like to express sincere gratitude to my adviser, Prof. Brian L. Evans for his continual support and commitment to my study and PhD research. What began as a casual conversation in the spring of 2009 has transformed into a challenging, rewarding, and mentally illuminating endeavour that has shaped many aspects of my life. I am indebted to him for the patience and guidance he has given me throughout my journey. In addition, I would like to thank the rest of my dissertation committee: Prof. Ross Baldick, Prof. Robert W. Heath, Jr., Prof. Russell Pinkston, and Prof. Preston S. Wilson. Their helpful suggestions and attention to detail have helped me improve the quality of my work and the clarity of its presentation.

I owe many thanks to the great mentors I have had along my path to the PhD. Namely, I would like to thank Dr. Ian C. Wong, Dr. Nikhil U. Kundargi, and Dr. Douglas Kim of National Instruments for the internship opportunities that motivated parts of my work presented here and the excellent collaboration we have. In addition, I am thankful to Dr. Khurram Waheed of Freescale Semiconductor for his guidance, mentoring, and patience working through laborious details with me. A special thanks to Dr. Terry Henderson and Dr. Keith Lent of the Applied Research Laboratories who played pivotal roles in my decision to pursue a PhD and colored my weeks with many hours

of thought-invoking conversation (and stories). Also, I would like to thank Mr. John Tucker of Cirrus Logic, Dr. Rajan Tandon and Dr. Jill Glass of Sandia National Laboratories, and Dr. Scott Teare of the New Mexico Institute of Mining and Technology for their contributions on my path of study that have helped me maintain excitement and curiosity toward my work.

I would also like to acknowledge my colleagues of The University of Texas at Austin including Dr. Kyle Wesson, Mr. Kenneth Perrine, Mr. Yousof Mortazavi, Dr. Marcel Nassar, Dr. Jing Lin, and Mr. Haiyu Huang. They helped keep me inspired, tackle challenges, and have undoubtedly contributed to my success. I extend many thanks to my close friends and girlfriend Claudia Shenoda for helping me maintain balance and peace in my life.

Last but not least, I would like to thank my dearest family members: my father Michael Nieman, my mother Julia Nieman, brother Derek Nieman, and grandparents Sharon and Lester Byington. I extend my dearest thanks for their love and unwaivering support throughout my life.

Space-Time-Frequency Methods for Interference-Limited Communication Systems

Karl Frazier Nieman, Ph.D.
The University of Texas at Austin, 2014

Supervisor: Brian L. Evans

Traditionally, noise in communication systems has been modeled as an additive, white Gaussian noise process with independent, identically distributed samples. Although this model accurately reflects thermal noise present in communication system electronics, it fails to capture the statistics of interference and other sources of noise, e.g. in unlicensed communication bands. Modern communication system designers must take into account interference and non-Gaussian noise to maximize efficiencies and capacities of current and future communication networks.

In this work, I develop new multi-dimensional signal processing methods to improve performance of communication systems in three applications areas: (*i*) underwater acoustic, (*ii*) powerline, and (*iii*) multi-antenna cellular. In underwater acoustic communications, I address impairments caused by strong, time-varying and Doppler-spread reverberations (self-interference) using adaptive space-time signal processing methods. I apply these methods

to array receivers with a large number of elements. In powerline communications, I address impairments caused by non-Gaussian noise arising from devices sharing the powerline. I develop and apply a cyclic adaptive modulation and coding scheme and a factor-graph-based impulsive noise mitigation method to improve signal quality and boost link throughput and robustness. In cellular communications, I develop a low-latency, high-throughput space-time-frequency processing framework used for large scale (up to 128 antenna) MIMO. This framework is used in the world's first 100-antenna MIMO system and processes up to 492 Gbps raw baseband samples in the uplink and downlink directions. My methods prove that multi-dimensional processing methods can be applied to increase communication system performance without sacrificing real-time requirements.

Table of Contents

List of Tables

List of Figures

List of Acronyms

16-QAM	sixteen point quadrature amplitude modulation
8-PSK	eight point phase shift keying
AC	alternating current
ACOMMS	(underwater) acoustic communications
ADC	analog to digital converter
AMP	approximate message passing
ARL	The University of Texas at Austin Applied Research Laboratory
AWGN	additive white Gaussian noise
BER	bit error rate
BPSK	binary (two point) phase shift keying
BS	basestation
CIR	channel impulse response
CORDIC	cordinate rotation digital computer
CP	cyclic prefix
CPU	central processing unit
CQI	channel quality index
CS	compressive sensing
CSIT	channel state information at the transmitter
D8PSK	differential eight phase shift keying
DAC	digital to analog converter
DBPSK	differential binary phase shift keying
DC	direct current
DFE	decision feedback equalizer
DFT	discrete Fourier transform
DMA	direct memory access
DQPSK	differential quadrature phase shift keying
DSP	digital signal processing
EVM	error vector magnitude
FCC	(US) Federal Communications Commission
FDD	frequency-divsion duplexing

FEC	forward error correcting code
FFT	fast Fourier transform
FIFO	first in first out (digital logic queue)
FIR	finite impulse response filter
FPGA	field-programmable gate array
GALS	globally asynchronous, locally synchronous MoC
GAMP	generalized approximate message passing
G3-PLC	G3 powerline communication standard [1]
GPS	global positioning system
GPSDO	global positioning system disciplined oscillator
ICI	inter-carrier interference
IDFT	inverse discrete Fourier transform
IEEE	Institute of Electronic and Electrical Engineers
IFFT	inverse fast Fourier transform
IO	input output in reference to digital circuit
ISI	inter-symbol interference
ITU	International Telecommunication Union
ITU G.9901-9904	ITU specification for NB-OFDM PLC transceivers [2]
iid	independent, identically distributed
LFM	linear frequency modulation
LMS	least mean squared
LOS	line of sight
LTE	3GPP Long-Term Evolution cellular standard [3]
LTE-A	LTE-Advanced [4]
LUT	look up table
LVDDM	LabVIEW DSP Design Module
MCLK	master clock
MCS	modulation and coding scheme
MCX	micro coaxial cable
micro-TCA	micro telecommunications computing architecture
MIMO	multiple-input, multiple-output
MMSE	minimum mean-squared error
MoC	model of computation
MRC	maximum ratio combining
MRT	maximum ratio transmission
MU-MIMO	multi-user MIMO

NB-OFDM PLC	narrowband OFDM powerline communications
NB-PLC	narrowband powerline communications
OCXO	oven controlled crystal oscillator
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OSNR	output SNR as measured at the output of the equalizer
P1901.2	IEEE P1901.2 powerline communication standard [5]
P2P	peer to peer (FPGA to FPGA) communication
PCI-Express	peripheral component interconnect express
PDF	probability density function
PLC	powerline communications
PN	pseudo noise
PRIME	PRIME Alliance powerline communication standard [6]
PSDF	parameterized synchronous dataflow MoC
PSK	phase shift keying
PXI	PCI extension for instrumentation
QPSK	quadrature phase shift keying
RAM	random access memory
RF	radio frequency
RIO	reconfigurable input output
RMC	rack mount controller
RMS	root-mean-squared
ROBO	robust (repetition-coded) DBPSK
RRH	remote radio head
RT	real-time
SC-FDMA	single-carrier frequency division multiple access
SCR	silicon-controlled rectifier
SDMA	spatial division multiple access
SDR	software-defined radio
SINR	signal-to-interference plus noise ratio
SM	spatial multiplexing
SNR	signal-to-noise ratio
STBC	space-time block code
TDD	time-division duplexing

UE	user equipment
USRP	universal software radio platform
USRP-RIO	universal software radio platform w/ reconfigurable input output
VI	virtual instrument
WARP	wireless open-access research platform
Wi-Fi	a wireless local area network technology
WSCS	wide-sense cyclostationary
ZF	zero-forcing

Chapter 1

Introduction

1.1 Motivation

Communication is the foundation of social relationships and the essential backbone of the interconnected world. Being the vehicle used to exchange information and ideas over distance and time, communication can take many forms. Social expression, human-to-human interactions, and machine-to-machine language are all forms of communication. Local interconnections such as these can be scaled to enable large-scale collaboration and information exchange between potentially many people or things. The importance of communication and its role throughout history cannot be overstated.

The advent of electrical communication using electrical currents in wires and electromagnetic waves in air has allowed distances to shrink, responsiveness to increase, and information paths to densify. Early examples of electrical communication include Ronalds' wired telegraph of 1823 and Marconi's intercontinental wireless radio transmission of 1901. In the early 20^{th} century, analog communication techniques were used to convey analog information such as raw audio (as in a radio broadcast) and/or analog television signals. Analog communication via wired or wireless transmission was the primary form of

human made communication for most of the 20^{th} century. Techniques used to send analog information are known as *analog modulation* techniques.

In the latter part of the 1900s, the predominant information conveyed over communication channels became digital bits and bytes, parcels of information that are readily digestible by digital computers. Like their analog counterparts, digital information can be conveyed via wired or wireless means using modulation techniques known as *digital modulation*. These techniques transform bits and bytes into waveform suitable for transmission in our analog world. These techniques and their use form the basis of digital communications that has enabled astonishing levels of interconnectivity on a global and extra-planetary scale, connecting people and things like never before.

In this work, I address challenges faced in modern day digital communication systems in in three application areas of interest: (*i*) underwater acoustic, (*ii*) powerline, and (*iii*) multi-antenna cellular communications. Although the physical transmission media in these three applications widely varies—i.e. longitudinal elastic compression waves in fluids, electrical currents on conductive wired transmission lines designed for power delivery, and electromagnetic waves in air using multiple antennas at the transmitter and/or receiver, respectively—similar, multi-dimensional signal processing methods can be applied. In this work, I confront challenges in these three application areas using new, multi-dimensional signal processing methods. These methods are used to improve communication performance quantified by rate, reliability, and/or efficiency. These performance metrics are discussed next.

1.2 Background

Digital communication relies upon the modulation, transmission, reception, and demodulation of digital information encoded for suitable transmission over the communication media. In its most fundamental form, data (bits) must be transferred reliably from transmitter to receiver. In an ideal communication channel, this information is perfectly received, leading to perfect reception and lossless communication. Unfortunately, in the practical setting, transmitted information is corrupted by noise, interference, and/or other distortions that may or may not be recoverable at the receiver.

Bit error rate (BER) quantifies this reliability of a communication link by measuring how many bits have been interpreted incorrectly at the receiver divided by the total number of bits received. If we assume a pure noise channel, BER decreases monotonically with increasing signal-to-noise ratio (SNR) since there is a lower probability of bit error as noise power wanes. Forward error correcting codes (FEC) can be applied to repair bit errors at the receiver by adding additional redundant information in the transmitted data. Applying FEC in this manner trades throughput for reliability.

Throughput as measured in bits per second (bits/s) defines the achieved bit rate of a communication link. In links with multiple users, the sum rate is the sum of all individual link throughputs. Latency (in units of time) is the duration of time that elapses between valid transmission and reception of data. Links can have high throughput and high latency, low throughput and low latency, or some mix and match of the two.

Efficiency as quantified in spectral efficiency quantifies how efficient the usage of spectrum is for band-limited digital communications. The units of spectral efficiency are bits per second per Hertz (bits/s/Hz) and define the achieved throughput for a given bandwidth. For instance, you would expect a high-fidelity receiver to have much higher efficiency than that of a poor quality receiver although they operate over the same communication band. Though the two operate over the same band, the high fidelity system can achieve higher throughput (and thus higher efficiency) since it is more sensitive to transmitted signals.

Although this background has been brief, I deemed it necessary prior to the formal introduction of the three application areas this dissertation will address. For further details about the above, refer to [7].

1.2.1 Underwater Acoustic Communications

Underwater acoustic communications (ACOMMS) is a classic example of a doubly-selective communication channel—i.e. channel coherence times are on the order of the channel delay spread. The nonlinear time-varying propagation phenomena, particularly in shallow water channels, lead to the channel capacity being limited by reverberation (self-interference) rather than noise. The spatial channel remains the most promising dimension for interference reduction and spectral efficiency increases. Although prior work has investigated this topic, several shortcomings are evident. Prior work is centered around (i) a relatively small number of spatial channels, (ii) highly-complex processing

methods, and/or (*iii*) simulated channels. In this work, I develop new space-time methods for ACOMMS systems with a large number of spatial channels. My approach uses a large array of closely-spaced receive hydrophones to spatially process the communication channel. My method enhances coherent (*desirable*) paths while suppressing incoherent, highly-spread, and diffuse (*undesirable*) paths, leading to substantial increases in spectral efficiency in the multi-user scenario (up to 10× previous state-of-the-art) while being capable of real-time performance.

1.2.2 Powerline Communications

Narrowband orthogonal frequency division multiplexing (OFDM) powerline communication (NB-OFDM PLC) systems are key components of current and future smart grids. NB-OFDM PLC systems enable smart metering, distributed control, and monitoring applications by enabling connectivity over existing power lines. In practice, performance of these systems is severely limited by impulsive, non-Gaussian additive noise. The sources of this noise include rectifiers, induction motors, and switching transients of powered devices. Two primary components of this noise are (*i*) cyclostationary noise with time-periodic statistics and (*ii*) asynchronous impulsive noise. In this work, I address both of these impairments from both the transmitter and receiver sides of the NB-OFDM PLC link. For cyclostationary noise, I quantify cyclic coherence of measured noise and propose a cyclic adaptive modulation and coding scheme to enhance reliability and boost throughput at a given SNR.

My method demonstrates operating point shifts of up to 25 dB in standard-compliant noise channels. For asynchronous impulsive noise, I design a real-time impulsive noise mitigation system based on approximate message passing (AMP) that achieves streaming throughput and recovers up to 8 dB SNR of impulsive noise over a wide SNR range. Both these methods can be applied independently or together to combat NB-OFDM PLC noise.

1.2.3 Multi-Antenna Cellular Communications

To satiate mobile data demand in a wireless spectrum crunch, space truly is the final frontier. Future wireless networks rely upon spatial reuse to scale network capacity and improve network service. Research in massive multiple-input, multiple-output (Massive MIMO) suggests that even further capacity increases can be achieved by exploiting channel reciprocity with an excess number of antennas at the basestation. Channel reciprocity assumes the uplink (handset to basestation) and downlink (basestation to handset) channels are the same as long as the system operates in time-division duplexing mode and uplink downlink switching is sufficiently fast to afford fresh channel state information. Initial measurements in Massive MIMO show promising results that could allow this technology to be applied in practice. Unfortunately, real-time testbeds capable of Massive MIMO are either not publicly available, not real-time, or both. In this work, I develop a flexible platform that supports prototyping up to 20 MHz bandwidth 128-antenna MIMO. The key challenges I address are (i) scaling data rates and interfaces to 128 antennas, (ii) provid-

ing low-latency data transfer to support reciprocity even at high-mobility, and (*iii*) synchronizing time and frequency over 128 antennas, all with commercial off-the-shelf hardware. This platform is used to form LuMaMi, the world's first 100-antenna MIMO system [8].

1.3 Dissertation Summary

1.3.1 Thesis Statement

In this dissertation, I defend the following statement:

Multi-dimensional signal processing methods can be applied to underwater acoustic, powerline, and multi-antenna cellular communications to dramatically enhance communication performance without sacrificing real-time requirements.

1.3.2 Summary of Contributions

The three primary contributions of this dissertation are summarized as follows:

1. **Space-time methods for underwater acoustic communications**

   In this contribution, I develop two methods for increased underwater acoustic communication performance in the presence of strong reverberation (self-interference) typical of shallow water channels:

   (a) a windowed Doppler compensation technique coupled with a multi-stage time-adaptive decision-feedback equalizer.

(b) a space-time filtering method coupled with a decision-feedback structure to spatially filter the communication channel.

The first method is used to process more complicated motion effects observed in experimental data. The second method is used to enhance coherent (*desirable*) paths while suppressing incoherent (*undesirable*) paths. These methods are combined to form a robust underwater acoustic communication receiver for shallow water channels that is scalable to large arrays.

Results from a shallow water test environment are presented. Experimental and simulated results agree, demonstrating reduction of interference in excess of 10 dB at the equalizer output. In addition to performance gains, computational complexity is reduced by using the smaller number of pre-processed array outputs of a many-element receiver. Finally, this method is used in conjunction with a multi-user scenario to demonstrate gains in sum spectral efficiencies an order of magnitude above prior state-of-the-art.

2. **Time-frequency methods for powerline communications**

In this contribution, I characterize cyclic statistics of powerline noise. I confirm that the noise is predominantly cyclostationary with dominant period equal to half the AC cycle. Other noise such as asynchronous impulsive noise and narrowband noise are observed. I use these results to develop two methods:

(a) an adaptive cyclic modulation and coding scheme over time and frequency

(b) an implementation of an asynchronous impulsive noise mitigation technique using a time-frequency signal structure

First, I demonstrate the utility of this adaptive cyclic modulation and coding scheme (MCS) adaptation using measured noise samples and the G3-PLC MCS set. My results show that this scheme is capable of offering up to a 2× increase in throughput using the current G3-PLC modulation set. In addition, operation point shifts of up to 25 dB are observed, leading to substantial gains in efficiency and/or throughput. Second, I design and implement a real-time OFDM receiver with AMP impulsive noise mitigation on FPGAs. My design flow is to (*i*) model the receiver in synchronous dataflow (SDF), (*ii*) convert the receiver to fixed-point data and arithmetic, and (*iii*) synthesize the fixed-point SDF model onto two Xilinx Vertex-5 FPGAs. In the third step, I generalize the SDF receiver model to a globally asynchronous, locally synchronous computational model in LabVIEW DSP Design Module. My OFDM transceiver testbed achieves real-time streaming throughput at G3-PLC rates and mitigates up to 8 dB SNR of impulsive noise.

3. **Space-time-frequency methods for multi-antenna cellular communications**

In order to make Massive MIMO prototyping more realizable, I

have developed a scalable platform for large antenna count MIMO using commercial off-the-shelf hardware. The key contributions of this design are

(a) scaling processing and interfaces to 20 MHz bandwidth with up to 128 antennas

(b) developing a low-latency routing and frame structure to support reciprocity in high-mobility channels

(c) synchronizing over time and frequency up to 128 antennas at the basestation

The design is capable of processing up to 491.5 Gbps of raw baseband samples (128 antennas $\times$ 120 MS/s $\times$ 32 bits/S) in both the uplink and downlink directions with a channel estimate to precoder turnaround time of less than 500 μs. The design is reconfigurable and flexible to accommodate various signal structures and system requirements for current and future Massive MIMO prototyping efforts. This platform is used in LuMaMi, the world's first 100-antenna MIMO testbed.

1.4 Organization

The remainder of this dissertation is organized as follows.

Chapter 2 presents space-time methods developed for underwater acoustic communications. It starts with a brief overview of the challenges of the shallow underwater acoustic communication channel. A physical channel model

is then presented along with prior work in high speed underwater communications. Two space-time methods are then developed. These methods are analyzed using simulation and experiments conducted in a shallow water test environment.

Chapter 3 presents measurements of noise in powerline communication (PLC) systems. Cyclic spectral analysis is applied to characterize the measurements. Next, a system model is discussed along with prior work in noise mitigation. Then, a comprehensive model that addresses the two primary noise components in narrowband PLC is presented. This model forms the basis of two time-frequency methods that are then developed. Finally, these methods are evaluated using measured and simulated noise and a real-time testbed.

Chapter 4 begins with a discussion of the current state-of-the art for large antenna count (Massive MIMO) wireless system prototyping. Next, a system model and a parameterized synchronous dataflow model of the processing in a Massive MIMO basestation are presented. Then, three key challenges of implementing a prototype system are addressed using space-time-frequency methods. Initial measured results are presented.

Chapter 5 concludes the dissertation with a summary of the presented work and outlines possible avenues of future research.

Chapter 2

Space-Time Methods for Wideband Underwater Acoustic Communication Systems

2.1 Motivation

Underwater acoustic communications (ACOMMS) is a classic example of a doubly-selective communication channel, where channel coherence times are on the order of the channel delay spread. The nonlinear time-varying propagation phenomena, particularly in shallow water channels, lead to the channel capacity being limited by reverberation (self-interference) rather than noise. The spatial channel remains the most promising dimension for interference reduction and channel capacity increases. Although prior work has investigated this topic, several shortcomings are evident. Prior work is centered around (*i*) a relatively small number of spatial channels, (*ii*) highly-complex processing methods, and/or (*iii*) simulated channels. In this work, I develop new space-time methods for ACOMMS systems with a large number of spatial channels. My approach uses a large array of closely-spaced receive hydrophones to spatially process the communication channel. My method enhances coherent (*desirable*) paths while suppressing incoherent, highly-spread, and diffuse (*undesirable*) paths, leading to substantial increases in channel capacity (up to 10× previous state-of-the-art) while being capable of real-time performance.

Table 2.1: Typical values for a short to medium range ACOMMS system.

range (km)	0.02 – 10
bandwidth (kHz)	1 – 100
center frequency (kHz)	5 – 100
ratio of attainable speed to propagation speed for typical mobile user	0.00 – 0.01

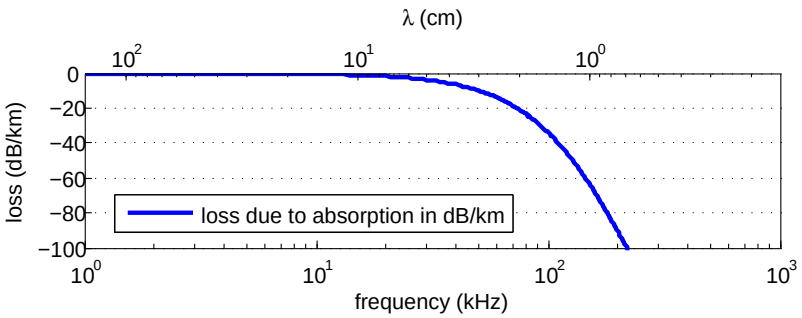

Figure 2.1: Acoustic loss due to absorption as calculated using formula in [9].

2.2 Background

Underwater acoustic communications (ACOMMS) is a technique for sending and receiving information wirelessly underwater using acoustic waves. Similar to terrestrial wireless signals that carry information using modulated electromagnetic waves, ACOMMS signals are subject to impairments such as time and frequency dispersion, noise, and interference. However, two key differences in propagation physics strongly differentiate ACOMMS from terrestrial wireless. First, underwater acoustic waves propagate ~200,000× slower than electromagnetic waves in free space (1,500 m/s for a longitudinal acoustic wave in 13 °C seawater versus 300,000,000 m/s for a transverse electromagnetic wave in free space). Second, ACOMMS channels are strongly band-limited, arising from various absorptive mechanisms including viscosity, heat conduction, strain relaxation, and boundary-layer effects. This leads to the pronounced lowpass channel response shown in Figure 2.1.

These two effects distinctly impact the design of ACOMMS systems. Table 2.1 lists common system parameters for short to medium range (high to medium frequency) ACOMMS systems. Such systems typically operate within a generous portion of the absorption-limited spectrum (see Figure 2.1) due to low ambient background noise and the lack of regulatory measures in many underwater communication scenarios [10]. This leads to many systems being inherently wideband with respect to their center frequencies. To illustrate, an ACOMMS system operating with 30 kHz of bandwidth centered at 30 kHz has a relative bandwidth of 1.0, whereas a 20 MHz LTE band (18 MHz bandwidth at 2.5 GHz) has a relative bandwidth of 0.0072 [3]. Thus narrowband assumptions common in terrestrial wireless systems cannot be applied.

A second resulting effect of the physical propagation environment arises from the relatively slow propagation speed of sound in water. Slow propagation speed increases the severity of Doppler-induced distortions in mobile communication scenarios. For instance, in the case of ACOMMS, it is physically possible to achieve motion of 1% of the propagation speed at a speed of 30 knots (15.43 m/s), a speed readily attainable by modern fast submersibles and surface ships. On the other hand, an SR-71 jet flying at its max recorded speed of Mach 3.4 (1157 m/s) achieves a mere 0.00034% of the free space electromagnetic propagation speed. In addition, slow sound speed also substantially lengthens the impulse response duration of practical ACOMMS channels, particularly in shallow water environments. Under such conditions, the root mean suare (RMS) delay spread is on the order of the channel coherence time,

Table 2.2: Key attributes of ACOMMS vs. a terrestrial wireless system.

	ACOMMS, $f_c = 30$ kHz, high mobility (30 kts)	LTE urban, $f_c = 2.5$ GHz, max speed (350 km/h) [11]
RMS delay spread	3.3 ms	2 μs
typical coherence time	1 ms	1.2 ms
Doppler dilation factor	0.01	3.24×10^{-7}
relative bandwidth	1.0 for 30 kHz bandwidth	0.0072 for 18 MHz bandwidth

making the shallow water ACOMMS channel doubly-selective in time and frequency. A summary of key differences between typical ACOMMS and LTE channels is provided in Table 2.2.

Modern, high data rate ACOMMS has been enabled using a number of advanced, time-adaptive processing methods. These methods dynamically isolate and track time-dispersed and Doppler-spread paths, returning coherency over a finite time-frequency span of the received signal [12, 13]. Though good performance has been achieved in real channels, coherency can only be supported over a fixed time-frequency span. This is caused by too much Doppler spread and/or time-dispersion. *Residual reverberation* (self-interference) outside this span limits the achievable capacity of ACOMMS systems. This manifests as inter-symbol interference (ISI) in single carrier systems and inter-carrier-interference (ICI) in multicarrier systems. Typical multicarrier systems use orthogonal frequency division multiplexing (OFDM) signals [14, 15].

Despite prior work (detailed in Section 2.2.2), an empirical *range-rate bound* is found to be 40 kbits/s/km over real channels. This bound has yet to be substantially broken. In this work, I apply new space-time processing methods to improve communication performance beyond this bound.

2.2.1 System Model

The acoustic pressure waveform at the receiver can be modeled as a superposition of propagation paths (eigen-rays) that are delayed and affected by wideband Doppler effects [16], resulting in a signal $y(t)$ that is doubly spread and scaled in both time and frequency,

$$y(t) = \sum_{i=0}^{N_p-1} a_i x\left(\frac{1}{d_i}t + \Delta_i\right) + n(t). \tag{2.1}$$

Here, N_p is the total number of paths, $n(t)$ is the noise amplitude at time t, and the remaining terms are characteristics of the i^{th} path—its amplitude a_i, its Doppler factor d_i, and its delay Δ_i. The i^{th} Doppler factor d_i can be approximated by projecting the relative velocity vector (between the receiver and the source) $\vec{v} = \vec{v}_r - \vec{v}_s$ onto the i^{th} position vector, $\vec{p}_i$, pointing from the receiver to the i^{th} image of the source. The the i^{th} delay d_i is computed as

$$d_i = 1 + \frac{\vec{v} \cdot \vec{p}_i}{|\vec{p}_i|\, c}. \tag{2.2}$$

Using the i^{th} Doppler factor, the length of the i^{th} path vector, and the speed of sound in water, c (assumed constant), the i^{th} delay can be computed,

$$\Delta_i = \frac{|\vec{p}_i|}{c d_i}. \tag{2.3}$$

The i^{th} path amplitude is reduced by reflection, geometrical spreading and absorption. Thus, the i^{th} path amplitude is the product of all N_r reflection coefficients, Γ_r, the wavelength dependent geometrical spreading expression

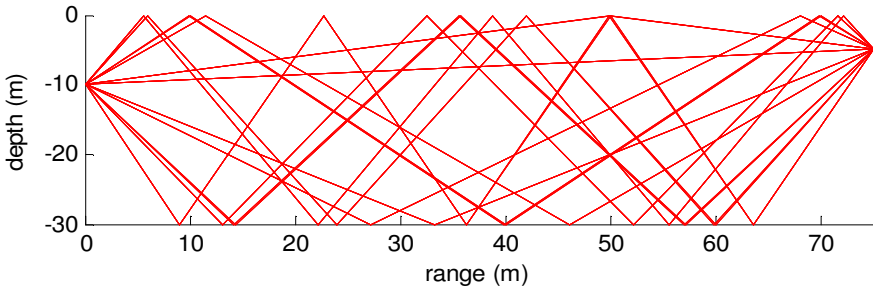

Figure 2.2: Realization of a three medium (air, water, sediment) ray-based channel model with number of paths $N_p = 11$. Here, a transmitter at 75 m range and 5 m depth communicates with a receiver at 10 m depth.

with geometry factor k (where $k = 1$ for spherical), and the absorption loss (assuming constant absorption coefficient α over frequency and time),

$$a_i = \left(\prod_{r=0}^{N_r - 1} \Gamma_r \right) \left(\frac{\lambda}{4\pi \left| \vec{p}_i \right|^k} \right) e^{-\alpha |\vec{p}_i|}. \tag{2.4}$$

Note that there are more complicated effects in physical channels that have been neglected—e.g. time-varying path parameters, wavelength-dependent absorption and sound speed, and Doppler effects that change over time due to platform acceleration [17]. These effects can be added to form a more complete physical ACOMMS channel model.

A realization of this channel model for a three medium (air, water, sediment) vertical channel with no refraction is shown in Figure 2.2. More complicated channel effects are shown in Figure 2.3, where space-time-frequency received signal strength distributions are shown for a shallow water ACOMMS channel of similar geometry. This four-dimensional channel model is based on acoustic data collected and analyzed in experiments listed in [17–20]. This serves as a qualitative model of measured four-dimensional channels.

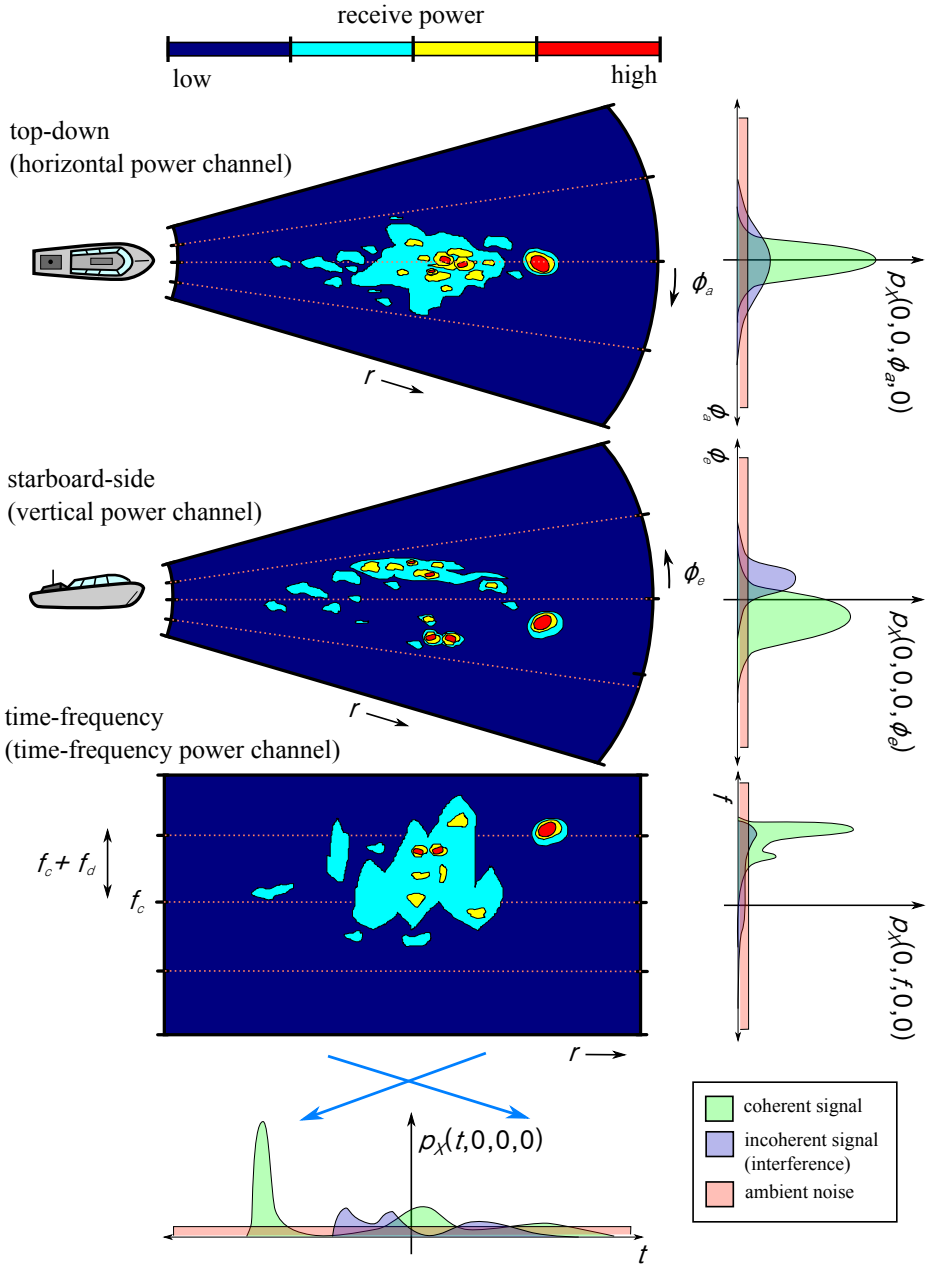

Figure 2.3: The power distribution of ACOMMS signals along four dimensions (from top to bottom) — azimuth angle, elevation angle, frequency, and down-range (time). Receive power along these dimensions is denoted by color mapping to the overhead color bar.

Figure 2.3 shows the four-dimensional ACOMMS channel based on measurements discussed in [17–20]. The four dimensions are visualized in three planes with relative receive signal strength shown in color. Marginal power distributions for three components of the received signal (coherent signal, incoherent signal, and ambient noise) are shown along each of the four dimensions (azimuth angle ϕ_a, elevation angle ϕ_e, frequency f, and time t) as $p_{\mathrm{X}}(0, 0, \phi_a, 0)$, $p_{\mathrm{X}}(0, 0, 0, \phi_e)$, $p_{\mathrm{X}}(0, f, 0, 0)$, and $p_{\mathrm{X}}(t, 0, 0, 0)$, respectively.

In the top-down channel, the receive power has a narrow distribution along the azimuth dimension owing to the unconstrained horizontal channel in large bodies of water. The received signal power has a large hot spot—i.e. the line-of-sight (LOS) path. Other scatterers appear closer in range but are difficult to resolve in azimuth angle. The majority of scattered energy is resolvable only in elevation angle. In the vertical channel, several distinct features are apparent. First, a wide swath of diffuse scattering is observed near the surface. Second, the specular line-of-sight (LOS) path (now clearly isolated spatially) can be seen at a lower elevation angle (i.e. it is lower in depth than the receiver). Third, a group of specular scatterers are located near the lake bed. The time-frequency statistics of these scatterers are revealed by separating the time and frequency distribution of the receive power (third plot from top). As seen in the time-frequency distribution, the specular (line-of-sight) path has a narrow Doppler distribution along with the scatterers from the lake bed. The wide swath of scatterers from the surface are highly spread in space and frequency, owing to surface choppiness and wave motion.

The potential for individual path resolution can be analyzed in the surrounding plots which show the marginal distribution of received signal energy collapsed along the four dimensions $p_{\mathrm{X}}(0, 0, \phi_a, 0)$, $p_{\mathrm{X}}(0, 0, 0, \phi_e)$, $p_{\mathrm{X}}(0, f, 0, 0)$, and $p_{\mathrm{X}}(t, 0, 0, 0)$ and binned into three categories: (i) the coherent LOS and coherent reverberations are shown in green, (ii) the incoherent signal (owing to too much Doppler spreading and/or rapid fading) is shown in blue, and (iii) ambient background noise is shown in red and is uniformly distributed (or white) along every dimension.

The marginal distributions reveal the following. First, the horizontal channel demonstrates a symmetric distribution vs. azimuth angle ϕ_a. The optimal spatial filter along this dimension is simply a matched-filter that is matched to the point spread function of the seeing aperture. Second, the vertical cross-section of the power distribution of the channel reveals that there are resolvable clusters of coherent and incoherent signals, a filtering scheme for which will be discussed later in Section 2.3.2. Third, in frequency, the receive power distribution is distributed according to a single-sided Jakes' Doppler spectrum. The distribution is single-sided owing to the fact that the receiving system is not omnidirectional and is only sensitive to scatterers in the forward half-space. Since the scatterers experience a Doppler-induced dilation proportional to the projection of the velocity vector onto the scattering ray, for many scatterers this is less than or equal to the Doppler shift. Finally, along the time axis, signals are similarly separable, owing the long delay spread of the ACOMMS channel and the wide bandwidth of the system.

2.2.2 Prior Work

Modern, high data rate ACOMMS has been made possible using a number of advanced, time-adaptive processing methods. These methods dynamically isolate and track time-dispersed and Doppler-spread paths, returning coherency to a finite time-frequency span [12, 13]. In Figure 2.3 this is observed as the deconvolution of the spread along the frequency and time axes in order to reduce inter-symbol interference (ISI) in single carrier systems and inter-carrier-interference (ICI) in multicarrier systems. Ideally, such processing could compensate for all time and frequency dispersion, resulting in a delta function in time and frequency of power equal to total received signal power. Unfortunately, practical limitations such as limited time-frequency resolution and finite precision arithmetic lead to estimation and compensation inaccuracies. In [21], the authors note that after processing and recombining a certain number of paths, subsequent results lead to negligible gains due to estimation and compensation inaccuracies. Uncompensated paths, or *residual reverberation* (self-interference), outside this span is often many orders of magnitude higher in power than ambient noise, leading to a reverberation-limited channel characterized by its signal-to-interference plus noise ratio (SINR).

In [10], the authors analyze achieved results of various ACOMMS systems. They find an empirical *range-rate bound* of 40 kbits/s/km that has yet to be substantially broken. This range-capacity bound is contributed by low SINR in reverberation-rich channels and absorption-limited bandwidths. A promising avenue for increased channel capacity has come from spatial pro-

Table 2.3: Prior empirical results for multi-channel ACOMMS systems.

Method	Number of Elements/ Array Geometry	Center Frequency (kHz)	Range (km)	Rate (kbps)	Bound (kbps)	Sum-Rate Capacity (bps/Hz)
Multi-Channel Adaptive Equalization [22]	8 vertical or horizontal line, multi-user	23	0.5-2	2.8	20	0.56
Channel Eigen Decomposition [21]	64 cross-beam	24	3.2	16	12.5	1.0
Spatial Filter then Equalize [23]	32 vertical line	1.2	10	0.4	4	1.0
OFDM [15]	8 vertical	25	1	24	40	2.0
Single-Carrier MIMO [24]	8 vertical receive, 2 vertical transmit	17	1-3	32	13.3	2.3

cessing using multiple transmit and receive elements [21–23]. Increased spatial degrees of freedom allow for further path separability and processing, improving signal and/or reducing interference and resulting in higher SINR. Most prior work concerns (*i*) a relatively small number of spatial channels and/or (*ii*) highly-complex processing methods that are not tractable for current real-time system designs. Hydrophone arrays used in these systems are typically one dimensional and consist of elements separated by many wavelengths (i.e. vertical line arrays suspended from a surface ship).

Table 2.3 summarizes key prior work in many-element ACOMMS systems, operating bands, array geometry, range, rate, the theoretical rate bound according to [10], and the achieved sum-rate capacity in bits-per-second per Hz. As observed in the table, though these techniques achieve good sum rates, they do not achieve rates substantially higher than the *range-rate bound*.

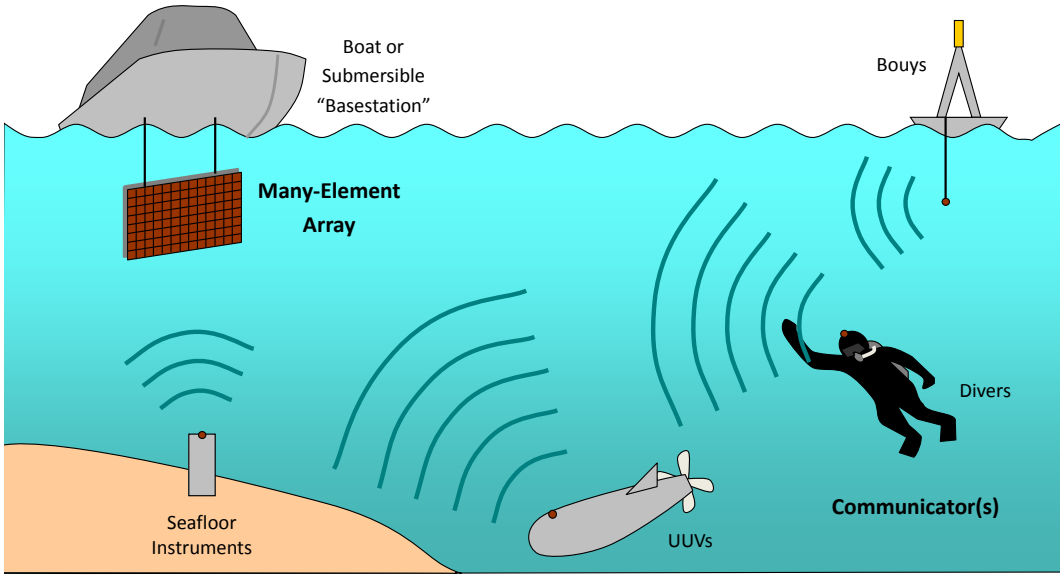

Figure 2.4: A multi-user ACOMMS scenario where potentially many users communicate with a many-element array mounted on a boat or submersible.

2.3 Large-Array Space-Time Processing

Using inference from the channel characteristics shown in Figure 2.3 and the works listed in Table 2.3, I have developed new space-time processing methods to improve ACOMMS performance beyond prior state-of-the-art. My methods are focused on the use case where a large, many-element basestation is used to communicate with potentially many single-element mobile units as shown in Figure 2.4. The large-array processing methods I develop consist of two components: (*i*) enhanced time-varying Doppler estimation and correction methods and (*ii*) computationally-tractable space-time filtering methods for shallow water ACOMMS channels. These contributions are discussed in Sections 2.3.1 and 2.3.2, respectively.

2.3.1 Windowed Doppler Estimation/Compensation

In mobile ACOMMS scenarios, relative motion between the transmitter and receiver manifests as Doppler distortions on the received signal. Since RMS delay spread is of the order of the coherence time (see Table 2.2), signaling often occurs over many coherence times. Doppler spread results in a number of independent and rapidly-varying paths that much be tracked and compensated at the receiver in order to be coherently recombined and improve SINR. First-order, or linear, Doppler distortion is typically estimated and corrected. This accounts for differences in clocks between transmitter and receiver in addition to bulk linear motion over the signaling interval. More complex motion in practice leads to higher order Doppler effects as seen in Figure 2.5. In this figure, time-domain QPSK symbols, which are ideally received as four 90°-spaced symbols are instead received with linear and higher order Doppler distortions. Figure 2.5 shows the received phase of each symbol after linear correction, which compensates for constant progressive phase rotation over the signaling duration. Higher order motion effects remain present.

To compensate for these higher order Doppler distortions, I present windowed piecewise estimation and correction methods to more accurately compensate for higher order Doppler effects. I apply these methods on data collected using a mobile transmitter with a stationary receiver in a shallow water test environment. Received data is processed using a MATLAB software receiver shown in Figure 2.6 using several competing decision-feedback equalizer (DFE) designs discussed in detail in [18].

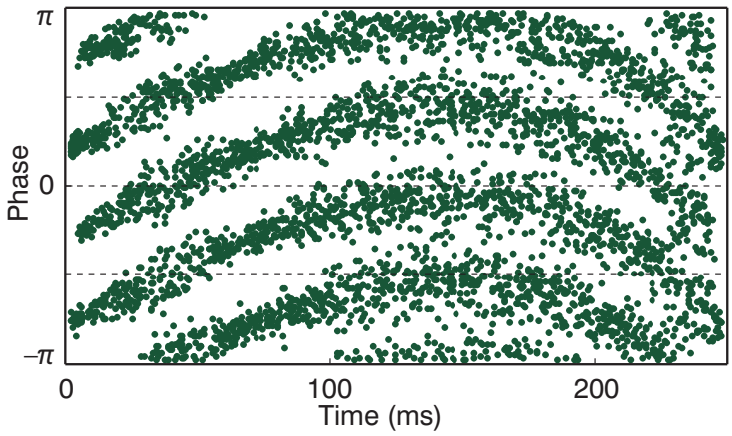

Figure 2.5: Typical residual Doppler effects on the phase of a received QPSK signal after a bulk Doppler offset correction. Voronoi regions for the maximum-likelihood detector are denoted by dashed lines. Here, maximum phase error is apparent at the peak offset near 150ms.

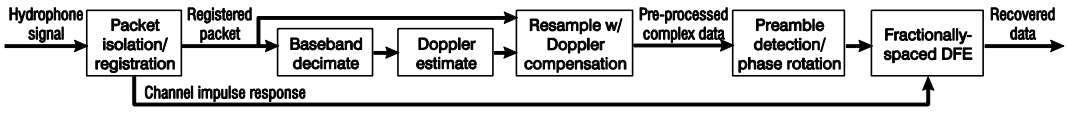

Figure 2.6: Software receiver used to analyze time-varying Doppler. Samples from the hydrophone are synchronized then decimated. A windowed Doppler estimate is then performed and compensation is applied. Doppler-compensated samples are then processed using a time-adaptive equalizer.

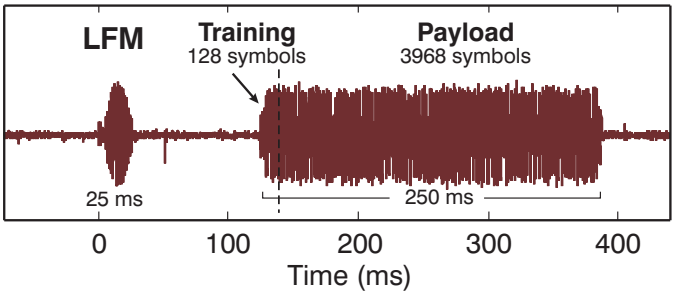

Figure 2.7: Waveform consisting of an LFM, training, and payload. The waveform pictured is sampled from a 230 m long shallow underwater channel.

The software receiver shown in Figure 2.6 operates on raw real-valued samples from one or multiple receive hydrophones. Packets are synchronized using a special synchronization signal, in this case a linear frequency modulation (LFM) sweep. This particular waveform has an autocorrelation function that is insensitive to frequency shifts [16], allowing packets to be detected even in high Doppler environments. This time-registers packets that are then decimated to baseband. A Doppler estimate is then performed accompanied by a time-varying resampling operation. This de-spread complex-valued signal is then input to a preamble detection block to correct for residual phase offsets to ensure the best initial conditions for equalization. The resulting samples are then input to a fractionally-spaced decision-feedback adaptive equalizer. Resulting symbols are then bit sliced to measure receiver performance.

The waveform used for testing is shown in Figure 2.7. The leading LFM precedes the modulated payload and preamble data. A guard period is then inserted in order to minimize self-interference of the synchronization signal with the payload data. Signaling over the 250 ms duration allows for large-scale components of the channel to remain relatively stationary while providing a means to track small-scale changes over the packet duration. The waveform shown was sampled at a 230 m long shallow water channel in the test environment shown in Figure 2.8. For this work, Doppler shift is estimated and compared using three methods: (*i*) self-referenced correlation [25], (*ii*) nonlinear PSK rectification (squaring) and windowed DFT [17], and (*iii*) pilot tone frequency over windowed intervals.

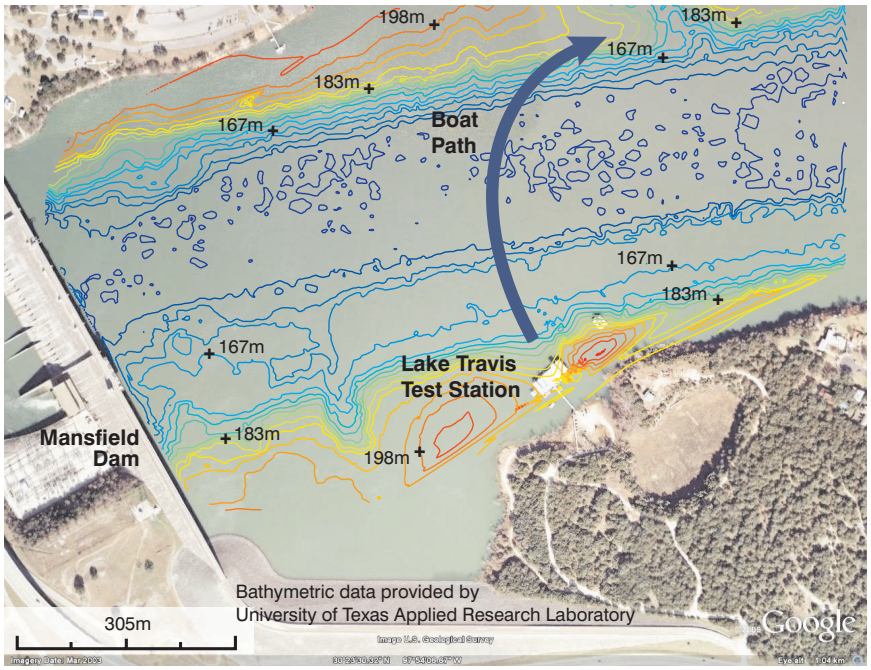

Figure 2.8: Aerial view of the Lake Travis test environment with lake bed elevations given above sea level. The water level is 198 m. The array receiver is located at the Lake Travis Test Station.

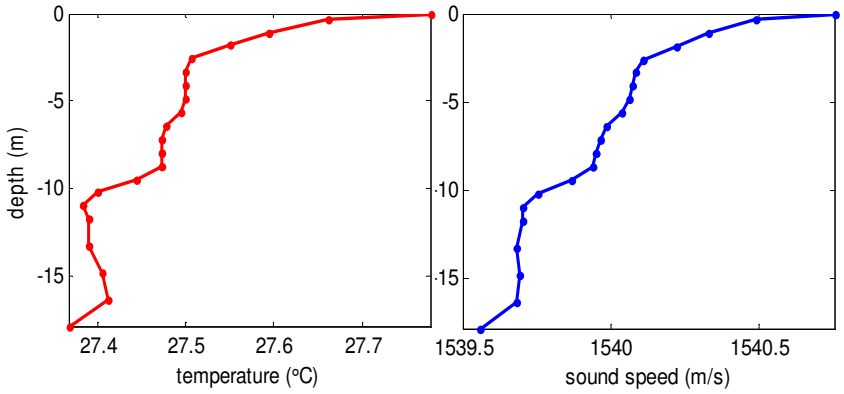

Figure 2.9: Thermal variation as a function of depth leads to the refractive sound velocity profile shown at right. These conditions (measured Nov 6, 2009 at Lake Travis), even though mildly downward refractive, are relatively favorable propagation conditions. Line of sight propagation can be achieved up to distances of 1.5 km or more.

2.3.1.1 Results from a Shallow Water Experiment

Data for this study was collected November 6, 2009 at the ARL Lake Travis Test Station. Lake Travis is located approximately 15 km northwest of downtown Austin, TX. The lake has an estimated maximum depth of 37 m [26]. As seen in Figure 2.8, the lake has a nearby dam and hilly former river bed terrain that contribute to high reverberation. Challenging, highly-refractive propagation conditions are typical during temperature extremes in July and January. Fortunately, during testing the propagation conditions were only mildly downward refractive as seen in Figure 2.9.

An omnidirectional transmitter was hung to various depths between 1 m and 8 m from a small research vessel. Transmissions were made at one of 9 positions listed in Table 2.4 (and shown in Figure 2.10), where the boat was either docked to the test station barge (Pos. 1), slowly drifting in the lake (Pos. 2-3), or in motion (Pos. 4-5). In particular, Pos. 4 involved a vertical, oscillating motion of about 0.5 Hz (from depths of about 5.7 m to 7.2 m) to simulate boat motion from heavy waves, and Pos. 5 involved the towing of the transducer at speeds of 5 km/h at varying depths no greater than 5 m.

The transducer on the boat was connected to an amplified DAC operating at a 500 kHz sample rate. This DAC was the IOTech Personal Daq/3000 [27]. At unity gain, the transducer placed about 1 W of acoustic power into the water, although the gain was manually configured at each position in order to ensure adequate signal level at the receiver (e.g. reasonable signal to electronic noise ratio and no clipping). The receiver, mounted at the test station, con-

Table 2.4: Boat and transducer positions.

Pos.	Range	Motion	TX Gain
1	15 m	Docked to barge	-13 dB
2	325 - 375 m	Free-floating	-7 dB
3	1235 - 1285 m	Free-floating	-3 dB
4	185 - 255 m	Simulated vertical "wave" motion	-7 dB
5	300 - 80 m	Towing at ∼3kts	-10 dB

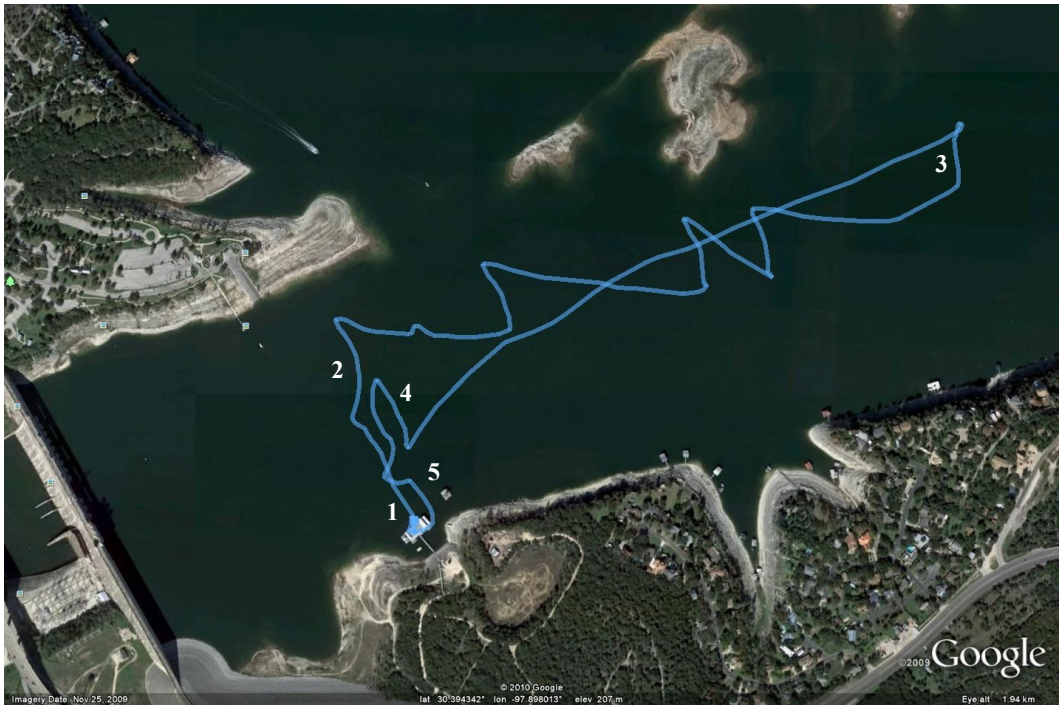

Figure 2.10: GPS map of typical transmitter path during experiment. Various transmit positions are enumerated where the mobile boat communicates back with the array mounted to the barge located at position 1.

sisted of a planar array of 5 directional hydrophones horizontal and vertical half-power beamwidths of approx. 45° and 10°, respectively. They were arranged in an L-shaped array with dimensions provided in [28]. The array was submerged to a depth of 4.6 m and rotated (when needed) to generally face the transmitter. The received signals were preamplified and sampled at 200 kHz by another 1 MHz multiplexed DAC [27]. Although the DAC sampled each channel sequentially in each 5 μs window, the sequences were upsampled and shifted to allow the samples to be treated as simultaneously coincident across all five channels. Only one channel was processed for analysis. Multi-channel spatial processing techniques applied are detailed in Section 2.3.2.

Figure 2.11 shows channel impulse magnitude for near and far ranges in line-of-sight conditions. Longer reverberation time constants and more complex power delay profiles are often observed as range increases. The loudest path is not always the earliest received path, especially for far range and non-line-of-sight conditions. Attenuation also varies with frequency and transmit power due to transducer sensitivity [29]. Note that the large reverberation tail in Position 3 is more challenging to process than that in Position 1.

To evaluate the different Doppler estimation and correction methods, different configurations listed in Table 2.5 were used to process the ~360 packets recorded during the test [28]. Configurations A-I correspond to different window sizes applied to the three methods (*i*) self-referenced correlation [25], (*ii*) nonlinear PSK rectification (squaring) and windowed DFT [17], and (*iii*) pilot tone frequency over windowed intervals.

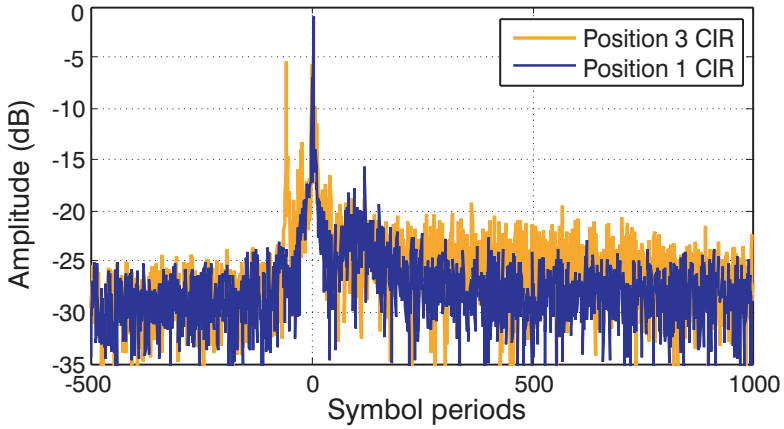

Figure 2.11: Channel impulse responses (CIR) magnitude for near and far ranges. Position 1 range is 30 m and Position 3 range is ∼1260 m. The large reverberation tail in Position 3 is more challenging to process than that in Position 1 due to its long dispersion.

Table 2.5: Doppler estimation configurations.

	A	**B**	**C**	**D**	**E**
Method	Moose	DFT	DFT	DFT	DFT
# Windows D =	1	1	2	4	8
	F	**G**	**H**	**I**	
Method	Pilot	Pilot	Pilot	Pilot	
# Windows D =	1	2	4	8	

Column headings of Figure 2.12 denote the use of static and adaptive decision-feedback equalizers defined in [18], while the main portion shows the results of each individual decoding operation corresponding with one sampled packet (grouped among individual receivers within the 5-receiver array). These groupings are then arranged sequentially according to each received packet, as 5 or 6 packets were transmitted at each position. Lighter shades signify

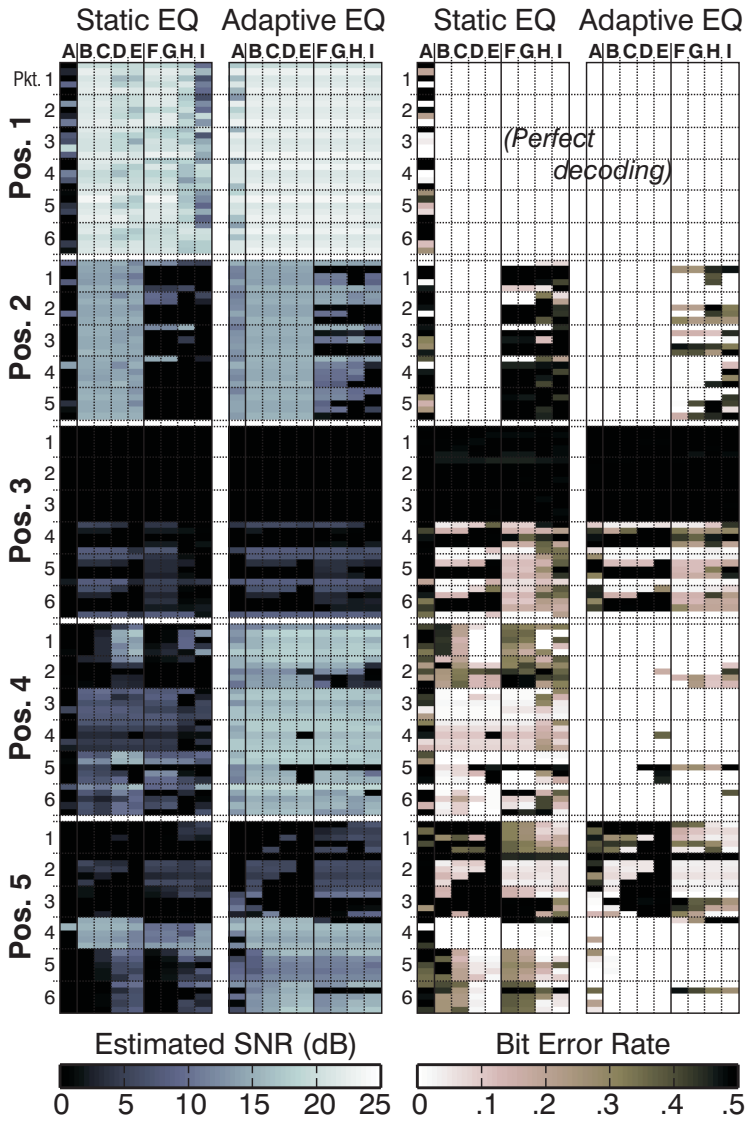

Figure 2.12: Received SINR and BER performance of QPSK single element receiver at five positions given in Table 2.4 using the configurations listed in Table 2.5.

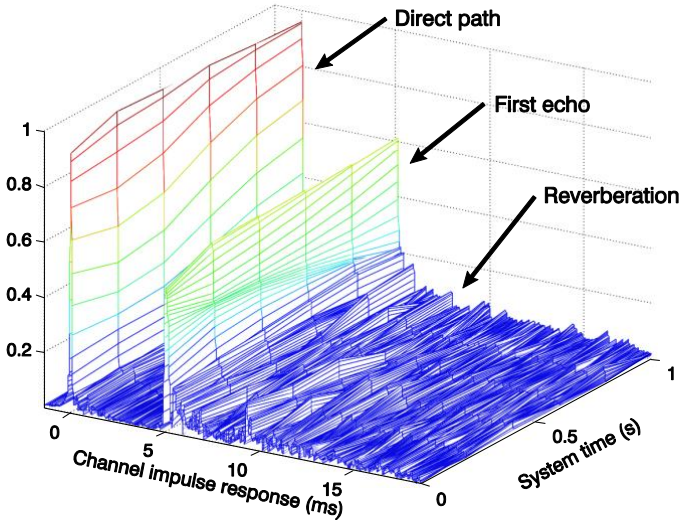

Figure 2.13: Channel magnitude response as a function of time at 250 m range. The direct path and first echo magnitude are approximately stationary over the system time whereas reverberant paths greatly vary over the system time.

successful decoding operations; darker shades signify failure. Figures of merit are estimated SINR at the equalizer output computed using inverse error-vector magnitude (EVM) and bit-error rate (BER). It is apparent in these results that the best performing (highest estimated SINR and lowest BER) correspond to the nonlinear PSK rectification (squaring) and windowed DFT method. Adaptive equalization applied to the less granular ($D = 1$) Doppler compensation method is most robust, though more granular ($D > 1$) methods assist the static equalizer with the main path shown in Figure 2.13. Figure 2.13 shows the evolution of the magnitude of various components of the channel impulse response magnitude as a function of time. Time-adaptive processing is then coupled with spatial filtering method discussed in the following section. These methods together provide robust performance in shallow water channels.

2.3.2 Wideband Monopulse Space-Time Filtering

Monopulse processing is a technique used to improve the angular precision of tracking systems to a value finer than the beamwidth of the seeing aperture. This has the advantage of allowing for more compact tracking systems since the beamwidth of an aperture is physically constrained by its cross-section in wavelengths (i.e. beamwidth within a given plane is $\sim 1.22\lambda/B$, where B is the baseline of the aperture in that plane). The wideband monopulse method, also known as the space-time filtered gradient method, is an extension of classical sum-and-difference monopulse for use in wideband systems [30–32]. A comprehensive review of direction finding methods and monopulse processing is given in [33].

A discrete array of sensors can be used to sample an acoustic field across a given aperture. It can be steered to give an output that is most sensitive for plane waves arriving from a particular direction by the following steps: (*i*) delay each element output such that plane waves from that direction arrive simultaneously and (*ii*) add those delayed outputs. The resulting angular response pattern can be tailored by applying weights prior to summing. For this reason, the resulting *beam output*, $s(t)$, is said to be formed by a delay-weight-sum process. This process is most approachable when considering an M-element uniformly-spaced line array of inter-element spacing D such as the one shown in Figure 2.14. Here, a beam can be steered toward a unit sinusoid of frequency f at angle θ_s by summing all elements outputs $x_i(t)$ with a progressive phase delay and corresponding weight w_i (chosen for desired

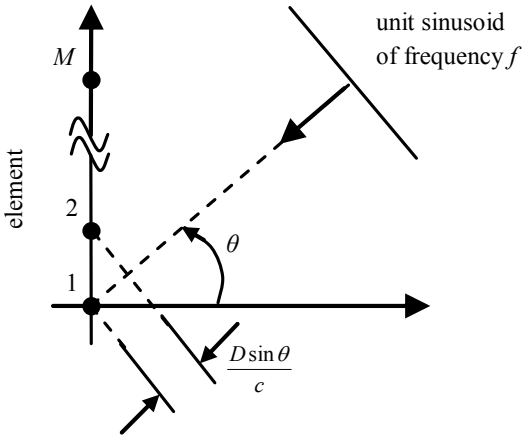

Figure 2.14: An M-element uniformly-spaced line array of inter-element spacing D is illuminated by a plane wave arriving from a source at angle θ from broadside. Elements can be progressively delayed by $D\sin\theta/c$ and summed to maximize response at angle θ.

beampatterns—e.g. edges tapered for reduced sidelobes [34]):

$$y(t) = \sum_{i=0}^{M-1} w_i x_i(t) e^{j2\pi f \frac{D\sin\theta_s}{c} i}. \tag{2.5}$$

For wideband signals, phase shifts proportional to frequency must be applied to achieve the required delay. This leads to the result:

$$y(t) = \int_{-\infty}^{\infty} \sum_{i=0}^{M-1} w_i X_i(f) e^{-j2\pi f \frac{D\sin\theta_s}{c} i} e^{j2\pi f t} \ \mathrm{d}f. \tag{2.6}$$

Here, $X_i(f)$ denotes the continuous Fourier transform of the element output $x_i(t)$. A discrete Fourier transform is commonly used in practice, as element outputs are typically digitally sampled after being pre-amplified and conditioned. Discrete-time processing methods are then used. Though many factors must be considered for practical array design (e.g. minimum spatial

sampling to avoid aliased response patterns, or grating lobes, element response, variability and coupling, etc.), the delay-weight-sum beamforming process can be applied in this manner.

A wide array needs many azimuthally-narrow, steered beams to cover the field of view (see Figure 2.16). To limit the complexity of vertical beamforming, wideband monopulse processing generates a pair of beam outputs, $s_0(t)$ and $s_1(t)$, as discussed in [15-17], with the following linear relationship:

$$s_1(t) \approx (\sin\theta - \sin\theta_s) s_0(t) \tag{2.7}$$

Here, θ is the elevation angle of the source and θ_s is the vertical beam steering angle. Two uses of (2.7) are (i) direction finding and (ii) null steering. These beam outputs $s_0(t)$ and $s_1(t)$ can be generated in five steps: (i) assume the array is a uniformly-spaced vertical stack of M identical sub-arrays forming identical azimuthal beams, (ii) form w_b, a vector of $M-1$ shading coefficients selected for desired vertical beam patterns (e.g. a Hanning window), (iii) form two M-element shading vectors $w_0 = [1,1] \otimes w_b$ and $w_1 = [1,-1] \otimes w_b$, where $\otimes$ denotes discrete convolution, (iv) form two beam outputs by applying w_0 and w_1 to the M sub-arrays using classical delay-weight-sum beamforming for vertical steering angle θ_s, and (v) apply filters $h_0(t)$ and $h_1(t)$ to those beam outputs, where $h_0(t)$ and $h_1(t)$ have the following relationship:

$$h_0(t) = \left(\frac{D}{2c}\right) \frac{\mathrm{d}}{\mathrm{d}t} h_1(t). \tag{2.8}$$

Here, D is the spacing between elements and c is the wave propagation speed in the medium of interest. These steps give the frequency domain relationship

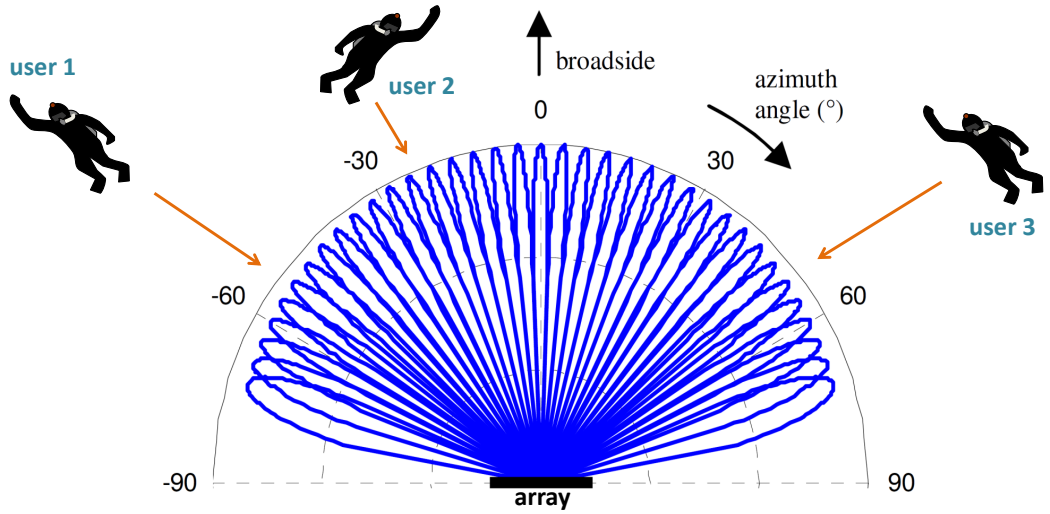

Figure 2.15: An overhead view of the response of a fan of narrow azimuthal beams produced by delay-weight-sum beamforming a horizontally-wide array (array centered on the lower axis shown in heavy black, 20 dB/div). The array serves multiple simultaneous users using spatial division multiple access.

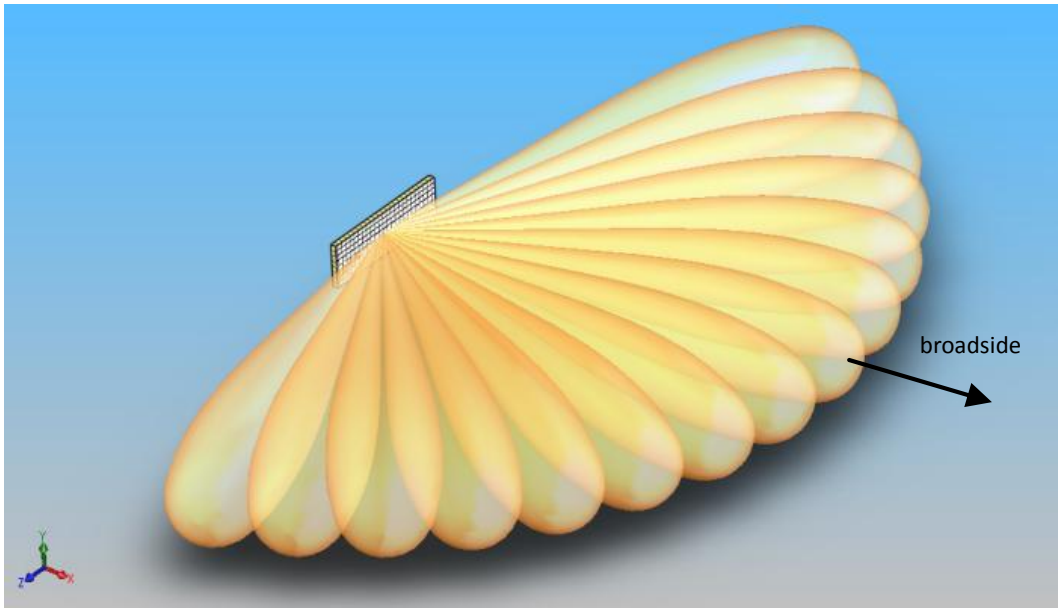

Figure 2.16: A 3-D rendering of a fan of azimuth beams.

(this justifies (2.7)):

$$\frac{S_1(f)}{S_0(f)} = \frac{j \tan\left[\pi f(D/c)(\sin\theta - \sin\theta_s)\right]}{j\pi f(D/c)} \approx \sin\theta - \sin\theta_s \tag{2.9}$$

A beam with a broadband (vertical) null at elevation angle θ_n can be formed through a linear combination of $s_1(t)$ and $s_0(t)$: $x(t) = s_1(t) - (\sin\theta_n \sin\theta_s) \approx \alpha s_0(t)$, where $\alpha = \sin\theta - \sin\theta_n$. Therefore, $s_1(t)$ and $s_0(t)$ can be fed as separate inputs to a linear, multichannel equalizer. The equalizer is then allowed to form this same linear combination (if it wants to), thereby nulling broadband interference at a dynamically varying θ_n. The theoretical wideband vertical beampatterns of a six-element vertical line array with $\lambda/2$ spacing and quality factor $Q = 4$ are shown in Figure 2.17. In this figure, $s_0(t)$ and $s_1(t)$ are linearly combined to steer a wideband null to $\theta_n = 10^\circ$ (i.e. $x(t) \approx s_1(t) - 0.17 s_0(t)$). This process can be extended to form additional monopulse pairs with the same relationship. For example, one can form three beam outputs, $s_0(t)$, $s_1(t)$, and $s_2(t)$ such that this linear relationship,

$$s_2(t) \approx (\sin\theta - \sin\theta_s) s_1(t) \tag{2.10}$$

applies in addition to (2.7).

Introducing this processing in ACOMMS for large array receivers serves three purposes. First, spatial dimension allows filtering of residual reverberation, restoring a large amount of SINR over simpler receivers. Second, reduction in baseband channel count allows for reduced computational complexity while still harvesting gains from spatial degrees of freedom and reverberation-suppression capabilities. Third, reduced channel count allows time-adaptive

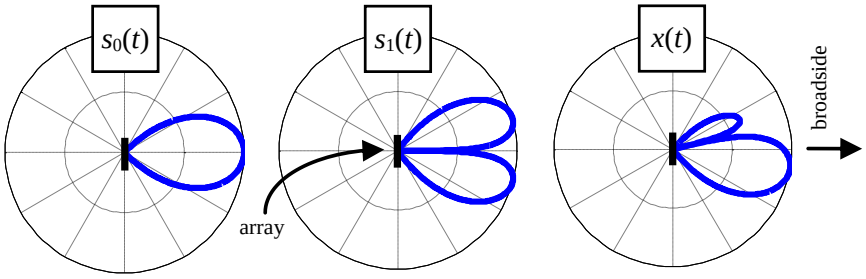

Figure 2.17: Broadband theoretical beam patterns showing angular sensitivity of a six-element vertical line array (heavy black) with $\lambda/2$ spacing and quality factor $Q = 4$. From left to right—$s_0(t)$, $s_1(t)$, and their linear combination, $x(t)$ for $\theta_n = 10°$ from broadside—are plotted in log scale (20 dB/div).

processing to operate with faster adaptation rates while improving performance of channel and Doppler estimates presented in Section 2.3.1. Increased adaptation rates allow for increased Doppler spread tolerance and/or larger, stable feedback structures.

The proposed system shown in Figure 2.18 is centered around a large hydrophone array. Front-end processing consists of substantial dimensionality reduction along the spatial dimension to a subset of beams pairs. Since the signal is highly specular in azimuth (horizontal) as shown in Figure 2.3, dimensionality can be immediately reduced using a specular matched filter—i.e. a delay-weight-sum beamformer designed for the horizontal aperture. This provides array gain to the LOS and vertical channel (improving its signal-to-ambient noise) and rejects off-azimuth angle incoherent surface scatterers (improving its signal-to-reverberation). The second stage of the front-end processing is the wideband monopulse spatial filter along the vertical dimen-

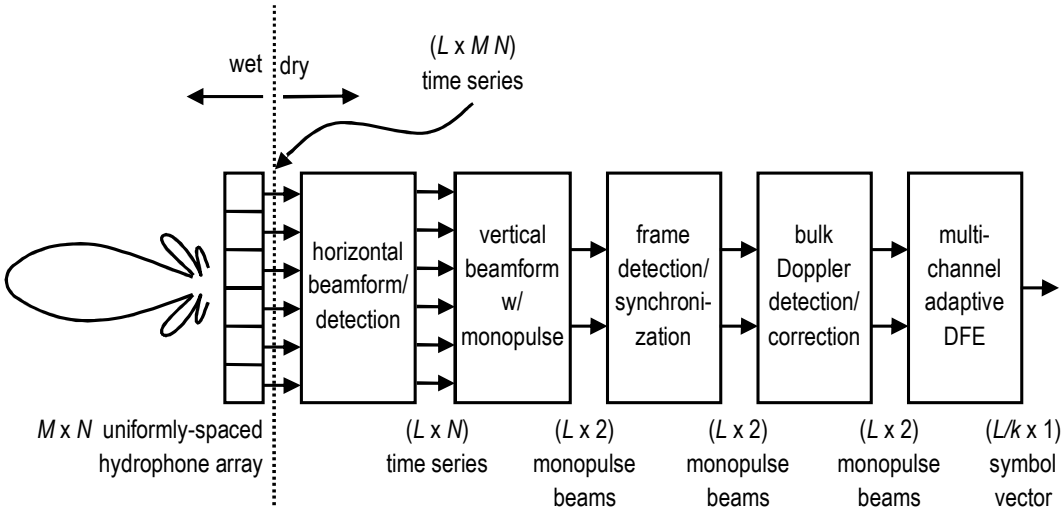

Figure 2.18: System block diagram of the wideband monopulse ACOMMS receiver. Samples from a $M \times N$ uniformly-spaced hydrophone array are horizontally processed using a beamformer. The array is then steered toward the most coherent azimuth direction and vertically processed to yield monopulse beam pair $s_0(t)$ and $s_1(t)$. Channel delay and Doppler distortion are compensated using adaptive tracking, and both channels are recombined using a multichannel adaptive DFE.

sion [19]. These beam outputs are used by the receiver to adaptively steer a wideband null in the elevation angle denoted θ_n using the linear relationship,

$$x(t) = s_1(t) - (\sin\theta_n - \sin\theta_s)s_0(t) \text{ for time } t. \tag{2.11}$$

This can be applied independently along different taps of the equalizer, forming potentially independent spatial nulls at each tap of the multichannel DFE as directed by the output performance of the equalizer decisions. This process can be extended to a number of vertical beam pairs, allowing multiple wideband spatial nulls to be formed. Up to two pairs (three beam outputs) were examined in simulation and experiment and compared against processing raw element outputs using a larger, more complex DFE spanning all channels.

To evaluate the performance of the monopulse ACOMMS receiver, a physical channel simulation based on (2.1) was constructed. For a narrow azimuthal beam, the shallow ACOMMS channel can be simplified into a 2-D vertical plane of water containing the receiver and the source. This plane is intersected by two horizontal boundaries (the surface and the bottom). Simplifying the physics allows the path parameters in (2.2, 2.3, 2.4) to be obtained. The r^{th} reflection coefficient can be calculated using the Rayleigh formula [35]:

$$\Gamma_r = \frac{z_{1,3}/z_2 - \sqrt{1-(n-1)\tan^2\phi_i}}{z_{1,3}/z_2 + \sqrt{1-(n-1)\tan^2\phi_i}} \tag{2.12}$$

Here, $z_{1,3}$ is the acoustic impedance of the surface or bottom media depending on which boundary the ray encounters (e.g. air or wet sand), z_2 is the acoustic impedance of the water, is the incident angle with respect to the normal to the boundary, and $n = (c_{1,3}/c_2)^2$, where $c_{1,3}$ is the sound propagation speed of the surface or bottom media and c_2 is the sound speed in water. Boundary parameters for the simulation are shown in Table 2.6. For this case, Γ_r is always real. Typical reflection coefficients using these values for the surface and bottom are about -1 and 0.2 at normal incidence, respectively.

The math can further be simplified by neglecting refractive effects within the propagation media by assuming it is isotropic (i.e. constant sound speed), as it is in Figure 2.9. This allows boundary incident angles to remain constant. Using this generalization, the point source transmitter is imaged by reflecting its position with respect to the boundaries N_r times for a path with N_r reflections, allowing $\vec{p}_i$ to be readily calculated. The maximum number

Table 2.6: Physical channel simulation specifications.

parameter	value(s) used
ranges (m)	25:25:700
transmitter depth (m)	10
receiver depth (m)	5
transmitter closing speed	**sim1 = 0/sim2 = 5**
sound speed in medium 1, c_1 (m/s)	343
acoustic impedance of medium 1, z_1 (Ns/m^3)	415
sound speed in medium 2, c_2 (m/s)	1500
acoustic impedance of medium 1, z_1 (Ns/m^3)	1,480,000
sound speed in medium 1, c_1 (m/s)	1500
acoustic impedance of medium 1, z_1 (Ns/m^3)	2,324,000 [35]
maximum number of reflections, N_{r-max}	1
ambient noise power, N_0 (dB re: direct path)	-20 at 700 m
DFE adaptation rate μ	**sim1 = 0, sim2 = 0.01**
number of (1/2)-spaced forward taps per channel	11
number of (1)-spaced decision feedback taps	5

of reflections, N_{r-max}, is a simulation input. The i^{th} delay, Δ_i, and Doppler factor, d_i, follow from $\vec{p}_i$ as described in Section 2.2.1. The propagation paths for a 75 m isotropic channel with $N_{r-max} = 5$ are shown in Figure 2.2.

Using this simulation, I model two challenging multipath scenarios:

1. Stationary transmitter broadcasting to a 6-element line receiver of the type discussed in Section 2.3.2. Range is varied from 25 to 700 m in 25 m steps. This simulation demonstrates that, using monopulse processing, the multichannel equalizer can control the linear combination of the monopulse pairs to suppress incoherent multipath within its time span.

2. Same as simulation 1 but with a roving transmitter travelling at modest speeds of 5 m/s. This simulation demonstrates that monopulse processing can be used to suppress incoherent Doppler-spread paths arriving within EQ coherence time but outside its coherence frequency.

2.3.2.1 Simulation 1 Stationary Transmitter and Receiver

For this simulation, the transmitter is at 5 m depth and the receiver line array is centered at 10 m depth. Both platforms are stationary, but the distance between the platforms is varied from 25 to 700 m in 25 m steps. The values for the remaining parameters are shown in Table 2.6. Four receiver configurations were used for the test: (i) 1-ch equalizer processing $s_0(t)$, (ii) 2-ch equalizer processing $s_0(t)$ and $s_1(t)$, (iii) 3-ch equalizer processing $s_0(t)$, $s_1(t)$, and $s_2(t)$, and (iv) a 6-ch equalizer operating on the element outputs. Equalizer OSNR was recorded for each run. Results are shown in Figure 2.19.

As seen in the figure, equalizer OSNR for all four configurations varies with range due to the different delays and strengths of multipath echoes of each geometric environment. Using the $s_0(t)$ output (a ~20° wide beam aimed at $\theta_s = 0°$), the surface and bottom echoes strongly interfere with the direct path at certain ranges, particularly at 100-200 m, 325 m, and 425 m. Using this configuration, the equalizer is only able to suppress ISI when the delay lies within the time-span of the forward or feedback filters. The longest delay is roughly a few hundred symbol periods observed at the closest range and the smallest delay is roughly a few symbol periods at the furthest range.

Using the wideband null made available from preprocessing the array, the second configuration is able to suppress arrivals from a given direction, resulting in substantially better performance than using $s_0(t)$ alone. At 25-150 m, the equalizer steers the null toward the surface. Beyond 150 m, the surface echo arrives within the visible span of the equalizer and is able to be

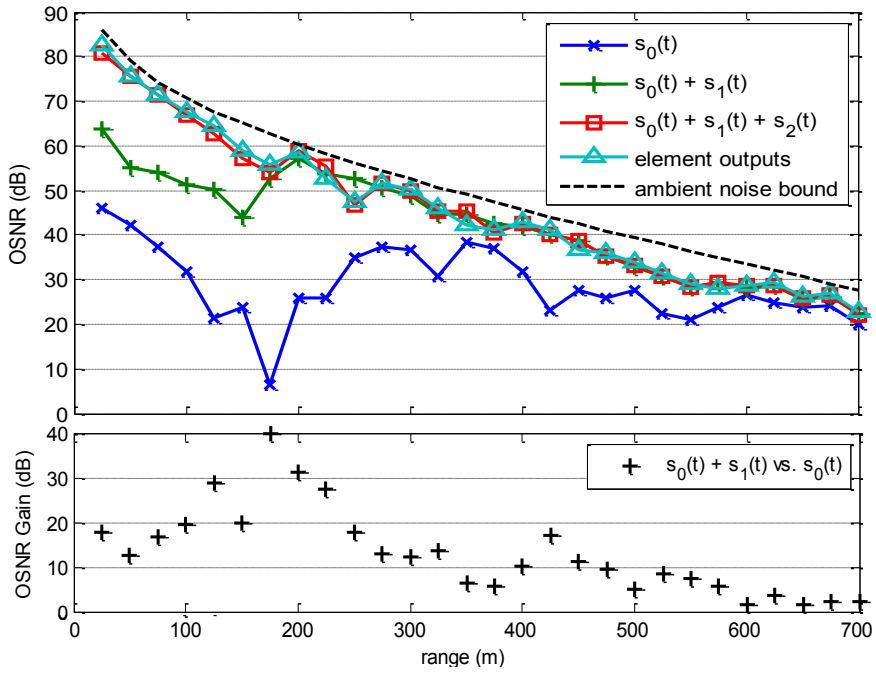

Figure 2.19: Equalizer output estimated signal to noise ratio (OSNR) vs. range for communication in a no Doppler physical channel model. The ambient noise bound (27.8 dB at 700 m) is denoted by the heavy dashed line. Monopulse channels $s_1(t)$ and $s_2(t)$ allow for nulls to be steered at interfering paths outside the coherent time-frequency span of the DFE. The 3-output monopulse (2 free nulls) performs nearly as well (and better at some ranges) as the 6-ch equalizer at 1/8th the computation cost.

recombined by the filter taps. The primary interferer beyond this range is the bottom echo. If the array is processed more extensively to yield $s_2(t)$ of (2.10), another free null is added. This allows the equalizer to null both the surface and bottom echoes at range < 150 m, reducing ISI and allowing the receiver to approach the ambient noise limit. Allowing the equalizer to operate directly on the element outputs results in a marginal increase in signal quality over the third configuration ($s_0(t) + s_1(t) + s_2(t)$) for most ranges (though not every) at a large increase in complexity—i.e. doubling the size of the equalizer increases the operations required to invert the coefficient initialization matrix by 8 ($O(N^3)$) and doubles the operations in the LMS adaptive stage.

2.3.2.2 Simulation 2 Roving Transmitter and Stationary Receiver

This simulation follows the same configuration as Scenario 1 except that the transmitter is now in motion at a closing speed of 5 m/s and the equalizer performs decision-directed adaptation over the course of the packet (with adaptation factor $\mu = 0.01$). As in Scenario 1, the increased number of monopulse channels allows for more rejection of incoherent multipath. In this case, multipath is incoherent due to Doppler differences with the direct path while using only a single compensation factor. The difference between the OSNR and the ambient noise bound decreases with range for most configurations since the difference between the Doppler distortion in the direct path and the multipath is reduced. This is most apparent in the similarity between the furthest ranges in Figure 2.20 and Figure 2.19, which are nearly identical.

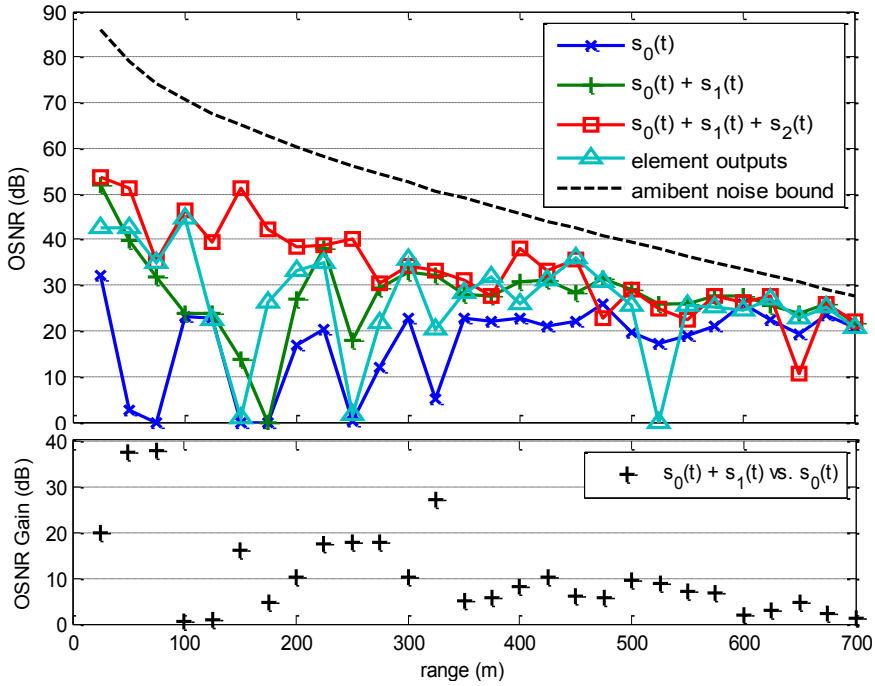

Figure 2.20: In this case, Doppler spread is severe, reducing the coherency of paths even when they lie within the time-span of the equalizer. Since the equalizer cannot track and coherently recombine highly Doppler-spread paths, it is most beneficial to suppress them with spatial filtering, increasing OSNR.

Another observation is that as range increases, the delay between the direct path and the surface and bottom reflection(s) is reduced, though their strength increases as incident angle approaches grazing. Using the linear combination of $s_0(t)$ and $s_1(t)$, the surface reflection is suppressed. Further, adding the $s_2(t)$ beam allows for a linear combination to suppress both interferers, reducing ISI and improving OSNR.

2.3.2.3 Results from a Large Array in Shallow Water

To validate simulated results, several data collection campaigns were conducted to evaluate the proposed space-time filtering methods [17–20]. Figure 2.8 shows an overhead view of the Lake Travis Test Station environment, showing elevations above sea level for underwater features as discussed in Section 2.3.1. Tests were conducted using many-element arrays of different geometries to analyze horizontal and vertical spatial processing methods. Tests were conducted under similar conditions of those presented in Section 2.3.1 with the large array mounted at the barge (Position 1).

Figure 2.21 shows the SINR gain of adding a monopulse pair ($s_0(t)$ + $s_1(t)$) vs. $s_0(t)$ for a six-element array of similar size to that simulated in Figure 2.19 and Figure 2.19. Each black plus corresponds to a decoded packet. The most apparent gains are seen for the 26 packets recorded at the 75 m range (Experiment #31-56). Gains of a few dB are seen throughout other ranges with some variability. Depths were gradually varied from 0-10 m then from 10-0 m at each range. Gains below zero dB are attributed to doubling

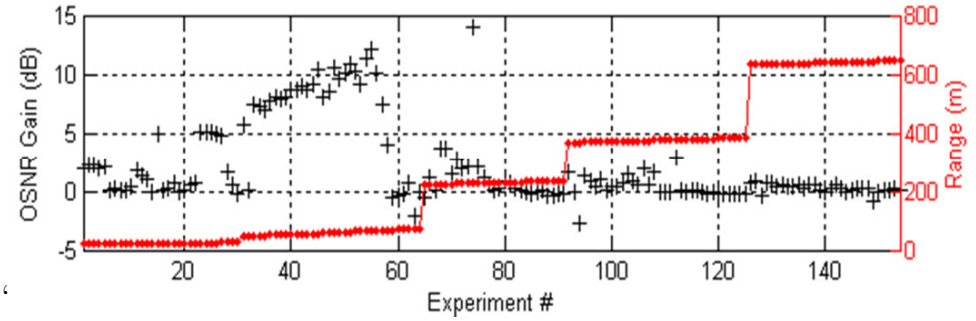

Figure 2.21: OSNR gain vs. range for monopulse ACOMMS receiver in a shallow water lake. Gain is given with respect to the OSNR of the $s_0(t)$ beam.

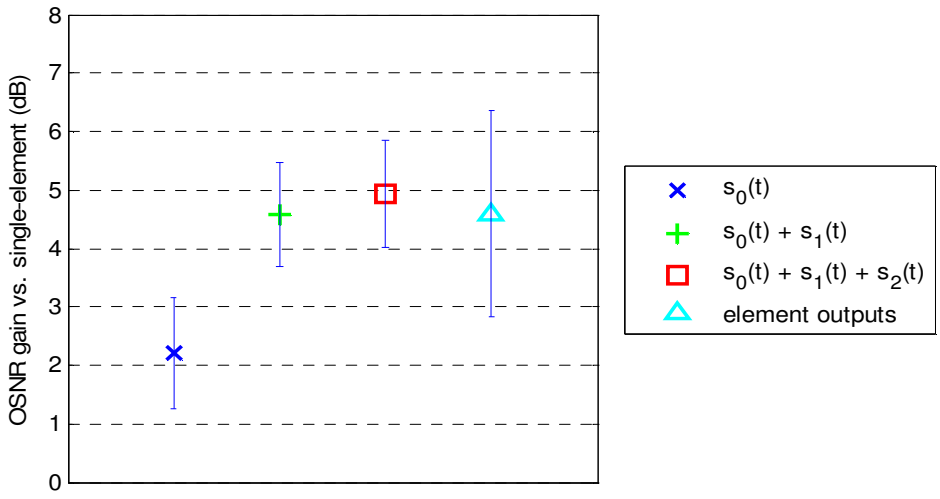

Figure 2.22: Spatial processing gains for an 8-element vertical array vs. a single element receiver (96 packets at various rates, Q = 0.8-3.1)

the number of equalizer taps while keeping the adaptation factor constant in the presence of uncompensated Doppler. It is apparent that observed gains are highly dependent on range and depth. This is in agreement with the simulations of Section 2.3.2.1 and 2.3.2.2. Figure 2.22 shows the mean SINR gain of an 8-element tall array vs. a single-element receiver. Several vertical processing methods were compared. Only packets with BER < 0.05 for all five methods were included. It is apparent in these mixed-condition experiments that the monopulse ACOMMS receiver provides gains over processing all channels with reduced computational complexity.

Additional capacity experiments were conducted in [20]. Large gains in SINR were realized for the horizontal processing which suppresses azimuthally-spread reverberation and provides LOS directionality gains. In one experiment, a subset of the fan of azimuth beams was processed to increase the capacity of the system. In this configuration, eight uncoordinated users communicated simultaneously with the large array. Seven simultaneous transmissions were decoded, achieving a sum spectral efficiency of 28 bps/Hz, nearly an order of magnitude of what previous state-of-the-art had demonstrated.

Though these are good initial results, care was taken to ensure that transmit powers from each transmission were received with roughly equal power by positioning them at similar ranges. This is a challenge in deployment and more sophisticated, adaptive schemes such as uplink power control must be used to ensure that receive power is approximately equal over all users.

Table 2.7: New empirical results using wideband monopulse system.

Method	Number of Elements/ Array Geometry	Center Frequency (kHz)	Range (km)	Rate (kbps)	Bound (kbps)	Sum-Rate Capacity (bps/Hz)
Prior State-of-the-Art [24]	8 vertical receive, 2 vertical transmit	17	1-3	32	13.3	2.3
Monopulse + SDMA [20]	2-D w/ hundreds, 7 simultaneous users	—	—	1400	—	28

2.4 Conclusion

In this chapter, I have presented two methods for increased underwater acoustic communication performance in the presence of strong reverberation (self-interference) typical of shallow water channels. The first method is a windowed Doppler compensation technique coupled with a time-adaptive decision-feedback equalizer. The second method is a space-time filtering method coupled with a decision-feedback structure to spatially filter the communication channel. This method is used to enhance coherent (*desirable*) paths while suppressing incoherent (*undesirable*) paths. These methods are combined to form a robust underwater acoustic communication receiver for shallow water channels that is scalable to large arrays.

Results from a shallow water test environment are presented. Experimental and simulated results agree, demonstrating reduction in interference in excess of 10 dB at the equalizer output. In addition to performance gains, computational complexity is reduced by using the smaller number of pre-processed array outputs of a many-element receiver. Finally, this method is used in conjunction with a multi-user scenario to demonstrate gains in sum spectral efficiencies an order of magnitude above prior state-of-the-art.

Chapter 3

Time-Frequency Methods for Powerline Communications

3.1 Motivation

Powerline communication (PLC) systems operate by coupling modulated, high-frequency signals onto wires and transmission lines used for electrical power delivery. Operating under these conditions subjects these signals to strong impulses and transient disturbances caused by the multitude of noisy devices connected to a power grid. Two primary components of this noise are: (i) cyclostationary noise with time-periodic statistics synchronous to half of the AC cycle and (ii) bursty asynchronous impulsive noise. Existing powerline standards only cursorily address these impairments through built-in mechanisms defined in G3-PLC, PRIME, IEEE P1901.2, and ITU G.9901-9904 [1, 2, 5, 6]. In this work, I first measure power line noise and investigate its time-frequency properties. I use these results to develop a cyclic adaptive coding and modulation method. With this method, I demonstrate up to 25 dB operating point improvements for the G3-PLC standard. Second, I implement a method to mitigate impulsive noise in OFDM systems. This method recovers up to 8 dB SNR of impulsive noise in real-time. Both methods can be applied independently to improve PLC system performance.

Table 3.1: Narrowband PLC bands by region [37].

Region	Band	Frequency (kHz)
US	FCC	1 – 100
Europe	CENELEC-A	3 – 95
	CENELEC-B	95 – 125
	CENELEC-C	125 – 140
	CENELEC-D	140 – 148.5
China	EPRI	3–90
Japan	ARIB	10–450

3.2 Background

Powerline communications (PLC) is a technique that enables digital communication over powerlines designed for electrical power delivery. The availability of large amounts of existing infrastructure, particularly in developed countries, makes PLC systems attractive for consumer and commercial use. Narrowband PLC (NB-PLC) systems enable communications up to 1 Mbps over distances up to several km using signals in the 3-500 kHz band (see Table 3.1). Typical uses of NB-PLC include distributed load balancing, smart meter reading, fault monitoring and detection, etc. [36]. NB-PLC communication networks can be formed between homes, offices, smart devices, and local utility substations. Utility substations are used as aggregation points and the interface to other data networks via wired or wireless backhaul.

Modern electric grids were designed for power delivery at the AC mains frequency, making them a less than ideal medium for digital communications. Noise present in NB-PLC bands can be attributed to switching loads and generation sources attached to the grid. Sources of synchronous noise can include semiconductor controlled rectifier (SCR)-based light dimmers and motor con-

trollers that power loads on and off at different phases of the AC cycle. Other switching circuits such as DC-DC converters can contribute asynchronous but periodic noise into power lines. The authors of [38] analyze the broadband noise emissions of different loads in an anechoic chamber using a 100 MS/s sampling system. They observe strong cyclic noise properties, particularly for loads with high current AC motors and/or SCR switching devices.

Figure 3.1 shows the time-frequency decomposition of powerline noise measured in an apartment building in central Austin, TX. This figure captures noise typical in NB-PLC deployment scenarios. As seen in the figure, noise in NB-PLC is largely spectrally- and time-varying, with individual component powers rising as much as 80 dB above background noise levels. Time-periodicity with the dominant period of one-half the AC mains cycle—i.e. 8.3 ms for 60 Hz (predominantly North and South America) and 10 ms for 50 Hz (predominantly Africa, Asia, Australia, and Europe)—is clearly seen as noted in prior work [39]. In addition, aperiodic impulsive components are observed.

3.2.1 NB-PLC Noise Sources

Additive powerline noise is generated by electrical devices connected to the power grid and by external noise and interference coupled to the power network via radiation and/or conduction. In the 3–500 kHz band, the powerline noise can generally be decomposed into four classes [41]:

1. *Spectrally-shaped background noise*. The background noise is the summation of numerous low-power noise sources. Its power spectral density

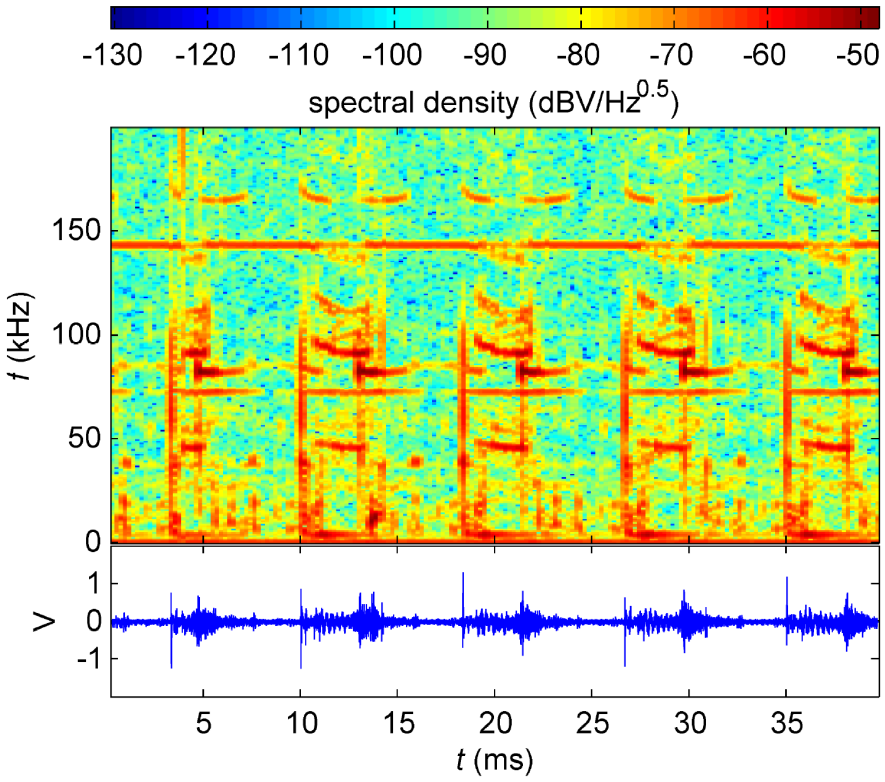

Figure 3.1: Noise power spectral density measured using Freescale NB-OFDM PLC modem at a residential apartment building in central Austin, TX. Note: (i) variations of noise spectral power over bands of interest (CENELEC-A 35.9-90.6 kHz, FCC 3-95 kHz) and (ii) spectral structure is periodic with period equal to one half the AC mains cycle (8.3 ms in the US) [40].

slowly varies over time (in minutes or hours), and exhibits a $1/f$-type decay due to the decreasing concentration of noise sources with frequency and the lowpass response of coupling devices such as transformers.

2. *Narrowband interference*. Broadcast stations in the long-wave bands (153–279 kHz) introduce narrowband interference in NB-PLC systems. The interference exhibits amplitude or frequency modulated sinusoids in the time domain, which correspond to harmonic clusters in the noise spectrum. The interference level generally varies slowly with daytime, and in some cases varies periodically with half the AC cycle, i.e., is cyclostationary. Fluorescent bulb ballasts are known narrowband sources.

3. *Cyclostationary noise*. The cyclostationary noise exhibits cyclostationarity in both the time and frequency domains, with time period equal to half of one AC cycle [42]. It typically appears in one of the two forms:

   (a) *Periodic impulsive noise synchronous to the mains frequency*. This type of noise consists of a series of isolated impulses of considerable duration and amplitude. The impulses have a repetition rate of 60/120 Hz (in the US) and always appear at the same instant of the AC cycle. They are typically caused by nonlinear power electronic devices, such as silicon controlled rectifiers and diodes, that switch on and off at a particular AC cycle phase.

   (b) *Periodic impulsive noise asynchronous to the mains frequency*. This noise component takes the form of impulse trains with repetition

rates unrelated to the mains frequency [43]. In addition to the high repetition frequencies, it can exhibit an underlying period equal to half the mains cycle. It is primarily injected by switching mode power supplies that operate at frequencies above 20 kHz. The impulses in such noise typically have much lower amplitudes and shorter durations than those in the synchronous impulsive noise, and hence are usually masked by other high-level noise components in the time domain. In the frequency domain, however, such noise can be easily identified as evenly spaced harmonic clusters with significantly higher power spectral density over the background noise.

4. *Asynchronous impulsive noise*. This type of noise consists of impulses of significant duration and amplitude. It typically arises from switching transients caused by the connection and disconnection of electrical devices and/or rapid load impedance changes at random times.

Recent field measurements on outdoor medium-voltage and low-voltage power lines have shown that cyclostationary noise, including periodic impulsive noise synchronous and asynchronous to the mains frequency, is the dominant noise component in the 3–500 kHz band [44]. In addition, extensive measurement campaigns taking place over the last several decades have revealed impulsive statistics, with spurious components sometimes reaching 40 dB above background noise levels [45–47]. A comprehensive summary of the types of powerline noise is provided in [36].

3.2.2 System Model

This work addresses the two dominant classes of NB-PLC noise enumerated above in Section 3.2.1 as classes 3 and 4: (*i*) cyclostationary noise and (*ii*) asynchronous impulsive noise. Cyclostationary noise such as that observed in Figure 3.1 can be modeled using a class of random processes known as wide-sense cyclostationary (WSCS) processes [48]. WSCS processes are random processes with periodic second-order statistics—i.e. its samples exhibit a power spectral density that is periodic with period N,

$$S_{xx}(l,\omega) = S_{xx}(l+kN,\omega) = \frac{1}{Nf_s}\left|\sum_{n=0}^{N-1} x_{n-l}e^{-j\omega n}\right|^2 \qquad \forall k \in \mathbb{Z}, \tag{3.1}$$

and, equivalently, an autocorrelation function that is periodic with period N,

$$R_{xx}(l) = R_{xx}(l+kN) = \sum_{n=0}^{N-1} x_n x_{n-l}^* \qquad \forall k \in \mathbb{Z}. \tag{3.2}$$

Here, f_s is the sample rate of the discrete time sequence, ω is its angular frequency, and $\mathbb{Z}$ denotes the set of all integers $(-\infty, \infty)$.

In [39], WSCS noise is modeled using a Gaussian sample distribution with time-periodic variance. This enables analytical solutions for capacity in cyclostationary noise channels. This model is further refined in the IEEE P1901.2 NB-PLC standard, which presents a more general model for WSCS noise using filtered additive, white Gaussian noise (AWGN) [5]. In this model, unit variance AWGN is input to a filter bank with outputs selected during sub-intervals of the cyclostationary period. The filters shape the AWGN input to the noise power spectral density shape within each sub-interval.

Table 3.2: Impulsive noise models for complex vector $\mathbf{X} = \{X^{(I)}, X^{(Q)}\}$ [49].

Model	Distribution characteristics	Notes
Symmetric Alpha Stable	Characteristic function: $\Phi_{X^{(I)},X^{(Q)}}(\omega_I, \omega_Q) = e^{-\sigma\left\|\sqrt{\omega_I^2+\omega_Q^2}\right\|^{\alpha}}$ Parameters: σ, α	Closed-form PDF does not exist except for $\alpha = 1$, $\alpha = 2$
Middleton Class A	PDF: $f_{X^{(I)},X^{(Q)}}(x^{(I)}, x^{(Q)}) = e^{-A}\delta(x^{(I)}, x^{(Q)}) + \sum_{m=1}^{\infty} \frac{e^{-A}A^m}{m!} e^{-\frac{(x^{(I)})^2+(x^{(Q)})^2}{\frac{2m\Omega_{2A}}{A}}}$ Parameters: A, Ω_{2A}	A particular form of the Gaussian mixture distribution
Gaussian Mixture	PDF: $f_{X^{(I)},X^{(Q)}}(x^{(I)}, x^{(Q)}) = p_0\delta(x^{(I)}, x^{(Q)}) + \sum_{l=1}^{L} p_l \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x^{(I)})^2+(x^{(Q)})^2}{2\sigma_l^2}}$ Parameters: p_l, σ_l^2, L	Mixture of Gaussians such that $\sum_{l=1}^{L} p_l = 1$

In addition to its cyclostationary components, NB-PLC channels are also subject to impulsive noise. This impulsive noise can be synchronous or asynchronous to the AC mains cycle arising from switching of devices connected to the powerline. Extensive measurement campaigns taking place over the last several decades have confirmed the presence of impulsive noise and have characterized its statistics [45–47]. Several common statistical models for impulsive noise and their respective probability density functions (PDFs) (if available) are listed in Table 3.2. These models are Symmetric Alpha Stable, Middleton Class A, and Gaussian Mixture. These models have been derived from empirical results and/or stochastic models of noise sources.

The system model for a conventional NB-OFDM PLC system is shown in Figure 3.2. This system uses OFDM signaling to provide robustness to narrowband noise and interference. An OFDM system with cyclic prefix diagonalizes the circulant channel matrix $\mathbf{H}$ using the discrete Fourier transform

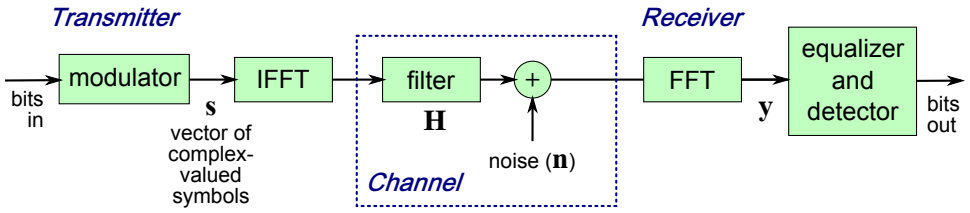

Figure 3.2: NB-OFDM PLC system operating over an additive noise channel. Modulated bits are transmitted through a channel consisting of filter $\mathbf{H}$ and noise vector $\mathbf{n}$. Received samples are processed with an FFT, cyclic prefix is removed, and received symbols are detected and converted to bits.

(DFT) and inverse DFT (IDFT) matrices $\mathbf{F}/\mathbf{F}^H$ (where ^{H} denotes Hermitian transpose):

$$\mathbf{y} = \mathbf{FHF}^H\mathbf{s} + \mathbf{Fn}. \tag{3.3}$$

Here, impulsive components of $\mathbf{n}$ contribute added noise in all subcarriers due to the spreading properties of the DFT.

The additive noise component $\mathbf{n}$ can be modeled as a superposition of three components: (i) background, (ii) cyclostationary, and (iii) impulsive noise such that $\mathbf{n} = \mathbf{b} + \mathbf{c} + \mathbf{x}$. Over a sufficiently short interval such as an OFDM symbol duration, samples can be assumed to be drawn from an *iid* Bernoulli (two-mode, L of Table 3.2 is 2) Gaussian Mixture distribution where γ_X and γ_B denote the impulse and background noise power ($\gamma_X \gg \gamma_B$), and π denotes impulse probability. The PDF can be expressed as:

$$GM(\pi, \gamma_B, \gamma_X) = \pi\mathcal{N}(0, \gamma_B + \gamma_X) + (1 - \pi)\mathcal{N}(0, \gamma_B). \tag{3.4}$$

Here, the $\mathcal{N}(0, \gamma_B)$ component includes contributions from the Gaussian background noise and cyclic noise over a small sub-interval of the AC cycle. The

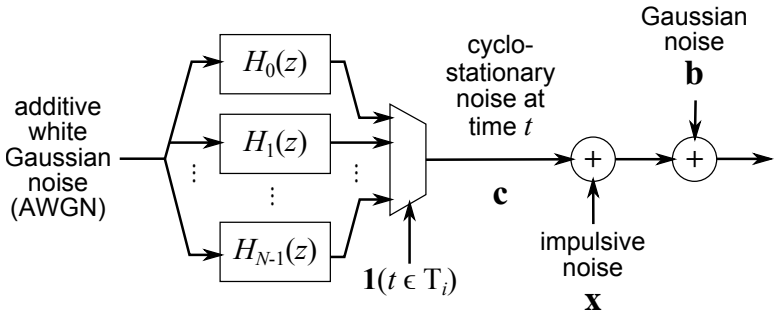

Figure 3.3: NB-PLC Noise model consisting of a superposition $\mathbf{n} = \mathbf{c} + \mathbf{x} + \mathbf{b}$ of cyclic noise $\mathbf{c}$ constructed from selecting noise power spectral density (PSD) $S_{cc,i}(t, \omega)$ in the i^{th} sub-interval T_i of a half AC cycle, asynchronous impulsive noise $\mathbf{x}$ derived from a Bernoulli-Gaussian mxture model, and background Gaussian process $\mathbf{b}$.

two-mode impulsive noise model lends itself to analysis due to its simple form yet retains relevant statistics for analysis. Often, higher mode Gaussian mixture distributions have rapidly diminishing returns, as even complex impulsive PDFs can be approximated by low-order models [50].

The cyclostationary component of the noise can be modeled according to the P1901.2 NB-PLC standard [5] as a linear, periodic time-varying random process. In particular, the cyclostationary noise can be generated by processing unit variance AWGN with a filter bank consisting of filters shaped to the PSD of noise in sub-intervals of the cyclostationary period. The i^{th} filter output is selected when the system time t modulo the cyclostationary period T lies within sub-interval T_i. Sub-intervals T_i are non-overlapping, contiguous, and span $(0, T)$—i.e. $\bigcap_{i=0}^{N-1} T_i = \emptyset$, $\bigcup_{i=0}^{N-1} T_i = (0, T)$. In P1091.2 noise models, the number of sub-intervals N is chosen to be 3. Figure 3.3 shows the complete noise model with background, cyclostationary, and impulsive components.

Table 3.3: Impulsive noise mitigation techniques.

Method	Low SNR	High SNR	Non-Parametric?	Computational Complexity
Nulling/ Clipping [51]		✓		Low
Iterative Decoding for OFDM Transmission [52]	✓	✓		High
Thresholded Least Squares/MMSE [53]		✓		Med
Sparse Bayesian Learning [54]	✓	✓	✓	High (matrix inversion)
l_1-norm Minimization [53]	✓	✓	✓	High
Approximate Message Passing (AMP) [50, 55]	✓	✓		Med

3.2.3 Prior Work

Using an impulsive noise model, communication systems can be designed for robustness and achievable rate in the presence of impulsive noise. These systems can provide vastly improved performance versus conventional systems designed for AWGN channels [51, 53]. Methods for impulsive noise mitigation applicable to NB-PLC include low-SNR techniques (impulse power much higher than background noise) such as nulling and thresholding [51]. More robust, wide-SNR techniques (impulse power may be closer to background noise) based on sparse (many samples near zero) reconstruction demonstrate promising performance but are typically burdened by high computational complexity [53, 56]. In addition to the methods above, other practical methods for robustness in the presence of impulsive noise include block forward error-correcting codes (FEC) and interleaving mechanisms defined in relevant PLC standards [1,2,5,6]. A summary of prior work for impulsive noise mitiga-

tion is provided in Table 3.3. Despite these techniques, throughput achieved in practice is often limited to a small fraction of advertised throughput, particularly in noisy powerline conditions.

The most promising wide-SNR approach is approximate message passing (AMP), which has been formulated as a low computational complexity, scalable technique which lends itself to efficient hardware implementation [57–59]. Recently, Nassar *et al.* proposed an AMP framework for joint impulsive noise mitigation and channel estimation in OFDM receivers [50]. This framework exploits the spreading of impulsive noise information across subcarriers by virtue of the DFT. Frequency-domain noise samples in null subcarriers are used to reconstruct time-domain noise using AMP.

Similarly, methods to combat cyclostationary noise exist in the standards [1,2,5,6]. This includes a time-frequency interleaver and FEC in G3-PLC that is able to correct for erasures of entire frequency bins and/or OFDM symbols each AC cycle using frame-length (many OFDM symbols) block codes. Additionally, a tone-mask procedure can be used to avoid subcarrier groups that may be hit by particularly high noise [1]. Though more effective at combating narrowband noise, this method can be used to avoid band-limited impulses. Other techniques include destroying periodic noise structure using a time-domain interleaver then applying one of the above wide-SNR impulsive noise mitigation methods [60]. The IEEE P1901.2 committee draft standard has adopted a noise model based on WSCS noise [5], though limited methods exist beyond FEC and interleaving to combat the noise.

3.3 Impulsive and Cyclostationary Noise Mitigation

In this work, I propose an enhanced NB-OFDM PLC system designed for robust performance in cyclostationary and impulsive noise channels. Figure 3.4 shows the proposed system model. The additive noise of the conventional system model in Figure 3.2 has been extended to include the comprehensive noise model of Figure 3.3. The proposed system has been modified by adding the new components shown in red. First, a cyclic modulation and coding functionality is added at the transmitter that maps input bits to complex-valued subcarriers. A cyclic de-mapper is added at the receiver to demap complex-valued subcarriers to bits according to the 2-D modulation coding scheme (MCS) map exchanged by the transmitter and receiver. Second, an impulsive noise estimation and mitigation block is added at the receiver to remove impulsive noise before the FFT is taken. The cyclic modulation and coding method and its performance results are discussed in Section 3.3.1. Section 3.3.2 details the implementation of a real-time impulsive noise mitigation block and its measured performance.

For this section, I consider the system parameters listed in Table 3.4. System parameters are identical to the G3-PLC specification [1]. In which, a 256-length real-valued FFT provides 128 complex-valued subcarriers, of which 36 are actively used. Null subcarriers (where zeros are transmitted) are used for impulsive noise detection and estimation. My proposed methods, though examined for this particlar configuration, can easily extend to other system configurations including other current and future PLC standards.

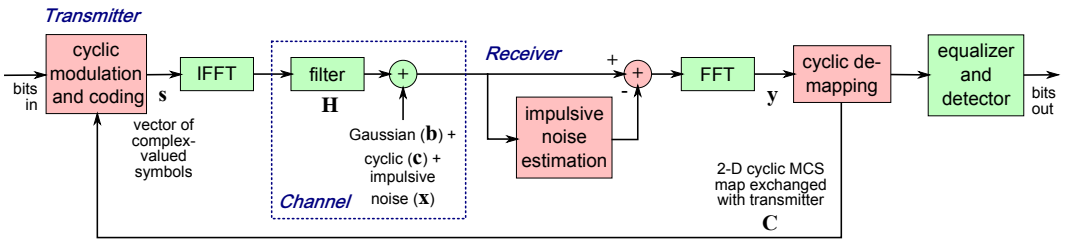

Figure 3.4: NB-OFDM PLC system with a cyclic and impulsive noise channel model. Cyclic adaptive modulation and coding and AMP impulsive noise mitigation are added (red) to increase channel capacity over conventional system (green). A cylic modulation coding scheme (MCS) map is exchanged between receiver and transmitter before the communication session begins.

Table 3.4: G3-PLC system parameters used for time-frequency methods [1].

System parameter	Value(s)
sampling frequency	$f_s = 0.4$ MHz
(real) FFT length	$N_{FFT} = 256$
number of overlapped samples	$N_{OVLP} = 8$
number of cyclic prefix samples	$N_{CP} = 30$
OFDM symbol period	$N_{OFDM} = N_{FFT} + N_{CP} - N_{OVLP} = 278$
available active subcarrier indices (indexed by zero)	23-58
firsts and last subcarrier frequencies	35.9 - 90.6 kHz
null subcarriers (indexing by zero)	0-22, 59-128
null subcarriers used by AMP (indexing by zero)	1-22, 59-100
OFDM symbols per half 60 Hz AC cycle	$L = \lceil 11.99041 \rceil = 12$
OFDM frame length	~15-190 ms
available modulation types (non-coherent mode)	ROBO6, ROBO4, DBPSK, DQPSK, D8PSK
available modulation types (coherent mode)	ROBO6, ROBO4, BPSK, QPSK, 8PSK

3.3.1 Cyclic Adaptive Modulation and Coding

To characterize the cyclostationary component of NB-PLC noise, I conducted a measurement campaign in central Austin, TX. Samples of powerline noise were collected at various locations on The University of Texas at Austin (UT-Austin) campus and surrounding area. Figure 3.5 shows a map of the locations where measurements were taken. Samples were collected using an NB-OFDM PLC modem manufactured by Freescale Semiconductor. The modem was configured to capture 16000 16-bit samples from the powerline at 400 kS/s. These samples were then equalized in the frequency domain using coefficients derived from a calibration procedure via the zero-forcing solution to a received 16000 length chirp sequence. Samples were post-processed using a class of analysis methods known as *cyclic spectral analysis* [61, 62].

3.3.1.1 Cyclic Spectral Analysis

A random process X can be characterized by first defining its *symmetric instantaneous auto-correlation function* to be

$$R_{XX}[n,l] = \mathbb{E}\left\{ X\left[n+\frac{l}{2}\right] X^*\left[n-\frac{l}{2}\right]\right\}. \tag{3.5}$$

Periodicity of $R_{XX}[n,l]$ implies that (3.5) can be expressed as a Fourier series

$$R_{XX}[n,l] = \sum_{\alpha_i \in A} R_{XX}\left[l,\alpha_i\right] e^{j2\pi\alpha_i n\Delta}, \tag{3.6}$$

where $A = \{\alpha_i\}$ is the set of cyclic frequencies α_i. By the additional application of a Fourier transform to the time-dimension n of $R_{XX}[n,l]$, this cyclic

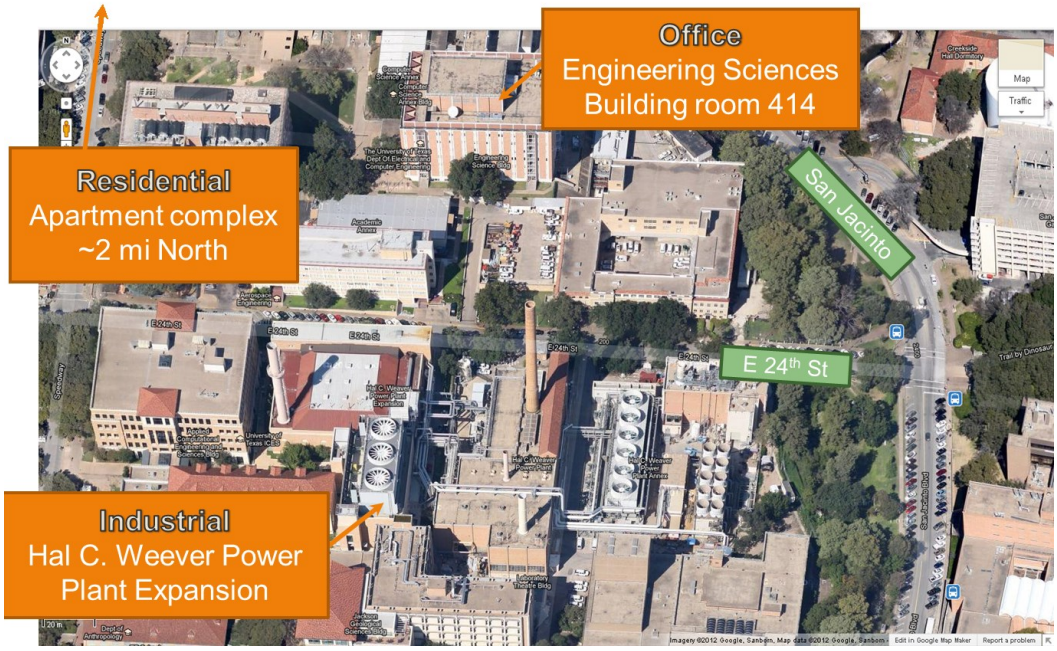

Figure 3.5: Overhead view of The University of Texas at Austin campus and locations where noise measurements were taken.

auto-correlation function can be transformed to yield a 2-D function of cyclic frequency α and frequency f:

$$S(\alpha, f) = \Delta^2 \sum_{n=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} R_{XX}[n, l]\, e^{-j2\pi\alpha n\Delta} e^{-j2\pi f l\Delta}. \tag{3.7}$$

Finally, the symmetric cyclic coherence function is defined as:

$$C_{YX}(f;\alpha) = \frac{\mathbb{E}\left\{\mathrm{d}Y\left(f+\frac{\alpha}{2}\right)\mathrm{d}X\left(f-\frac{\alpha}{2}\right)\right\}}{\sqrt{\mathbb{E}\left\{\left|\mathrm{d}Y\left(f+\frac{\alpha}{2}\right)\right|^2\right\}\mathbb{E}\left\{\left|\mathrm{d}X\left(f-\frac{\alpha}{2}\right)\right|^2\right\}}}. \tag{3.8}$$

Further details about this method and its use can be found in [61].

To confirm the cyclostationary structure of the noise, samples collected using the Freescale modem were then analyzed using discrete-time implemen-

tations of (3.7) and (3.8) in MATLAB. The cyclic frequencies examined are bins 1 through 150 (indexing by zero) corresponding to cyclic frequencies of 25 Hz to 3750 Hz with a 25 Hz step size.

Figure 3.6 shows the cyclic spectral density and the cyclic coherence of noise samples collected in the living room of an apartment in central Austin, TX. The apartment is part of a two-story building with 16 units and is lit using compact fluorescent bulbs. The time-domain samples and spectrogram of this noise is shown in Figure 3.1. The cyclic coherence in Figure 3.6 reveals periodic ridges along the α axis. The majority of noise power is located below 10 kHz and 90 to 100 kHz. The cyclic spectral density is relatively constant as a function of cycle frequency α with a slight decay with increasing α.

The cyclic spectral coherence of Figure 3.6 confirms that the noise at frequencies 4, 45, 58-70, 75, 80-85, 90-120, and 133-140 kHz are all periodic with twice the AC frequency of 60 Hz as indicated by the peak ridge along the $\alpha = 120$ Hz line. A noise source at 166 kHz is found to be periodic at four times the AC frequency. The large asynchronous impulse at $t = 14$ ms with spectral support from 10-14 kHz produces a long streak in the coherence over all α, revealing that it is not periodic and was a one-time event with respect to the system time. Further analysis of noise collected at this location and other locations of Figure 3.5 are provided in [40].

The system settings of Table 3.4 reveal that G3-PLC OFDM frames can span many AC cycles, subjecting transmitted signals to many cyclically coherent instances of the cyclostationary noise. The large spectral and tempo-

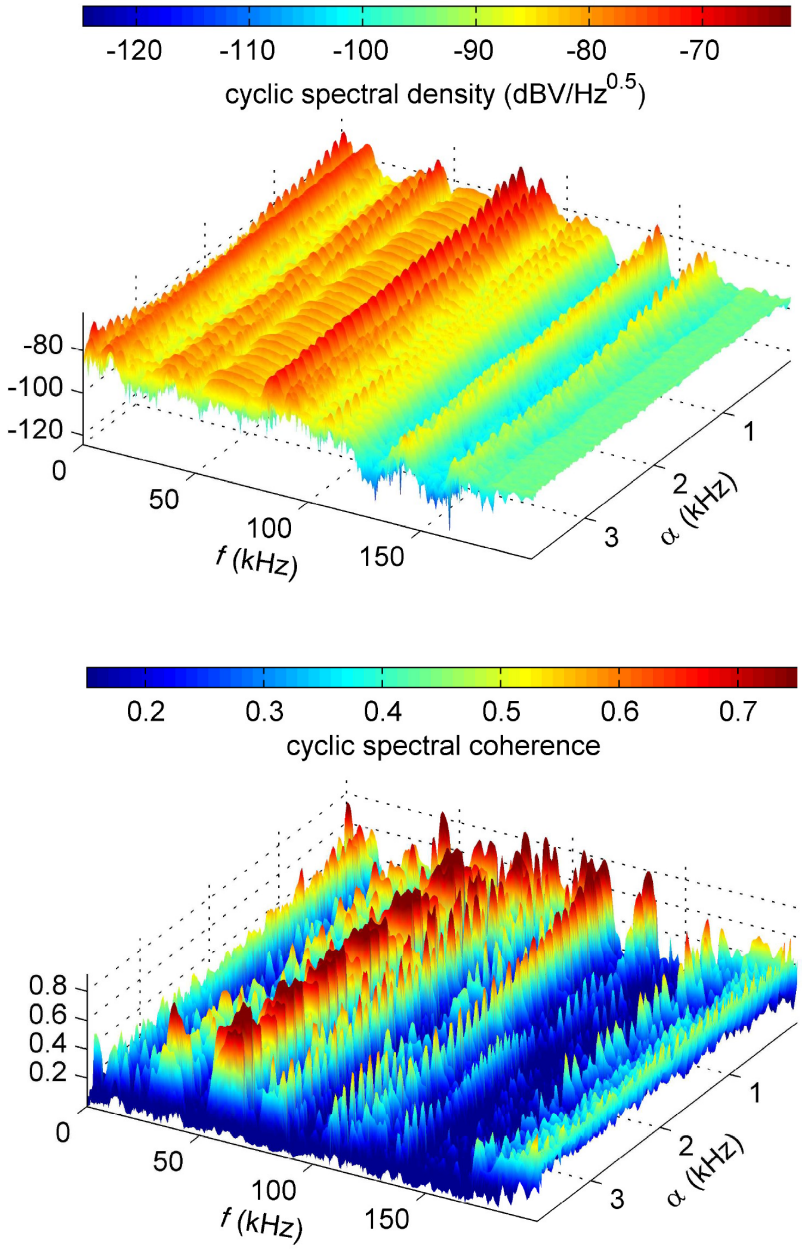

Figure 3.6: Cyclic spectral density (top) and cylic spectral coherence (bottom) of residential powerline noise in Figure 3.1. Despite the highly complex spectral structure, many spectral components exhibit high coherence along the $\alpha = 120$ Hz cycle frequency. Spikes occur at harmonics of this α since an impulse train is an eigenfunction of the Fourier transform.

ral variations of noise power over an AC cycle and the cyclic nature of the noise imply that a cyclic adaptive modulation and coding scheme could be used to improve system throughput. G3-PLC currently supports static modulation with tone mask over an OFDM frame but only allows for a fixed group of six subcarriers to be masked over the duration of the OFDM frame. Adaptive bit loading has been suggested previously for PLC [63] and could be tailored to cyclic noise shape and coherence times discussed above.

I extend these mechanisms to a cyclic adaptive modulation and coding scheme (MCS). For analysis, I consider the G3-PLC system listed in Table 3.4. This method is supported using an *expanded* tone map message that specifies a cyclic MCS map for the transmitter. This exchange is depicted as the feedback signal in Figure 3.4. For G3-PLC in North America, this *expanded* tone map message is a 36×12 (subcarrier $\times$ symbol) field of three-bit entries that index the MCS to use for half the AC cycle (~12 symbols for 60 Hz and ~10 for 50 Hz). This 2-D cyclic MCS map is exchanged during the link negotiation procedure defined in [1] using estimates of the SNR per subcarrier over the cyclostationary period. Once this map is exchanged, the transmitter modulates data, and the receiver de-maps subcarriers according its values throughout the communication session. Cycle phase is synchronized between the transmitter and the receiver using a zero-crossing detector for AC cycle phase that is built into G3-PLC and P1901.2 [1, 5]. An example cyclic map for MCS choices in G3-PLC is shown in Figure 3.7. The map can be interpolated for non-integer L using invertibility properties of the STFT as suggested in [64].

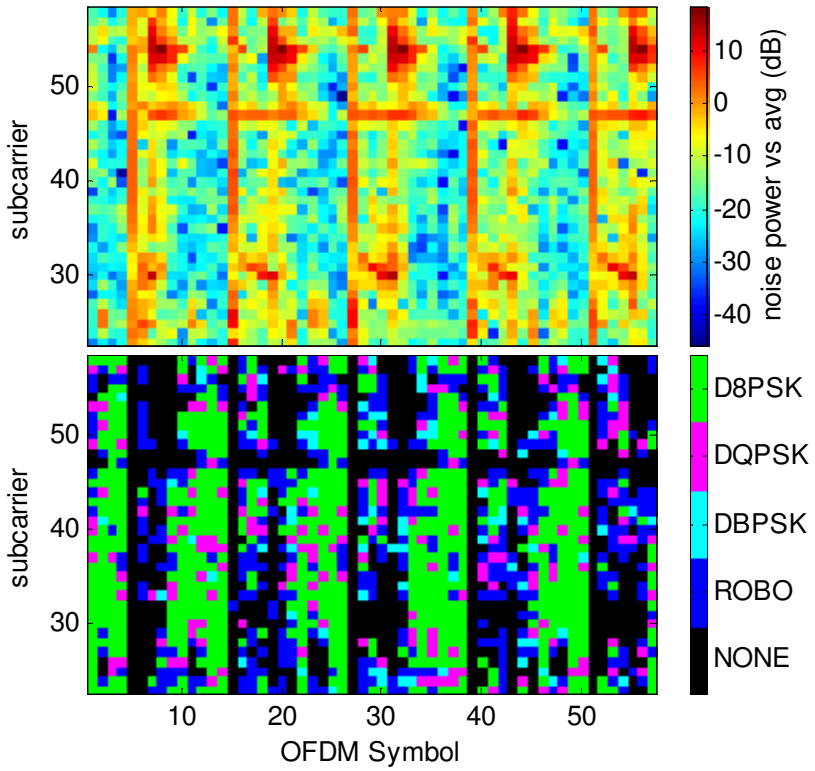

Figure 3.7: (top) In-band noise power for the channel shown in Figure 3.1 and (bottom) modulation coding scheme (MCS) map using the specifications of the G3-PLC standard listed in Table 3.4. The MCS map of dimensions $K \times L$ spans the K active subcarriers and the L OFDM symbols that span a cycle of the WSCS noise [40]. Here, $K = 36$ and $L = 12$.

For this analysis, I assume the noise measured is inversely proportional to the SNR—i.e. the channel is unit magnitude for all subcarriers. I assume a target bit error rate (BER) $p_b = 10^{-2}$ is used to make MCS choices according to the constrained rate-optimal MCS map $\mathbf{C}^\star$ solving

$$\mathbf{C}^\star = \arg\max_{\mathbf{C}\in\mathcal{C}} R(\mathbf{C})\mathbf{1}\left(Q(\mathbf{C},\mathbf{S}) \leq p_b\right). \tag{3.9}$$

Here, $R(\mathbf{C})$ is the sum rate assuming MCS selection $\mathbf{C}$, and $Q(\mathbf{C},\mathbf{S})$ estimates theoretical BER with $\mathbf{C}$ and $\mathbf{S}$, the 2-D time-frequency SNR over a cyclostationary period T. $\mathcal{C}$ is the space of all possible MCS choices. $Q(\mathbf{C},\mathbf{S})$ can be derived analytically for standard modulation types using nearest-neighbor union bound (see [65]) or can be implemented as a look-up table of BER vs. SNR. The available MCSs currently supported in G3-PLC from most robust to highest throughput are ROBO (4x repetition-coded DBPSK signal), DBPSK, DQPSK, and D8PSK. Figure 3.7(bottom) shows an instantiation of $\mathbf{C}^\star$ based on the SNR over the signaling band as seen in Figure 3.7(top).

Figure 3.8 compares the throughput of static modulation allocations with tone mask versus a cyclic adaptive modulation scheme based on the observations of the first 12 symbols (corresponding to ~half the AC cycle). A tone mask is applied to disable tones that do not meet the SNR threshold of a given modulation (i.e. $p_b \leq 10^{-2}$) for at least half the OFDM frame duration. This causes the higher rate D8PSK to underperform the lower rate DQPSK. The cyclic adaptive modulation scheme approaches the genie assisted (perfect noise information for all time), which is displayed as a dashed line. Substantially higher gains could be realized with a larger modulation set.

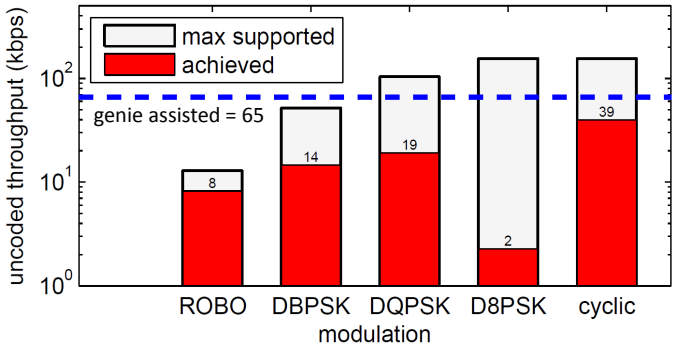

Figure 3.8: Throughput gains realized by an NB-OFDM PLC system employing cyclic adaptive MCS scheme over the G3-PLC CENELEC-A band. Throughput is boosted 2× versus the conventional G3-PLC operation using the next-best rate-optimal choice DQPSK with tone map.

3.3.1.2 G3-PLC Cylostationary Noise Simulator

To further evaluate the proposed method, simulations were constructed using a G3-PLC compliant simulator built in MATLAB using a reconfigurable channel model. Of interest are impulsive noise models accepted into the IEEE P1901.2 standard [5] and a model used to simulate narrowband interference. The link consists of an encoder, modulator, channel model, demodulator, and decoder with settings listed in Table 3.5. Performance metrics used are BER, block error rate (BLER), and packet length (in units of OFDM symbols). Three typical PLC channels are evaluated: (*i*) mild noise profile, (*ii*) moderate noise profile, and (*iii*) moderate noise profile plus narrowband noise. Narrowband noise is representative of interference from other PLC devices, jammers, or ingress radio stations or narrowband noise generated by a non-communicating source connected to the powerline.

Table 3.5: G3-PLC link simulation parameters.

Parameter	Setting
user data N_{PDSU}	64 bytes
target BER p_b	$\{10^{-1}, 10^{-2}\}$
inner code	convolutional (rate = 1/2, $K = 7$)
outer code	Reed-Solomon ($N = 255$, $K = \{239, 247\}$, $T = \{8, 4\}$) (depends on modulation type)
interleaver type	G3-PLC interleaver configured for D8PSK
low SNR frame size limit	max length = 128, per block bit-error minimizing allocation

3.3.1.3 Mild Noise Profile

The mild noise profile refers to cyclostationary noise that can potentially corrupt up to one OFDM symbol per half AC cycle—i.e. the impulse duration is less than the duration of an OFDM symbol N_{OFDM}/fs. For this case, the IEEE P1901.2 noise model for impulse profile *LV3* is used [5]. Figure 3.9 shows the communication performance of the proposed method over existing methods in the G3-PLC standard that are implemented in the simulator. As shown in the figure, raw BER for DBPSK, DQPSK, and D8PSK all display the typical monotonically decreasing BER curves. As expected, DBPSK is the most robust modulation of the three, and each MCS is separated by about 4 dB for a given operating point (SNR where BER is fixed).

Over the same channel, the cyclic adaptive MCS algorithm is used with target BERs of $p_b = 10^{-1}$ and $p_b = 10^{-2}$. As seen in the figure, operating points are improved in the low SNR regime. For $p_b = 10^{-1}$, the improvements occur at a lower SNR and remain constant at a level of 5×10^{-2} BER until it approaches the D8PSK curve at high SNR (since further MCS choices are not available). The BER curve traces the D8PSK curve at high SNR since all

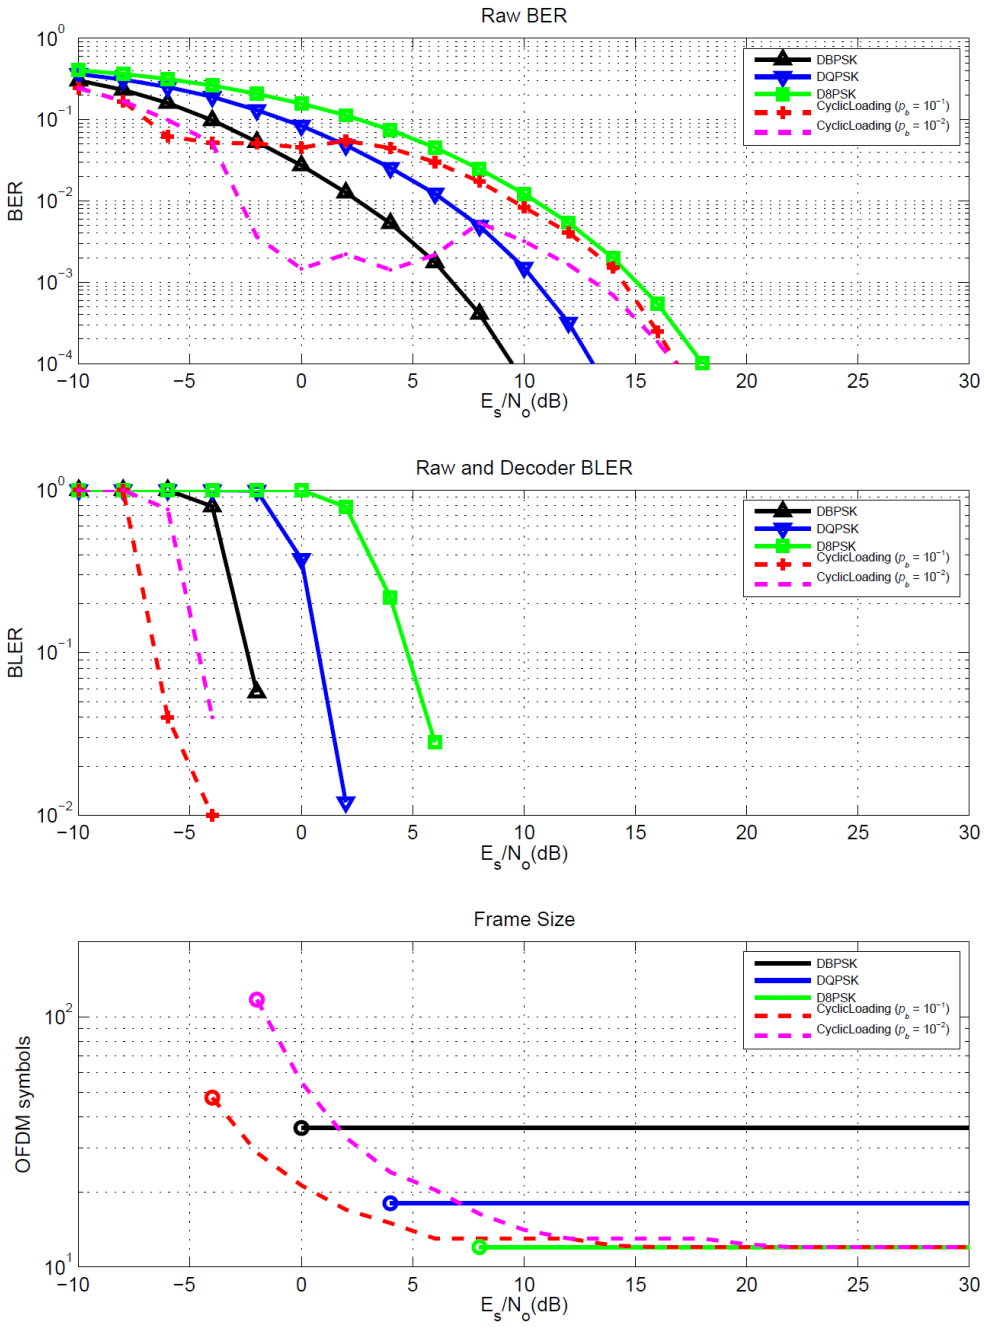

Figure 3.9: Mild cyclostationary noise *LV3*. Using the proposed method (dashed lines), receiver sensitivity is increased 4 dB for $p_b = 10^{-1}$ and 2 dB for $p_b = 10^{-2}$. Payload length tracks static modulation schemes but provides up to 3× compression versus DBPSK since bit density is dynamically increased.

subcarriers will be populated with the highest modulation (D8PSK) at high SNR. The $p_b = 10^{-2}$ curve follows similar trends. The key differences are that, due to the higher target bit error, the BER knee occurs at a higher SNR since there are fewer subcarriers that meet the minimum selection threshold; however, once met, the algorithm will settle at a lower BER between 1×10^{-3} to 5×10^{-3} until approaching the D8PSK curve at high SNR.

Though raw BER curves are useful for qualitative comparison, a more informative metric for true link performance is block error rate (BLER) which measures the error rate of a frame of data after FEC decoding (either all bits were decoded correctly or not). Typically, BLER drops dramatically once a minimum number of bit errors is reached since the FEC allows correction of bit errors below some threshold. As seen in the BLER plot, the three MCSs are once again separated by ~4 dB in the same order as the BER curves. The proposed method shifts the knee of the BLER curves by 2 dB for the $p_b = 10^{-2}$ and 4 dB for the $p_b = 10^{-1}$ configurations, demonstrating an improved operating point of 2 and 4 dB, respectively. An additional metric of interest is the resulting payload size in OFDM symbols since the proposed method varies length depending on the bits allocated per subcarrier over the OFDM frame. The lower plot shows payload length as a function of SNR. A circle denotes the minimum SNR required to achieve a successful operating SNR (BLER $\leq 10^{-2}$). Payload length monotonically decreases with increasing SNR by virtue of increased bit density. In the asymptotic limit, payload length converges to the highest static MCS (D8PSK).

3.3.1.4 Moderate Noise Profile

Moderate impulsive noise is noise that corrupts two OFDM symbols per half AC cycle. For this case, noise profile *LV5m* is used. As shown in Figure 3.10, the static modulation schemes DBPSK, DQPSK, and D8PSK suffer from severe degradations in their BER. Typical BER operating points are shifted by up to 15 dB SNR from those in the mild impulsive noise case.

For this noise environment, cyclic loading provides more dramatic improvements than those observed in the mild impulsive noise case. This can be attributed to larger variations of noise power spectral density over the AC cycle since there is more power in the impulsive portions of the noise. This is more optimal in the presence of colored noise versus spreading the information uniformly (which is what is done using a static modulation scheme with the G3-PLC interleaver). In terms of BER, operating points are improved up to 20 dB in this case. In terms of BLER, operating points are improved between 10 and 12 dB for $p_b = 10^{-2}$ and $p_b = 10^{-1}$ at an operating point of BLER = 10^{-2}. minimum SNR operating point is much further to the left than those for the static modulation schemes.

Payload length is once again longer than the static schemes at low SNR; however, at high SNR, payload length is reduced to increased bit density. For this noise case, there is a more dramatic difference in payload length between the static and cyclic modulation schemes. Once again, the upper limit in terms of payload length compression is 3×, which is achieved vs. DBPSK at high SNR.

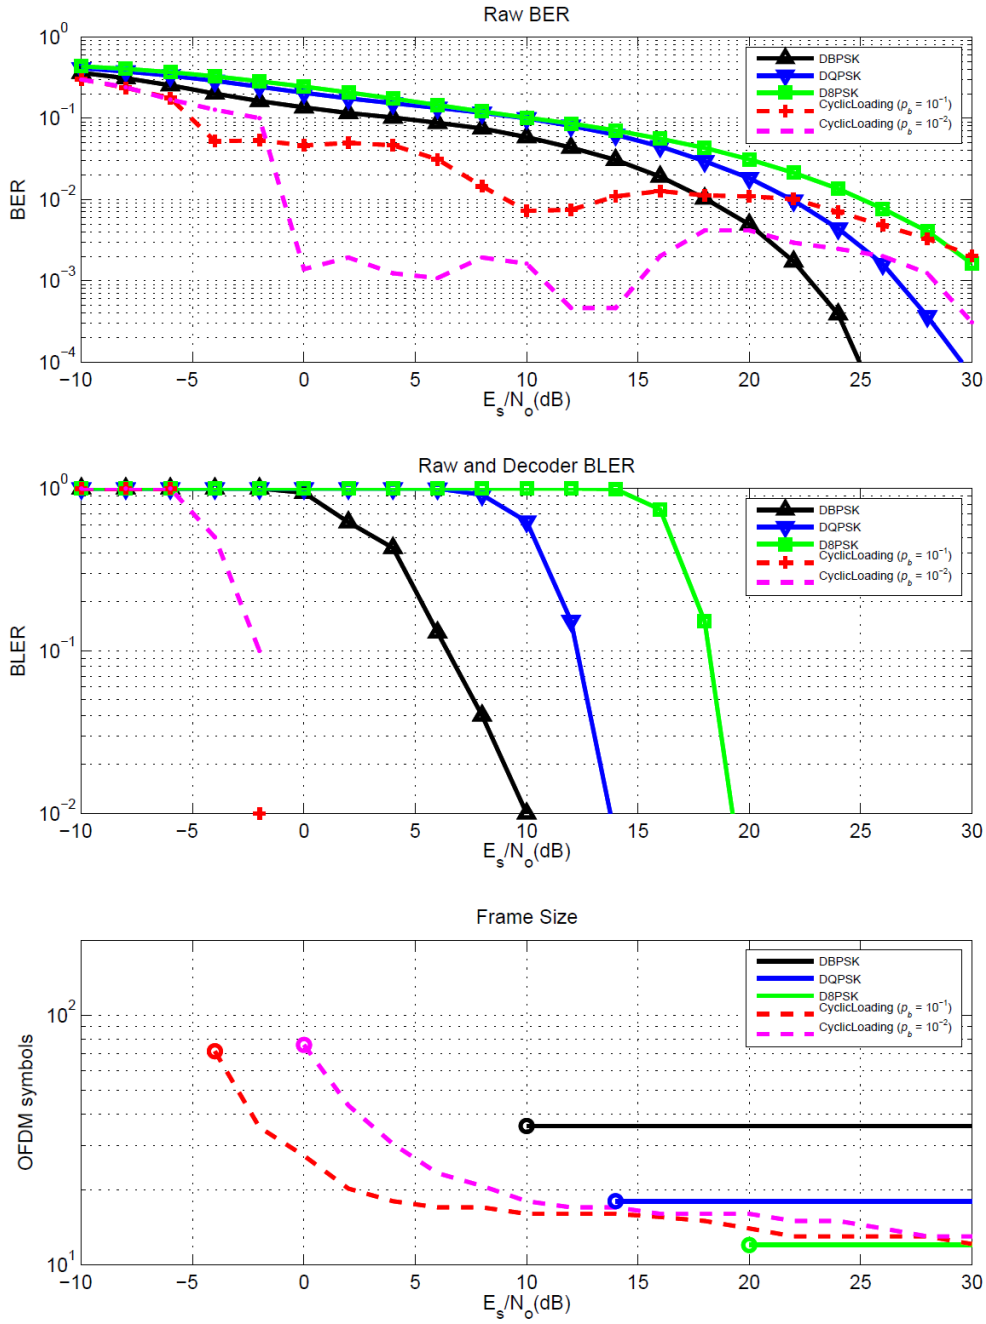

Figure 3.10: Moderate cyclostationary noise *LV5m*. All three modulation schemes are severely impaired in this case and are not able to achieve a BLER of 10^{-2} until 10 dB SNR or higher. Using cyclic bit loading, target BLER is achieved at -4 dB SNR for $p_b = 10^{-1}$ and at -2 dB for $p_b = 10^{-2}$.

3.3.1.5 Moderate Noise Profile + Narrowband Noise

In this noise channel, narrowband interference was introduced in the center of the band for noise profile *LV5m*. The narrowband source has a frequency of 63.28125 kHz (centered between frequency bins 40 and 41) and a strength of 0 dB interference to signal ratio (ISR). In this case, the narrowband interferer corrupts data in frequency bins 40 and 41 for all SNR values and has leakage into other subcarriers (since it is not centered on an FFT bin).

As shown in Figure 3.11, BER is substantially degraded across all SNRs for each static MCS. The interleaver for DBPSK spreads the errors created in frequency bins 40 and 41 across many bytes. These errors are not able to be corrected using the Reed-Solomon (255,239) code, revealing a weakness of the G3-PLC interleaver. For DQPSK and D8PSK, errors are spread across fewer bytes, affording a higher probability of correction using the G3-PLC FEC.

Using cyclic adaptive MCS selection, raw BER is substantially improved over static modulation schemes. BER operating points are able to be improved up to 30 or more dB. Decoding performance is also dramatically improved. Since bad subcarriers are able to be avoided, issues with the DBPSK interleaver are circumvented. For a BLER operating point of 10^{-2}, cyclic loading is able to improve sensitivity by 26 dB for $p_b = 10^{-1}$ and by 24 dB for $p_b = 10^{-2}$ versus DQPSK (since DBPSK is not able to achieve BLER $\leq 10^{-2}$). To achieve improvements in BLER, payload length is increased in the low SNR regime. Since DQPSK is the only static modulation scheme to achieve BLER $\leq 10^{-2}$, it is the only reference point for payload length comparison.

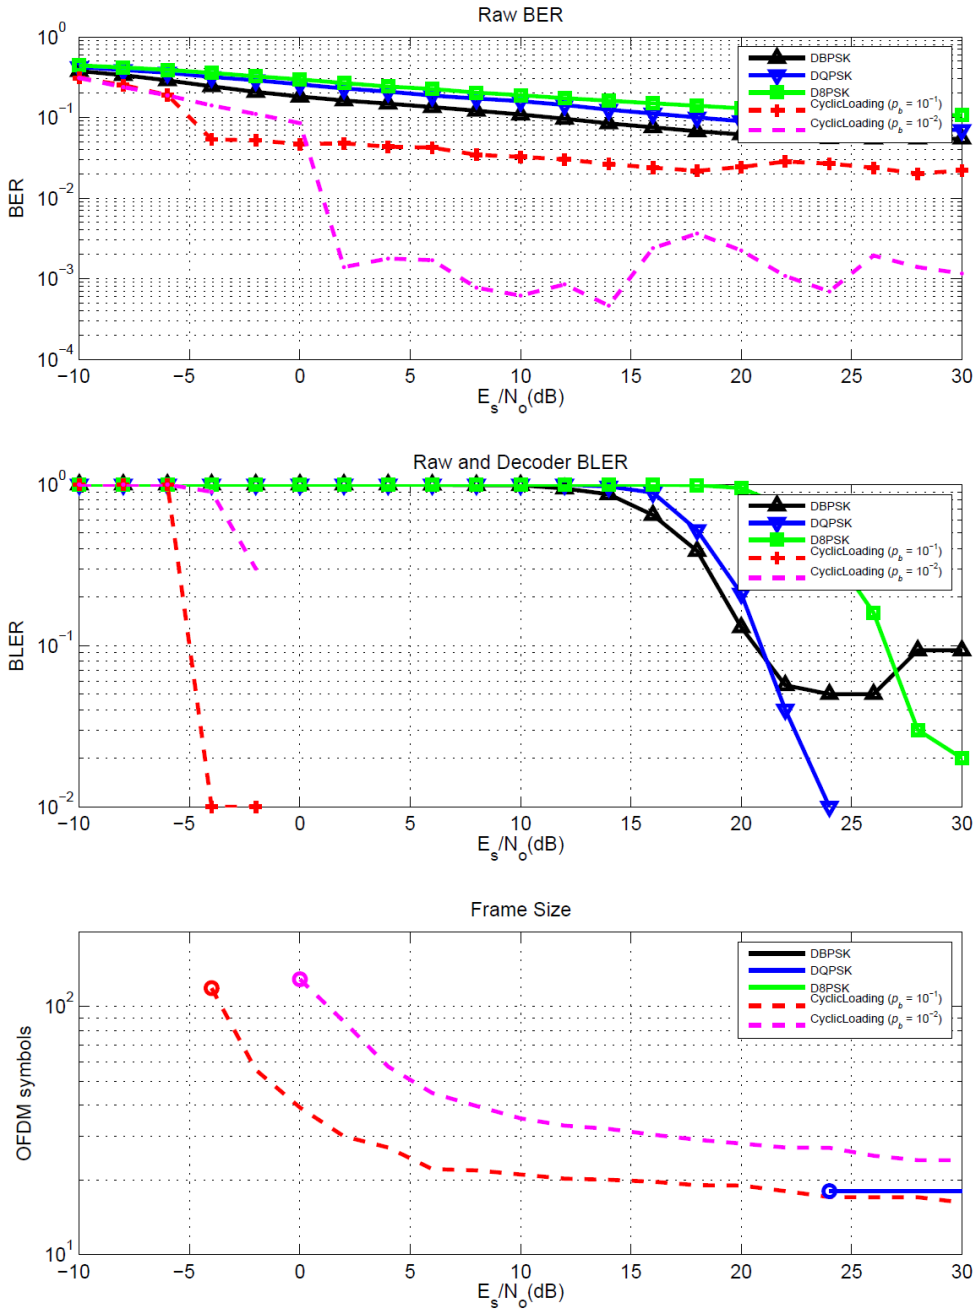

Figure 3.11: Moderate cyclostationary noise *LV5m* plus narrowband noise. Large groups of subcarriers are affected causing substantial reduction in BER and BLER. Cyclic adaptive MCS selection provides drastic improvements in decodability in these channel conditions.

3.3.2 Implementation of the AMP OFDM Receiver

To address asynchronous impulsive noise in NB-OFDM PLC channels, a compressive sensing approach is used. This approach takes out-of-band frequency domain samples of the noise in null subcarriers (those specific to G3-PLC are defined in Table 3.4). Null subcarriers located at the band edge are common in many OFDM systems. These null subcarriers are used to ease front-end filter design requirements. The compressive sensing approach used here is approximate message passing (AMP), whose benefits are discussed in Section 3.2.3. Namely, this approach is chosen because it separates an estimation problem involving matrix inversions into purely scalar operations, lending itself to efficient hardware implementation [57–59]. The framework used here was proposed in [50] for joint impulsive noise mitigation and channel estimation in OFDM systems. This framework exploits the spreading of impulsive noise information across subcarriers by virtue of the DFT. Frequency-domain noise samples in null subcarriers are used to reconstruct time-domain noise using AMP. As referenced in Table 3.4, G3-PLC's baseband signaling has 92 such null subcarriers, of which I use 64 (subcarriers 1-22, 59-100).

3.3.2.1 AMP Formulation for NB-OFDM PLC

The AMP algorithm takes as inputs the samples of the null subcarriers y_{Ω_i} from the set Ω where $|\Omega| = M < N$. M is number of null subcarriers used, and N is the total number of subcarriers and thus the FFT size of the OFDM system. The process of sampling null subcarriers can be implemented

efficiently by indexing the corresponding entries of the FFT output. This can be implemented efficiently as an enhanced FFT loadout procedure. I express this as $\mathbf{I}_\Omega$, a subset of the Ω rows of the $N \times N$ identity matrix $\mathbf{I}$. Table 3.6 outlines the explicit computational steps of the AMP formulation for this system model. Detailed derivations are provided in [50]. The inputs to this algorithm are the M-length vector of null subcarrier samples $\mathbf{y}_\Omega$, background and impulsive noise variance (γ_B, γ_X), and impulse probability π. The output is the N-length time-domain estimate of the impulsive noise $\hat{x}_j(t+1)$, where t here denotes iteration count.

The first step of AMP is an initialization step. The second, or output linear step, involves the forward FFT operation where a_{ij} denotes the ij^{th} entry of the transform matrix, the DFT in this case. My target architecture, the Xilinx Virtex-5 FPGA, implements the forward FFT as $X_k = \sum_{n=0}^{N-1} x_n \omega_N^{nk}$ and the IFFT as $x_n = (1/N) \sum_{k=0}^{N-1} X_k \omega_N^{-nk}$, where $\omega_N = e^{-j2\pi/N}$. The scaling by N and M factors are used to balance the scaling performed in hardware. Step 3 is the output non-linear step which consists of simple scalar operations. Step 4, the input linear step, involves the IFFT. Step 5 requires the highest depth of sequential operations, the most prohibitive of which is the exponential since its output has a high dynamic range. This was averted using an approximation discussed in Section 3.3.2.3. The divide also presents challenges for real-time operation. Steps 2-5 are repeated until convergence. In my implementation, I use 4 iterations to balance performance with throughput constraints.

Table 3.6: AMP impulsive noise mitigation for Bernoilli-Gaussian noise

Step	Calculation
1	$\hat{x}_j = 0$ $\tau_j^x(t) = (1-\pi)\gamma_X, \qquad \forall j \in \{1,...,N\}$ $\hat{s}_i(t) = 0, \qquad \forall i \in \{1,...,M\}$
2	$\tau_i^p(t) = \sum_j \lvert a_{ij}\rvert^2 \tau_j^x(t) = \sum_j \tau_j^x(t)$ $\hat{p}_i(t) = \sum_j a_{ij}\hat{x}_j(t) - \tau_i^p(t)\hat{s}_i(t-1)$ $= \mathbf{I}_\Omega \text{FFT}(\hat{x}(t)) - \tau_i^p(t)\hat{s}_i(t-1)$
3	$\tau_i^s = \frac{1}{\gamma_B + \tau^p}, \qquad \forall i \in \{1,...,M\}$ $\hat{s}_i(t) = \tau_i^s(t)\,(y_{\Omega_i} - \hat{p}_i(t))$
4	$\tau_j^r(t) = N\left[\sum_i \lvert a_{ij}\rvert^2 \tau_i^s(t)\right]^{-1}$ $\hat{r}_j(t) = \hat{x}_j(t) + \tau_j^r(t)\sum_i (a_{ij}^* \hat{s}_i(t))$ $= \hat{x}_j(t) + \tau_j^r(t)\text{IFFT}\left(\mathbf{I}_\Omega^T \hat{x}(t)\right)$
5	$\eta_j = \frac{\pi}{1-\pi}\frac{\tau_j^r(t)}{\gamma_I+\tau_j^r(t)}\exp\left(\frac{\gamma_I \lVert\hat{r}_j(t)\rVert^2}{\tau_j^r(t)\left(\gamma_I+\tau_j^r(t)\right)}\right)$ $\rho_j = \frac{\eta_j}{1+\eta_j}$ $\hat{x}_j(t+1) = \frac{\gamma_I}{\gamma_I+\tau_j^r(t)}\rho_j \hat{r}_j(t)$ $\tau_j^x(t+1) = \eta\left[\frac{\tau_j^r(t)\gamma_I}{\gamma_I+\tau_j^r(t)}\right.$ $\left. + \left(\frac{\gamma_I}{\gamma_I+\tau_j^r(t)}\right)^2 \lVert\hat{r}_j(t)\rVert^2(1-\eta_j)\right]$
*	repeat steps 2-5 until convergence, resulting time-domain noise estimate: $\hat{x}_j(t+1)$

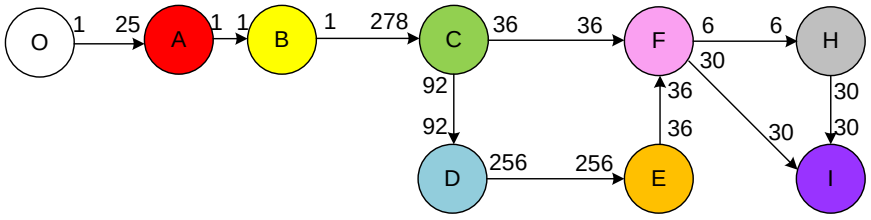

Figure 3.12: Synchronous dataflow (SDF) model of the OFDM receiver with AMP. Nodes are tasks and edges represent data dependencies. When task C executes, it consumes 278 samples on its input, and produces 92 samples on its bottom output and 36 samples on its right output.

3.3.2.2 Synchronous Dataflow Model

The streaming operation of the AMP-enhanced OFDM receiver can be modeled in synchronous data flow (SDF) as shown in Figure 3.12. Nodes are application tasks and edges are first-in first-out (FIFO) queues that represent data dependencies. Each task produces and consumes a fixed number of samples in each execution. The tasks in Figure 3.12 correspond to (A) sample rate conversion, (B) time and frequency offset correction, (C) FFT and CP removal, (D) AMP noise estimation, (E) FFT, (F) noise subtraction, (H) channel estimation, and (I) channel equalization. The deterministic properties of SDF enable fully-static analysis and synthesis of the system behavior. In particular, a periodic, static schedule can be determined by analysis of the SDF. For example, a periodic schedule for the SDF graph in Figure 3.12 is (278A)(278B)CDEFHI. In Section 3.3.2.3, I convert this SDF to a globally asynchronous, locally synchronous (GALS) model of computation using LabVIEW DSP Design Module. GALS affords more flexibility in hardware timing while still providing benefits of SDF.

3.3.2.3 Mapping to Hardware

Before proceeding with the hardware implementation, the AMP impulsive noise mitigation algorithm in Table 3.6 was mapped from floating-point to fixed-point. This typically involves determining the dynamic range of each variable and assigning an appropriate fixed-point representation of it that guarantees a specified level of performance. Although in-range data values, or the

dynamic range captured by the fixed-point representation, correlates to performance, more complicated dependencies on variable sizings were observed. To address this, I simulated the algorithm using the `fi` and `NumericScope` data types, both part of MATLAB's Fixed-Point Toolbox, as follows: (*i*) for each run of the algorithm I log the variables values, (*ii*) then, using the `NumericScope` I appropriately size these variables, (*iii*) finally, using the `fi` data type I simulate the algorithm with the variable sizing found in step (*ii*) and ensure target performance is met. The above procedure was performed for a fixed set of input parameters and input noise statistics—i.e., $\gamma_B = 0.0025$, $\gamma_X = 0.25$, and $\pi = 0.1$. The performance metric for the MATLAB simulation was the reduction in impulsive noise power over a large number of trials using simulated noise with matched parameters. This fixed-point mapping performs within 0.5 dB of the double precision floating-point MATLAB version of the algorithm for this parameter set.

The resulting variable sizings for the AMP implementation are shown in Table 3.7. My fixed-point data type representation convention is $SN_W.N_I$, where S is replaced by 'U' for unsigned and 'I' for signed, N_W is the wordlength in bits, and N_I is the integer width (or scaling by 2^{N_I}). The fractional width is $N_W - N_I$, and the sign bit is absorbed in the integer N_I. By targeting 90% or higher in-range data values, most variables were able to be sized within 16-bit wordlengths, an efficient size to utilize the single-cycle digital signal processing (DSP) 48 blocks in the Xilinx Virtex-5 FPGA.

An OFDM transceiver using the fixed-point version of the AMP algo-

rithm was implemented across several Xilinx Virtex-5 FPGAs. The intended hardware mapping was to one FPGA transmitter and one FPGA receiver. After initial sizing estimates, the AMP algorithm and the OFDM receiver were not able to fit within a single FPGA. The receiver was partitioned across two FPGAs that share data across the PXI-Express (PXIe) bus, an instrumentation-targeted version of the PCI-Express bus developed by National Instruments. Hardware mapping for the three FPGAs is detailed in the following section and presented graphically in Figure 3.13.

As shown in the figure, the OFDM transmitter interleaves modulated 2 $\times$ I8.1 complex symbols with a conjugate-symmetric pair of data and reference symbols. Reference symbols are encoded as quadrature phase shift keying (QPSK) signals and are interleaved between every 5 data subcarriers. These symbols are used for channel estimation at the receiver. Conjugate-symmetry is enforced in order to make the output purely real-valued, and coherent QPSK is used as the data subcarrier modulation.

The receiver consists of two FPGAs. The first FPGA, 'G3RX', performs front-end processing—i.e. resampling and time/frequency synchronization. Resampling is performed using a multi-rate, two-stage finite impulse response (FIR) filter whose coefficients were designed using a LabVIEW multi-rate filter design tool. Synchronization is performed by correlating the signal with a delayed version of itself. Since the cyclic prefix (CP) is periodic, its autocorrelation exhibits a peak at a lag equal to the the symbol duration. The location of this peak and its resulting phase can be used for time and carrier

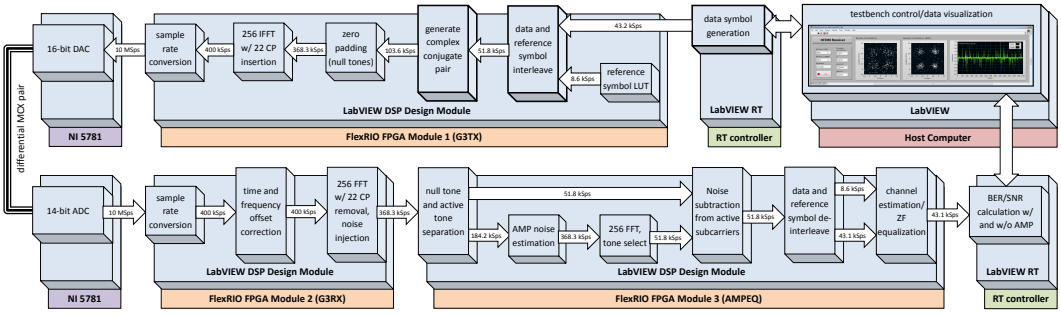

Figure 3.13: A block diagram of the AMP G3-PLC OFDM test system. Physical hardware blocks are separately colored, and software environments are shown next in the hierarchy. Functional blocks are the highest hierachical level with throughput rates and dataflow direction shown between them.

Table 3.7: Fixed-point word sizes for AMP parameters and variables.

Inputs	Complex?	Length	Representation
π		1	U16.0
γ_X, γ_B		1	U16.-2
y_i	✓	128	2 × I16.1
Intermediates	**Complex?**	**Length**	**Representation**
$\hat{x}_j(t)$	✓	256	2 × I16.1
$\tau_j^x(t)$		256	U16.-2
$\hat{s}_i(t)$	✓	128	2 × I16.6
$\tau_i^p(t)$		1	U16.-5
$\hat{p}_i(t)$		128	2 × I16.0
τ_i^s		1	U16.0
$\tau_j^r(t)$		1	U16.-2
$\hat{r}_j(t)$	✓	256	2 × I16.2
η_j		256	U16.6
ρ_j		256	U16.6
Outputs	**Complex?**	**Length**	**Representation**
$\hat{x}_j(t+1)$	✓	256	2 × I16.1

frequency offset synchronization [66]. Frequency offset correction is performed using a discrete cosine generator synthesized to 0.1 Hz phase increments using Xilinx CoreGen. The second receiver FPGA, 'AMPEQ', performs the AMP algorithm, channel estimation, and equalization. The AMP algorithm uses frequency-domain null subcarrier samples and parameters as input from 'G3RX'. Reference symbols are de-interleaved and used for channel estimation. Equalization is performed using a zero-forcing (ZF) equalizer.

In parameterizing the AMP algorithm for G3-PLC signaling, M and N are chosen to be 128 and 256, respectively. $M = 128$ instead of 64 (the number of null subcarriers) because I implement the full complex version of AMP based on 256-length complex-valued FFTs/IFFTs, whereas G3-PLC uses a 256-length real-valued FFT. In my implementation, I use null subcarriers 1-22 and 59-100 and their negative frequency pairs 234-255 and 156-197. The exponential function in Step 5 is approximated using a sixth-order Taylor series expansion. This was done to avert dynamic range issues when rescaling the output of the Xilinx CORDIC computation, which requires the inputs be normalized to $[0, 1]$. The ρ_j calculation in Step 5 in Table 3.6 involves a 16-bit division 256 times per iteration. This operation was parallelized into two streams, resulting in a 8656 cycle reduction in execution time at the expense of a small increase in resource utilization.

Figure 3.13 shows the system block diagram of the AMP-Enhanced OFDM transceiver testbed. The testbed consists of two National Instruments PXIe-1082 chasses, one for the TX and one for the RX. Both chasses make use

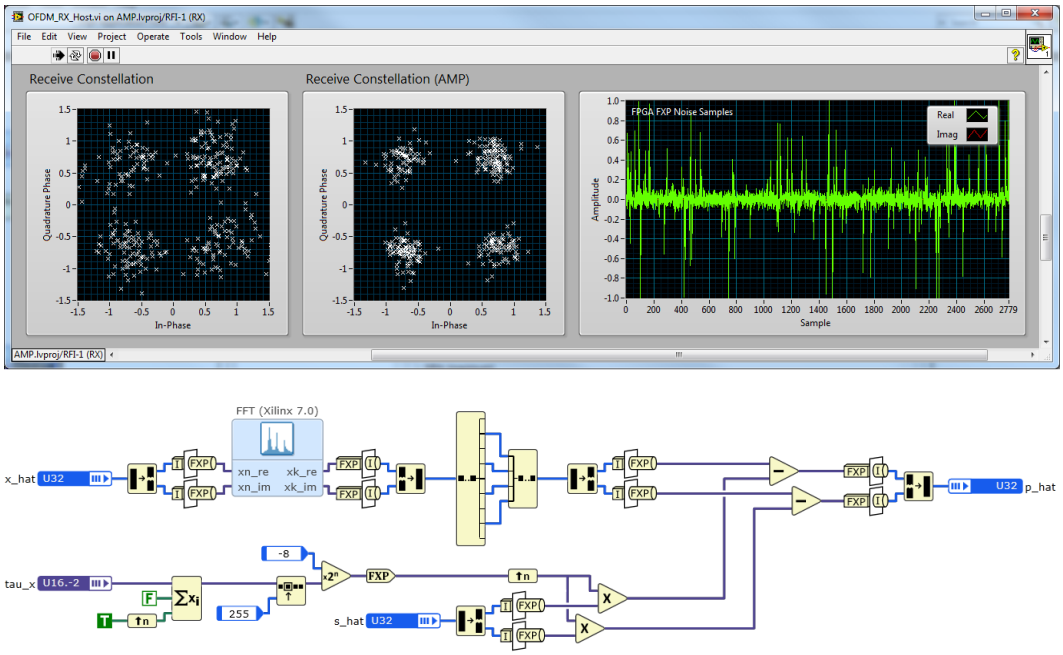

Figure 3.14: (top) The receiver host VI, showing input impulsive noise and an equalized constellation for the conventional and AMP-enhanced OFDM receivers over one OFDM frame. (bottom) The LabVIEW DSP Design Module graphical representation of the computations in Step 2 of Table 3.6.

Table 3.8: Resource utilization for each FPGA.

Resource Name	Total Available	FPGA1 *G3TX*	FPGA2 *G3RX*	FPGA3 *AMPEQ*
Total slices	14720	32.6%	64.0%	94.2%
Slice registers	58880	15.8%	39.3%	59.0%
Slice LUTs	58880	17.6%	42.4%	71.4%
DSP48s	640	2.0%	7.3%	27.3%
Block RAMs	244	7.8%	18.4%	29.1%
MCLK (MHz)	-	125	40	40

of an onboard PXIe-8133 1.73 GHz Quad-Core (Intel Core i7-820QM) PXI Express controller running a real-time operating system. These PXI controllers have been targeted to the deployment of LabVIEW Real-Time (RT) applications which can execute deterministic dataflow computations at granularity on the order of ~1 ms. Each of the two systems communicates via gigabit Ethernet back to a host PC for high-level, non-deterministic performance analysis and visualization. Each of the two RT chasses are configured as follows:

- **TX chassis** is fitted with a single NI PXIe-7965R FlexRIO FPGA Module named 'G3TX' featuring a Virtex-5 SX95T FPGA that is used for interleaving data/reference symbols, OFDM modulation (IFFT), appending cyclic prefix (CP), and upsampling to the 10 MS/s NI-5781 digital-to-analog converter (DAC) sample rate. The host RT chassis generates random data symbols to feed the input symbol direct memory acces (DMA) first-in, first-out (FIFO) buffer. Modulated and upconverted samples are clocked out of the FlexRIO and into the NI-5781 adapted module. Theses samples are then passed across two micro-coaxial (MCX) 50Ω differential pairs to the RX NI-5781 (although the output of the quadrature component is always zero by virtue of the conjugate-symmetry enforced in 'G3TX').

- **RX chassis** is fitted with two NI PXIe-7965Rs. The first FPGA module, 'G3RX', is configured for the front-end receiver processing—i.e. downsampling, synchronization, frequency-offset estimation, OFDM demodu-

lation (FFT), CP removal, and noise injection. The second FPGA module, 'AMPEQ' is passed frequency-domain samples from 'G3RX' using a peer-to-peer (P2P) stream over the PXIe backplane. These samples, in addition to parameters γ_X, γ_B, and π are passed from the RX host controller and are used as inputs to the AMP algorithm.

LVDDM is a high-level FPGA synthesis tool developed by National Instruments. LVDDM takes as an input a high-level data flow diagram of DSP computations (see Figure 3.14) and translates them to a LabVIEW FPGA register-transfer level representation. This representation is then mapped to a Verilog hardware description language product that is input to the Xilinx Integrated Software Environment FPGA compiler. LVDDM provides direct interfaces to Xilinx CoreGen features, allowing for the implementation of highly-optimized FPGA constructs including FFTs, discrete fourier transforms, direct digital synthesizers, and multi-rate FIR filter implementations.

The resource utilization and master clock (MCLK) of each of the three FPGAs used in the test system is shown in Table 3.8. Resource utilization is highest in 'AMPEQ', the FPGA with AMP processing and equalization. Overall resource usage could be reduced through further computation and buffer size optimization. Additionally, the use of real- vs. complex-valued FFTs could afford at least 2× reduction in resource utilization. Complex form was used so that the algorithm could be applied to other projects using complex signaling.

3.3.2.4 Real-Time Performance Benchmarks

To evaluate the real-time performance of this implementation, a bit-error-rate (BER) testbench was constructed in LabVIEW RT running on each embedded PXI controller. Equalized symbols, input impulsive noise, error vector magnitude (EVM), and trailing bit-error-rates (BER) are plotted for both the AMP-enhanced and conventional OFDM receiver (see top of Figure 3.14). Blocks of 6 OFDM symbols were processed at a time. Using this framework, a test array of input and output data bytes were transmitted accross the OFDM link using synthesized Gaussian mixture noise input to the receiver. Received symbols were demodulated and compared for errors. Using the test environment, the receiver was configured with a fixed π, $\gamma_X = 0.25$, and γ_B was set to be a multiple of γ_X listed in Figure 3.15. Transmit gain was swept over a 20 dB range in order to profile BER vs. SNR. Here, SNR is the transmit power over background noise power, γ_B.

The results of the BER analysis are shown in Figure 3.15. As shown in the figure, by mitigating impulsive noise with AMP, SNR levels of up to 8 dB are recovered and gains are realized over a wide SNR range. The AMP and conventional receivers have nearly identical BER curves in non-impulsive, or $\pi = 0$, noise. Injected noise was synthesized using a match to AMP input parameters; however, estimation techniques discussed in [50] could be used to learn the parameters adaptively using amortization over OFDM symbols. These results indicate that this algorithm could be targeted for implementation in other OFDM receivers beyond G3-PLC systems.

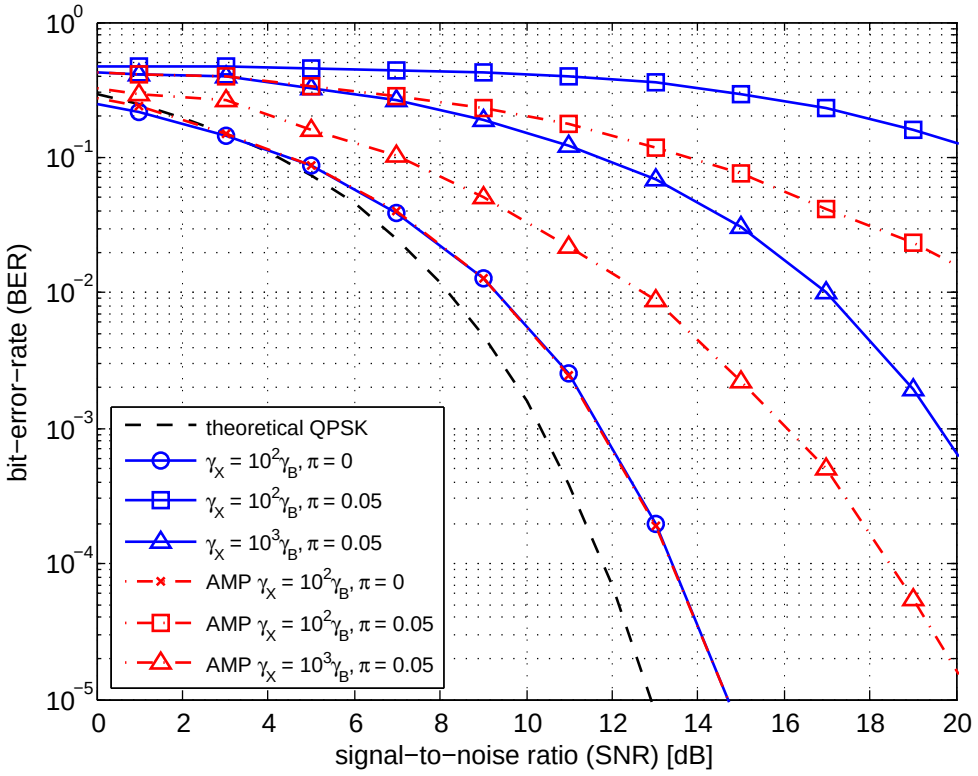

Figure 3.15: BER plots of received symbols with and without AMP algorithm for impulse-free noise and for impulsive noise with $\pi = 0.05$ (5% probability of impulse) using impulses 20 and 30 dB above the background level.

3.4 Conclusion

In this work, I have characterized cyclic spectral statistics of power-line noise. It is confirmed that the noise is predominantly cyclostationary with dominant period equal to half the AC cycle. Other noise such as asynchronous impulsive noise and narrowband noise are observed. I use these results to develop two methods: (*i*) an adaptive cyclic modulation and coding scheme method over time and frequency and (*ii*) an implementation of an asynchronous impulsive noise mitigation technique.

First, I demonstrate the utility of this adaptive cyclic MCS adaptation using measured noise samples and the G3-PLC MCS set. My results show that this scheme is capable of offering up to a 2× increase in throughput using the current G3-PLC modulation set in addition to up to 25 dB operating point improvements. Second, I design and implement a real-time OFDM receiver with AMP impulsive noise mitigation on FPGAs. My design flow was to (*i*) model the receiver in synchronous dataflow (SDF), (*ii*) convert the receiver to fixed-point data and arithmetic, and (*iii*) synthesize the fixed-point SDF model onto two Xilinx Vertex-5 FPGAs. In the third step, I generalize the SDF receiver model to a globally asynchronous, locally synchronous computational model in LabVIEW DSP Design Module. My OFDM transceiver testbed achieves real-time streaming throughput at G3-PLC rates and mitigates up to 8 dB SNR of impulsive noise.

These two methods can be used independently to improve robustness and performance in the presence of two dominant NB-PLC noise sources.

Chapter 4

Space-Time-Frequency Methods for Multi-Antenna Cellular Communication Systems

4.1 Motivation

To satiate mobile data demand in a wireless spectrum crunch, space truly is the final frontier. Future wireless networks rely upon spatial reuse to scale network capacity and improve network service. Research in massive multiple-input, multiple-output (Massive MIMO) suggests that even further capacity increases can be achieved by exploiting channel reciprocity with an excess number of antennas at the basestation. Initial measurements show promising results that could allow this technology to be applied in practice. Unfortunately, real-time testbeds capable of Massive MIMO are either not publicly available, not real-time, or both. In this work, I develop an flexible platform that supports prototyping up to 20 MHz bandwidth 128-antenna MIMO. They key challenges I address are (*i*) scaling data rates and interfaces to 128 antennas, (*ii*) providing low-latency data transfer to support reciprocity even at high-mobility, and (*iii*) synchronizing time and frequency over 128 antennas, all with commercial off-the-shelf hardware. This platform is used to form LuMaMi, the world's first 100-antenna MIMO system [8].

4.2 Background

The continuing pursuit of increased data rates and communication performance in cellular systems has led to the expansion of communication signal processing to the spatial dimension. Spatial processing techniques are broadly referred to as multiple-input multiple-output (MIMO). For wireless channels, MIMO involves the use of multiple transmit and/or receive antennas. Since its emergence in the mid 1990s [67] and perhaps earlier [68], MIMO technologies have been successful in practice, leading to deployments in Wi-Fi and cellular systems. More specifically, MIMO methods introduced in 802.11n support up to four transmit and four receive antennas [69]. Cellular systems based on LTE Release 8 support up to four antennas in the basestation per sector and one at the mobile station [70]. These numbers are increased to eight and four, respectively, in the soon-to-be-deployed LTE-Advanced (LTE-A), or LTE Release 10 [4]. Multiple antennas in these systems can be used to increase link robustness using space-time block codes (STBC) or data rate by applying spatial-multiplexing (SM). The success of SM requires a favorable propagation environment and typically high SNR [71].

MIMO techniques can be extended beyond point-to-point to multi-user applications using multi-user MIMO (MU-MIMO). MU-MIMO can be used to separate users in space, allowing for further network densification and increased capacity. MU-MIMO provides higher guarantees for spatial multiplexing than point-to-point system but inherits challenges in near-far power control and time/frequency synchronization. MU-MIMO modes have

Table 4.1: Parameters in MIMO cellular systems

Parameter	Variable	Typical Values
# of user equipment antennas	N_{UE}	1-8
# of basestation antennas	N_{BS}	1-8
# of user equipments in MU-MIMO	K	1-8

Table 4.2: Scaling Factors of Channel Estimation Overhead in MIMO Systems

Type	Downlink	Uplink
point-to-point MIMO	N_{UE}	N_{BS}
MU-MIMO	KN_{UE}	N_{BS}
Massive MIMO (assume $N_{UE} = 1$)	K	K

been provisioned as part of the LTE and LTE-A standards. Typical values for MIMO parameters in cellular systems are summarized in Table 4.1.

In MIMO, the high dimensionality of the channel imposes additional overhead in channel estimation. Channel estimation is required for coherent communications and scales with the number of antennas. As the number of antennas increase, so does the amount of overhead needed to estimate the higher dimensional channel since orthogonal estimates are needed. For example in LTE [70], going from 1 to 4 antennas increases channel estimation overhead from 4.8% to 14.3%. This limitation can be overcome using reciprocity-based MIMO, which exploits reciprocity in uplink and downlink signaling to ease channel estimation overhead. Relaxing channel estimation requirements allows the number of antennas on one side of the link to grow without bound as shown in Table 4.2 [72]. Such large-dimensional, or *Massive MIMO* systems provide distinct performance advantages over traditional MIMO by scaling the number of antennas by an order of magnitude or more [73–75].

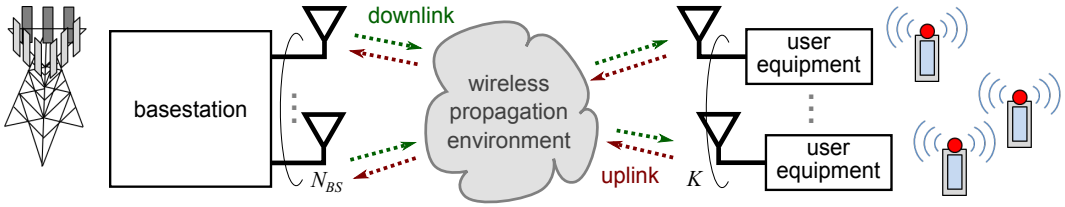

Figure 4.1: A Massive MIMO basestation with N_{BS} antennas serving K simultaneous, single-antenna users where $N_{BS} \gg K$.

4.2.1 System Model

Consider the Massive MIMO system shown in Figure 4.1. Bidirectional communication over the wireless channel is supported in uplink (user $\rightarrow$ basestaion) and downlink (basestation $\rightarrow$ user) directions. In this system, N_{BS} basestation antennas communicate with the K simultaneous, single-antenna users. Received signals in the downlink and uplink directions, $\mathbf{y}_{\mathrm{d}}$ and $\mathbf{y}_{\mathrm{u}}$, can be modeled as

$$\mathbf{y}_{\mathrm{d}} = \sqrt{\rho_{\mathrm{d}}}\mathbf{G}\mathbf{s}_{\mathrm{d}} + \mathbf{w}_{\mathrm{d}} \tag{4.1}$$

$$\mathbf{y}_{\mathrm{u}} = \sqrt{\rho_{\mathrm{u}}}\mathbf{G}^{\mathrm{T}}\mathbf{s}_{\mathrm{u}} + \mathbf{w}_{\mathrm{u}}. \tag{4.2}$$

Here, the $\mathbf{G}$ represents the complex-valued channel matrix that is the product of slow and fast-varying components and is of $K \times N_{BS}$ dimensions. $\mathbf{G}$ can be expressed as a product of $K \times N_{BS}$ matrix $\mathbf{H}$ that accounts for small scale fading (changes over intervals of a wavelength or less) and a $K \times K$ matrix $\mathbf{D}_{\eta}^{1/2}$, whose diagonal elements constitute a $K \times 1$ vector η of large-scale fading coefficients over the wireless path to each user,

$$\mathbf{G} = \mathbf{D}_{\eta}^{1/2}\mathbf{H}. \tag{4.3}$$

Here, the vector η accounts for both path loss and shadowing effects. As typical in MIMO systems, we assume a power constraint such that $\mathbb{E}\left\{||\mathbf{s}_{\mathrm{d}}||^2\right\} = \mathbb{E}\left\{||\mathbf{s}_{\mathrm{u}}||^2\right\} = 1$. The quantities ρ_{d} and ρ_{u} are proportional to the ratio of power divided by the noise variance, where $\mathbf{w}_{\mathrm{d}}$ and $\mathbf{w}_{\mathrm{u}}$ are the vectors of noise with independent samples drawn from a $\mathcal{CN}(0,1)$ distribution.

Many options are available for precoding the downlink signal $\mathbf{s}_{\mathrm{d}}$ and include linear and non-linear techniques. Linear precoding generates $\mathbf{s}_{\mathrm{d}}$ as

$$\mathbf{s}_{\mathrm{d}} = \mathbf{F}\mathbf{x}_{\mathrm{d}}. \tag{4.4}$$

Here, the downlink signal $\mathbf{s}_{\mathrm{d}}$ is formed by applying the downlink precoder $\mathbf{F}$ to the signals x_{d}^k intended for each of the K users, where $\mathbf{x}_{\mathrm{d}} \in \mathbb{C}^{K\times 1}$ and $\mathbb{E}\left\{||\mathbf{x}_{\mathrm{d}}||^2\right\} = 1$. Particular to this system, I consider zero-forcing (ZF) and minimum-mean squared-error (MMSE) precoders:

$$\mathbf{F}_{\mathrm{ZF}} = \alpha_{\mathrm{ZF}}\mathbf{G}^{\dagger} = \mathbf{G}^H\left(\mathbf{G}\mathbf{G}^H\right)^{-1}, \tag{4.5}$$

$$\mathbf{F}_{\mathrm{MMSE}} = \alpha_{\mathrm{MMSE}}\mathbf{G}^H\left(\mathbf{G}\mathbf{G}^H + \beta\mathbf{I}\right)^{-1}. \tag{4.6}$$

Here, β is a regularization factor applied to improve the quality of the inverse in rank-deficient channels. Scale factors α_{ZF} and α_{MMSE} are applied to ensure that the uplink power constraint is met—i.e. $\mathbb{E}\left\{\mathrm{tr}\left(\mathbf{F}_{\mathrm{ZF}}\mathbf{F}_{\mathrm{ZF}}{}^H\right)\right\} = \mathbb{E}\left\{\mathrm{tr}\left(\mathbf{F}_{\mathrm{MMSE}}\mathbf{F}_{\mathrm{MMSE}}{}^H\right)\right\} = 1$.

Different precoding strategies can be applied to maximize the sum rate of the system. The sum rate $\mathcal{R}$ achieved by the multi-user channel is

$$\mathcal{R}\left(\rho_{\mathrm{d}}, \mathbf{G}, \gamma\right) = \log_2 \det\left(\mathbf{I}_{N_{BS}} + \rho_{\mathrm{d}}\mathbf{G}\mathbf{D}_{\gamma}\mathbf{G}^H\right) \tag{4.7}$$

and can be maximized by solving the constrained optimization problem

$$\gamma^{\star} = \underset{\gamma \in \mathbb{R}^{N_{BS} \times 1}}{\arg\max}\, \mathcal{R}\left(\rho_{\mathrm{d}}, \mathbf{G}, \gamma\right) \quad \text{where} \quad \sum_{i=0}^{N_{BS}-1} \gamma_i = 1, \quad \gamma_i \geq 0 \ \forall i \tag{4.8}$$

Under Massive MIMO limits, where $N_{BS} \to \infty$, the columns of $\mathbf{G}$ become nearly orthogonal [72]. This allows a much simpler precoding method known as maximum ratio transmission (MRT) to be applied while still achieving maximum sum rate in the asymptotic limit [73]:

$$\mathbf{F}_{\mathrm{MRT}} = \alpha_{\mathrm{MRT}} \mathbf{G}^{H} \tag{4.9}$$

This approximation eliminates the need for matrix inversion and reduces complexity of downlink precoding. Precoded downlink symbols can be computed by normalizing for the number of basestation antennas N_{BS}, the large-scale fading coefficients η, and the channel $\mathbf{G}$ as

$$\mathbf{s}_{\mathrm{d}} = \frac{1}{\sqrt{N_{BS}} \mathrm{tr}\left(\mathbf{D}_{\eta}\right)} \mathbf{D}_{\eta}^{-1/2} \mathbf{G}^{H} \mathbf{x}_{\mathrm{d}}. \tag{4.10}$$

The MRT precoder is preferred due to its ease of implementation; however, benefits of ZF and MMSE are seen as long as the number of antennas remains finite. Namely, ZF is cited for its inter-user interference suppression capabilities and MMSE is cited for better performance in low rank channels (the probability of which are vanishingly small as N_{BS} grows) [72, 76]. Applying these precoders in Massive MIMO simulations have led to impressive gains, with sum rates in excess of 100 bits per channel use [73, 74]. In addition to increased spectral efficiencies, Massive MIMO has also demonstrated decreased interference, gains in radiated energy efficiencies, and reduced requirements per RF channel by an order of magnitude or more versus conventional MIMO.

4.2.2 Channel State Acquisition and Processing

In this system, channel reciprocity allows $\mathbf{G}$ to be estimated at the basestation using uplink pilots. This estimate, $\hat{\mathbf{G}}$ can be used to to compute a precoder for the downlink channel of the forms in provided in (4.5), (4.6), and (4.10) so long as the channel state information at the transmitter (CSIT) remains valid between its acquisition and use. In particular, this condition requires the next downlink signaling interval to occur within the coherence time of the last uplink channel measurement. For LTE cellular systems, a coherence time of ~1 ms is typical as listed in Table 2.2.

Typical processing for a Massive MIMO basestation is shown in Figure 4.2. In this diagram, the basestation processes N_{BS} channels that consist of N_{BS} antenna ports connected to N_{BS} OFDM transceivers. The entire processing flow is as follows. First, uplink signals are demodulated from RF and digitized using an analog-to-digital (A/D) converter and associated RF hardware. Second, the digital samples are then downconverted to the desired sampling rate that is some fraction $\frac{1}{S}$ times the A/D sample rate. The typical OFDM signal processing chain is then applied—i.e. the cylic prefix (CP) is removed, the FFT is taken after serial-to-parallel conversion, guard subcarriers are removed, and uplink data and pilots are deallocated according to the resource map, or schedule, observed by the link.

Pilot and data symbols for all N_{BS} receive chains are then passed to the channel estimator. Channel estimates are then used to decode the K uplink data streams and precode the K downlink data streams. The K downlink

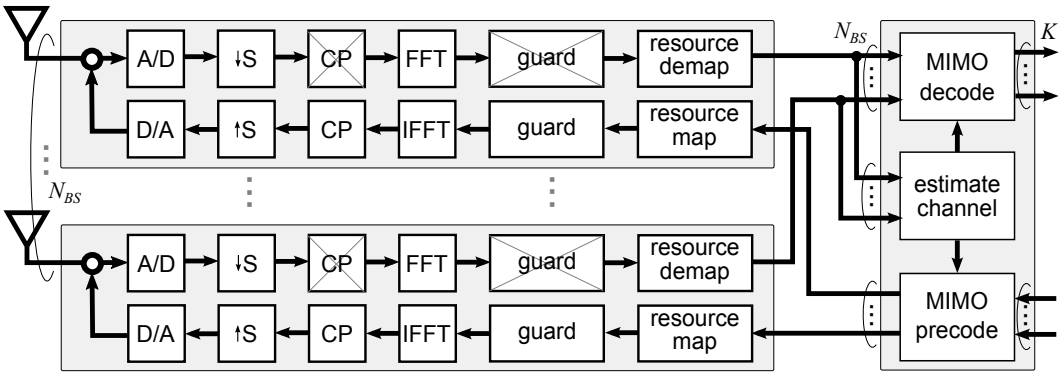

Figure 4.2: Processing performed by an OFDM Massive MIMO basestation. N_{BS} antennas of synchronized uplink baseband samples are acquired by an analog-to-digital (A/D) converter and processed using and OFDM receiver (CP removal, FFT, guard subcarrier removal, resource demap) then passed to a MIMO detector and channel estimator. Channel estimates are used to precode downlink data. Precoded symbols are then distributed to the N_{BS} OFDM tansmitters (resource map, guard subcarrier insertion, IFFT, add CP) and transmitted out the antenna ports.

data streams are precoded using the ZF precoder shown in (4.5) or the MRT result shown in (4.10). Precoded data for the N_{BS} antennas is then input to a resource mapper, guard symbols are inserted, serial-to-parallel conversion takes place, the IFFT is taken, CP is added, samples are digitally upconverted by S times the baseband rate, and digital samples are converted to analog using a digital-to-analog (D/A) converter and conditioned for transmission out of the basestation antenna ports.

Though this signal processing flow is standard in reciprocity-based MIMO, challenges arise in scaling N_{BS} beyond conventional systems (see Table 4.1). Practical Massive MIMO requires $N_{BS} \geq 64$ for the asymptotic results to begin to apply, imposing a high cost to prototype such systems.

4.2.3 Prior Work

Impressive theoretical results have prompted interest in experimentally validating Massive MIMO. Ideally, the processing shown in Figure 4.2 could be specified and readily deployed in hardware. Unfortunately, scaling antennas and data requirements an order of magnitude or more beyond classical systems (see Table 4.1) has led to challenges in availability of prototyping hardware. Large equipment vendors and infrastructure companies have invested in this area, particularly with the upcoming work on 5G wireless definition and standardization, of which, Massive MIMO has been named a key candidate. Though many researchers have remained secretive on their internal plans and prototyping efforts, work from several groups has been made public.

Table 4.3 lists existing many-antenna testbeds as of writing. The first system is a channel sounding system used at Lund University [73]. This testbed is used to measure the wireless channel with a large number of antennas to validate theoretical gains. Measurements were taken over long duration (hours) where channel is assumed constant and post-processed. Their results confirm that a large fraction of theoretical gains can be harvested in practice, granted that a system is capable of reciprocity and real-time performance. Second, researchers at Rice University [77] have constructed a testbed to perform Massive MIMO processing in real time. In [77], they present experimental results based on SINR measured at UE w/ high latency (100 ms) beamforming over 0.625 MHz of bandwidth. They use these measurements to compute channel capacity and confirm the theoretical gains. They mention that they are currently

Table 4.3: Existing many-antenna testbeds

Group	Band (GHz)	Hardware Platform	Number of Antennas at Basestation	Number of Users	Real-Time MIMO Processing?
Lund University [73]	2.6	Network Analyzer	128 cylindrical array	6	No
Rice University [77]	2.4	WARP boards, PowerPC	8×8 = 64 planar array	15	No
Samsung [78]	28	Proprietary w/ Freescale DSPs	8×8 = 64 planar array	1	Yes
Nutaq and Sherbooke Université [79]	0.3-3.8	25 Virtex-6 MicroTCA boards	100	?	?

working on a lower latency, higher bandwidth system. Third, researchers at Samsung [78] made public their work in many-antenna MIMO systems. Their testbed is targeted at millimeter wave bands but can be applied to cellular band applications. Unfortunately, many details about this system are omitted in the press release, though it is mentioned that a throughput of 1 Gbps is achieved at 2 km range. Fourth, Nutaq publicly announced a partnership with the Université de Sherbrooke to prototype Massive MIMO. Their website literature lists support for up to 100 antennas using 25 Xilinx Virtex-6 FPGA boards, and one press release mentions a 16-antenna prototype [79]. Real-time processing is likely supported but not explicitly discussed.

Despite prior work in large scale MIMO prototyping, many shortcomings are evident. Namely, existing testbeds are either proprietary, non-real-time, or both. This severely impedes researchers mapping theory to practice.

4.3 128-Antenna MIMO Platform

To address this limitation, I have developed an extensible platform to realize up to 20 MHz bandwidth 128-antenna MIMO. This platform consists of exclusively commercial off-the-shelf hardware, making it accessible and modifiable. Key challenges of this system and thus my contributions include (*i*) scaling data rates and interfaces to 128 antennas, (*ii*) providing low-latency data transfer to support high-mobility, and (*iii*) performing time and frequency synchronization over 128 antennas, all with readily available commercial off-the-shelf components. This platform forms the basis of LuMaMi, the world's first 100-antenna real-time Massive MIMO system [8].

4.3.1 Scaling to 128 Antennas

To perform uplink MIMO at the basestation, data from all N_{BS} antennas must be aggregated into the MIMO decoder as shown in Figure 4.2. Likewise, in the downlink direction, K streams of source data must be precoded into subcarriers for N_{BS} antennas. For an OFDM-based system, this results in a peak aggregate data rate in uplink and downlink directions of

$$R_{tot} = N_{BS} \sum_{i=0}^{N_S-1} \frac{N_{SC}^D}{N_{FFT} + L_i} f_s \quad \text{samples/s.} \tag{4.11}$$

This peak value assumes the schedule is fully allocated in uplink and downlink. Designing a system based on peak rate allows support for arbitrary frame schedules defined in Section 4.3.2. Table 4.4 lists the system parameters of the 128-antenna platform. The system is designed to support up to 10 simultaneous users via up to 128-antenna basestation. The signal structure follows that

of 20 MHz LTE [3] with a 30.72 MS/s baseband sample rate, 2048 FFT size, and 1200 occupied subcarriers. The system deviates from LTE in that OFDM signaling is used in both the uplink and downlink. To allow for reciprocity-based MIMO, received samples must be processed, uplink channels must be measured, and downlink signals must be precoded and transmitted within the coherence time of the channel, taken to be 1 ms for mobile LTE [3]. The processing time between the reception of uplink pilot signals, extraction of pilots, channel estimation, precoding, and transmission is known as T_t, uplink pilot receive to downlink transmit turnaround time, and is specified to be less than 500 μs to ensure that the channel used for the precoder is sufficiently fresh.

Using the parameters in Table 4.4 and (4.11) results in an aggregated rate of 2.15 GS/s for subcarriers in the uplink and downlink directions. For practical data types of 32 bits per complex sample (16 bits for real and imaginary component) results in 68.8 Gbits/s in each direction that need to be routed and processed in the system. This high throughput and the tight latency constraint necessitate dedicated hardware processing. For prototyping purposes, the best fit is a field-programmable gate array (FPGA), which allows reconfigurability, high throughput, and low-latency processing in hardware using a sea of reconfigurable hardware elements including random access memories (RAMs), digital signal processing blocks (DSP), look up tables (LUTs), registers, as well as digital input/output (IO).

Dedicated hardware resources are used to meet throughput and latency requirements of Table 4.4. Unfortunately, large MIMO detectors and many

channels of processing can easily exceed the size of many current state-of-the-art FPGAs. For reference, a reference design produced by Xilinx achieves 2×2 MIMO over 20 MHz with the full LTE physical layer for uplink and downlink. Even with only 2×2 MIMO, the resource usage is 544 DSP blocks, 151,800 LUTs, and 923 18 kbit block RAMs. This is an appreciable percentage of even large FPGA targets such as the Kintex-7 XC7K410T, where this translates to 544/1540 = 35.3% of available DSPs, 151,800/254,200 = 59.7% of available LUTs, and 923/1590 = 58.1% of available 18 kbit block RAMs [80]. To meet system rate requirements and finite resource constraints, the signal processing is broken into subbands since after OFDM processing, the signals are separable and mutually orthogonal in frequency.

Current software defined radio (SDR) platforms provide a user configurable processing element (an FPGA in some cases), reconfigurable RF transceivers, and a communication bus to transfer data to/from the SDR device. Distributing MIMO processing in this manner requires the formation of a hierarchical network that interconnects processing elements of SDR devices. Figure 4.3 shows a hierarchical overview of a scalable distributed MIMO system. In this system, a central controller manages distributed processing across a high speed digital communication network formed between processing elements within SDRs. Hubs allow switching and routing signals over the network.

Many options exist for commercial, off-the-shelf SDR solutions. The National Instruments (NI) USRP-2943R is a good choice because it features

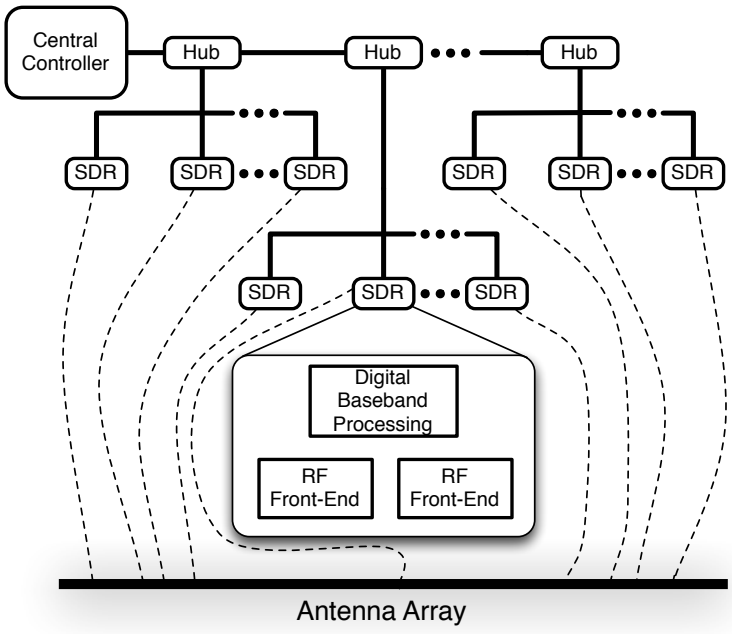

Figure 4.3: The hierarchical overview of an SDR-based distributed Massive MIMO system. A central controller manages distributed processing across a high speed digital communication network formed between SDR processing elements. Hubs allow switching and routing signals over the network [8].

an onboard Kintex-7 XC7K410T, 2 channels of 40 MHz RF bandwidth tunable from 1.2 to 6 GHz, and a Gen 1 x4 (830 MB/s per direction) PCI-Expess Interface. An additional benefit of this device is that the full system can be constructed using parts from one vendor (NI) and programmed using a common tool and design flow (LabVIEW). This contrasts other SDR platforms that are marked by disparate hardware and/or software design flows. Table 4.6 summarizes the relevant specifications of the USRP-2943R SDR. Notably, a high-speed direct memory access (DMA) interface is used to move data between SDRs and across the processing network without CPU intervention.

Using this SDR, a compatible hub is the NI PXIe-1085 chassis, an instrumentation chassis that provides a backplane of Gen 2 x8 PCI-Express switches that support up to 3.2 GB/s bidirectional bandwidth per switch (there are two per chassis). The chassis is interfaced to the SDRs using a cabled PCI-Express (MXI) interface card. The controller is chosen to be the RMC-8354 which provides a high-bandwidth Gen 2 x16 PCI-Express interface that supports up to 5.6 GB/s transfers. The chassis are linked with a Gen 2 x16 PCI-Express interface. The parts list for N_{BS} antennas is shown in Table 4.5.

Working within the constraints listed in Table 4.6 using the architecture shown in Figure 4.3, the processing is subdivided and mapped into groups of eight SDR (eight FPGA) subsystems, each of which form a modular unit that of 16 antennas that can be linked to make up to a 128-antenna system. The process of dividing bandwidth is further distributed in order to meet DMA channel and IO limitations identified in Table 4.6. Signal processing is then

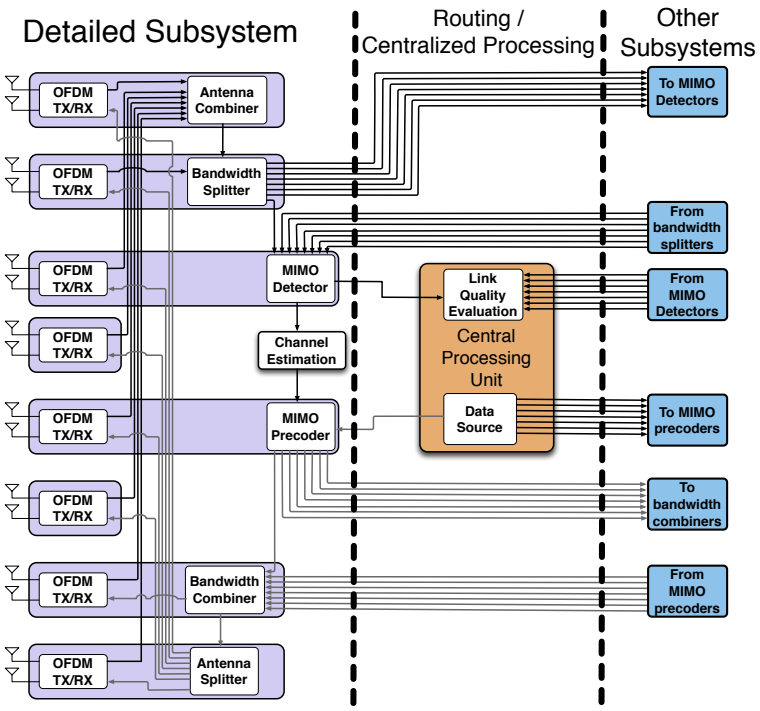

Figure 4.4: An 8-FPGA (16-antenna) subsystem used to build up to a 128-antenna system. Each arc represents data paths inter- or intra-FPGA. Inter-FPGA paths use a DMA interface implemented over PXI-Express.

mapped to FPGAs and the controller as shown in Figure 4.4. Each blue block corresponds to one of the eight FPGAs within an eight-FPGA (16-antenna) subsystem. The red block represents the controller, and gray blocks represent processing that performed within and outside of the subsystem. Processing taking place outside the subsystem is not encapsulated by red or blue area. Each arc between blocks represents inter- or intra-FPGA data and their respectable directions that must be passed between processing blocks. More details about this implementation are provided in [81].

Table 4.4: Specifications for Massive MIMO 128-antenna testbed

Parameter	Variable	Value
# of BS antennas	N_{BS}	128
# of UEs	N_{UE}	1-10
Baseband sample rate (MS/s)	f_s	30.72
FFT size	N_{FFT}	2048
Data subcarriers (excluding DC)	N_{SC}^{D}	1200
Guard subcarriers	N_{SC}^{G}	848
Occupied channel bandwidth (MHz)	B	18.015
OFDM symbols per radio slot	N_S	7
Uplink symbols per radio slot	N_{UL}	0-7
Downlink symbols per radio slot	N_{DL}	$\leq 7 - N_{UL}$
Modulation	—	OFDMA (downlink) OFDMA (uplink)
Length of CP	L_i	160 (symbol 0), 144 (symbols 1-6)
Uplink pilot receive to downlink transmit turnaround time	T_t	$<500\mu$s

Table 4.5: Basestation hardware for N_{BS}-antenna testbed

Description	Part No.	N_{BS}		
		64	100	128
NI USRP-2943R 2-Antenna SDR w/ Xilinx Kintex-7 FPGA (1.2 - 6GHz)	783148-01	32	50	64
NI PXIe-8374, x4, MXI-Express Daisy-Chain Copper Interface	781820-04	32	50	64
x4 MXI cables (7m)	779500-07	32	50	64
NI PXIe-1085, 18-Slot 3U PXI Express Chassis	781813-01	2	3	4
NI PCIe-8388, x16 Gen2 MXI-Express for PXI Express Interface	781762-01	1	1	1
NI PXIe-8388,1 Port, MXI-Express for PXI Express Interface	781760-01	1	1	1
NI PXIe-8389, 2 Port, MXI-Express for PXI Express Interface	781760-01	1	2	3
x16 MXI-express Cables (3m, 170 for 1m)	781763-03	2	3	4
NI RMC-8354 1U Controller, Core i7-860, Win7, 64-Bit	781652-04	1	1	1
Octoclock 8-Channel Clock Distribution System	782978-01	4	7	8
Octoclock 8-Channel Clock Distribution System w/ GPSDO	781037-01	1	1	1
PXIe-6674T Timing and Synchronization Module	781037-01	1	1	1

Table 4.6: Detailed specifications for USRP-2943R software defined radio.

System Parameter	Value
Center frequency	1.2 - 6.0 GHz
RF bandwidth	40 MHz
Number of RF channels per device	2
ADC sampling rate	120 MS/s
ADC resolution	14 bit
DAC sampling rate	400 MS/s
DAC resolution	16 bit
Onboard FPGA	Xilinx Kintex-7 XC7K410T
Digital backend interface	PCI-Express Gen 1 x4
Max bidirectional peer-to-peer bandwidth per device	830 MB/s
Max unique DMA endpoints per device	16

4.3.1.1 Parameterized SDF Model of 128-Antenna Basestation

To further analyze processing dependencies and aid in mapping the system to hardware, the processing graph is modeled using a parametrized synchronous dataflow (PSDF) representation. A PSDF model allows the analysis of processing dependencies and the determination of a quasi-static schedule for a given parameter space. PSDF is more flexible than synchronous dataflow (SDF) since it allows dynamic, run-time configurable parameters. In addition, the quasi-static schedule determined from PSDF allows proper sizing of buffers to guarantee that overflows and/or underflows do not occur. In the PSDF model of the eight-FPGA subsystem, parameterization is limited to the number of uplink and downlink subcarriers generated per slot N_{UL} and N_{DL} as listed in Table 4.4. The schedule that N_{UL} and N_{DL} are derived from can vary at run-time, so provisions must be made to accommodate a range of values. This PSDF model can be extended across all subsystems.

Figure 4.5 shows a parameterized dataflow model of that same 8-FPGA subsystem shown in 4.4. As shown in the model, each PCIe-2943R processes two channels of antenna data. Each token count corresponds to a two-antenna pair of complex baseband samples processed according to signaling rates listed in Table 4.6. The tasks common to each FPGA in one subsystem according to Figure 4.5 correspond to (A) downsampling from 120 MS/s to 30.72 MS/s, (B) performing FFT and CP and guard subcarrier removal, (C) deallocating N_{UL} scheduled uplink symbols (including uplink pilot symbol(s)), (M) allocating N_{DL} downlink data symbols, (N) IFFT and CP and guard subcarrier

addition, and (O) upsampling from 30.72 MS/s to 120 MS/s.

Not common to each FPGA are distributed MIMO processing tasks that are modeled in Figure 4.5. These tasks take subcarrier samples from each FPGA and perform the following tasks. In the antenna combiner (D), samples from FPGA nodes 0,2-7 are aggregated into a 7× rate stream that is input to (E). In the bandwidth splitter (E), these samples are then split across bandwidth into subbands for distributed MIMO processing. A typical subband distribution (and that reflected in Figure 4.5) is a 1/8 split. These samples are then transferred to 8 channel estimators/MIMO detectors (F,G), one of which lies within the i^{th} subsystem (F_i,G_i) and seven lie outside the i^{th} subsystem ($\text{F}_{j\neq i}$,$\text{G}_{j\neq i}$). The MIMO detector produces K streams of output samples for N_{BS} antenna samples.

Channel estimates from (F) are then transferred to (J_i), which performs MIMO precoding over a subband for the i^{th} FPGA subsystem using symbols from the host for K streams. The bandwidth combiner (K), takes in eight 100 MB/s streams to generate transmit subcarriers for the eight FPGA susbsytem. These subcarriers are then transferred to (L), the antenna splitter, and directly to the transmitter of FPGA 6 of the subsystem. Samples sources from the x64 controller are downlink data along with controls and indicators for status monitoring the communication link. Data consumed by the x64 controller include uplink data destined for basestation consumption.

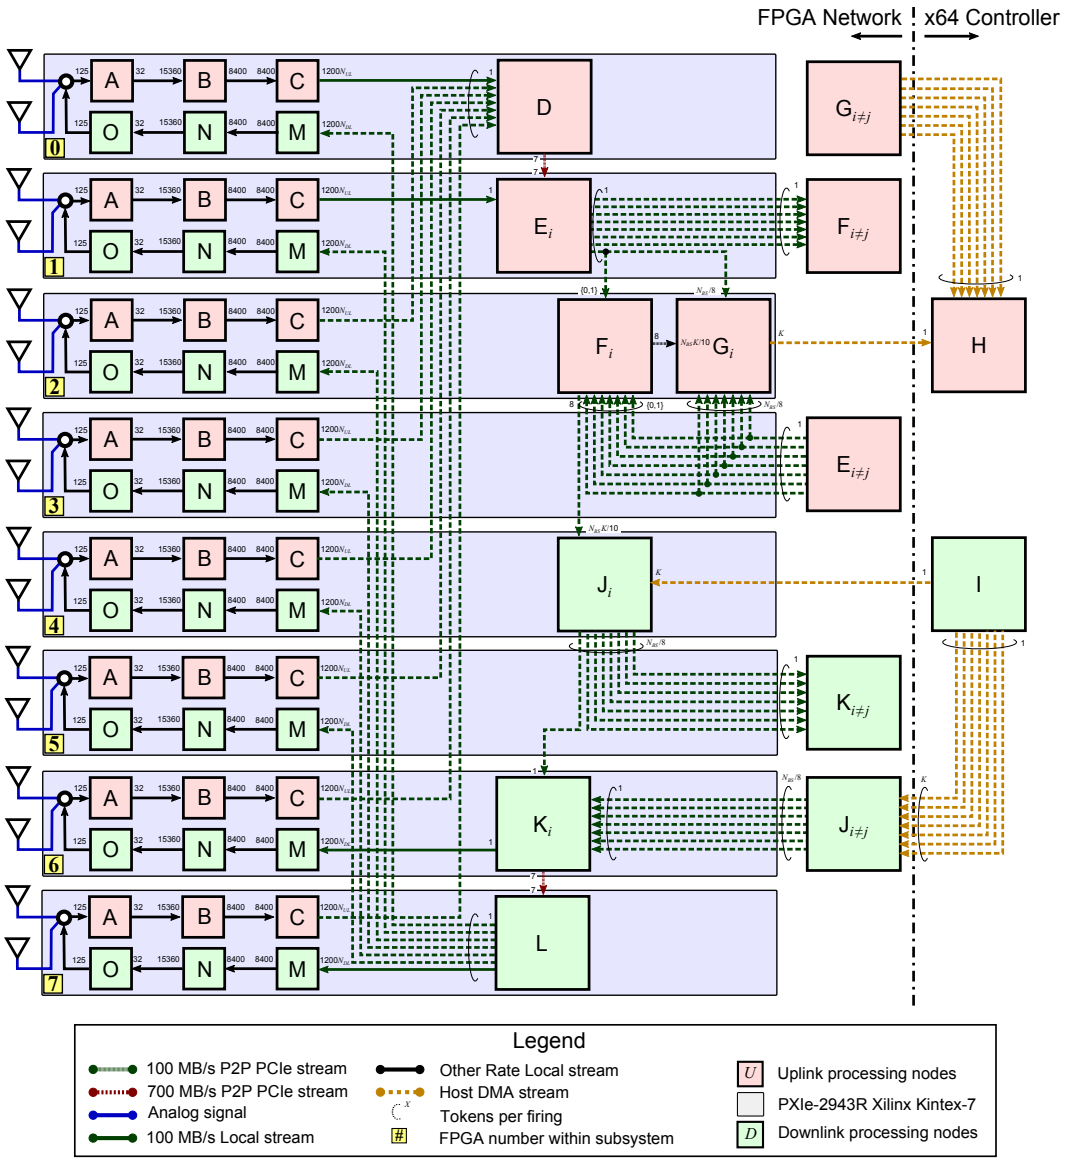

Figure 4.5: Parametrized synchronous dataflow model of the i^{th} 8-FPGA (16-antenna) subsystem with firings per node indicated by black numerical script. In the critical signal path, uplink processes ABC are executed in each FPGA, providing samples to process chain DEFGHIJKL which provide samples to MNO in the transmit direction. The number of uplink and downlink symbols is parameterized.

Table 4.7: Final DMA usage and IO rates for 128-antenna platform.

FPGA	DMA usage				approx. IO rates (MB/s)	
	in	out	RF	total	in	out
0	7	1	3	11	705.6	705.6
1	7	2	3	12	806.4	806.4
2	1	1	3	5	100.8	100.8
3	9	3	3	15	806.4	250
4	3	9	3	15	250	806.4
5	1	1	3	5	100.8	100.8
6	2	7	3	12	806.4	806.4
7	1	7	3	11	705.6	705.6

The peak rate for each stream of subcarriers per antenna in bytes/s is

$$R^{B/s}_{antenna} = W^{B}_{sample}\frac{R_{tot}}{N_{BS}} = W^{B}_{sample}\sum_{i=0}^{N_S-1}\frac{N^{D}_{SC}}{N_{FFT}+L_i}f_s \quad \text{bytes/s.} \qquad (4.12)$$

for a word length in bytes per sample of W^{B}_{sample}. Using $W^{B}_{sample} = 3$ bytes/sample (12 bits for real and imaginary part) and system parameters defined in Table 4.4, this implies peak rates of 50.4 MB/s per antenna and 100.8 MB/s per USRP-2943R. This balances IO limitations in the graph. Samples in the actual system are transferred across interfaces in blocks of 4 since the bus is most efficient operating with data units that can fit evenly within 64-bit chunks.

Table 4.7 shows the final DMA usage and IO rates for the 128-antenna platform according to rates shown in Figure 4.5. All values are within the limits defined in Table 4.6.

4.3.2 Latency Minimization

To meet channel coherence requirements with a $T_t \leq 500$ μs, uplink downlink signaling must be modified from that of LTE, which provisions TDD

configurations with switching intervals of 5-10 ms [3]. To meet this requirement, a frame structure is proposed that has a TDD period of 0.5 ms. As shown in Figure 4.6, the 10 ms radio frame is subdivided into 0.5 ms slots, of which there are two types—SYNC and DATA. The SYNC slot type is a special frame that provides a synchronization signal for UEs to synchronize to the basestation. The DATA slot type contains an uplink and a downlnk interval separated by guard bands. The first OFDM symbol in this slot is an uplink pilot signal that contains orthogonal pilots for each of K users. Each user transmits only in his subcarriers. The critical latency path is from the reception of the uplink pilots to when the first precoded downlink signals need to be transmitted. In this figure, this is shown to be 3 OFDM symbols, or 214 μs. The coherence interval spans from the start of the pilot until the end of the last downlink symbol, or 6 OFDM symbols (429 μs).

The processing and signal routing modeled in Figure 4.5 has been architected to provide low-latency in the signal path from channel estimate to MIMO precoder shown in Figure 4.2. This turnaround time must meet the frame schedule shown in Figure 4.6. This leaves 214 μs for total latency in the critical path Δ, which is the sum of latencies elaborated in (4.13). These latencies include RX front-end delay Δ_f^{rx}, OFDM RX (CP removal, FFT, guard subcarrier removal) Δ_o^{rx}, channel estimate calculation Δ_e, precoder calculation Δ_p, OFDM TX (guard subcarrier interleave, IFFT, CP addition) Δ_o^{tx}, and TX front-end delay Δ_f^{tx}. Additional sources of latency include data routing, packing, and unpacking durations ϕ as well as latency for each hop across the PCIe

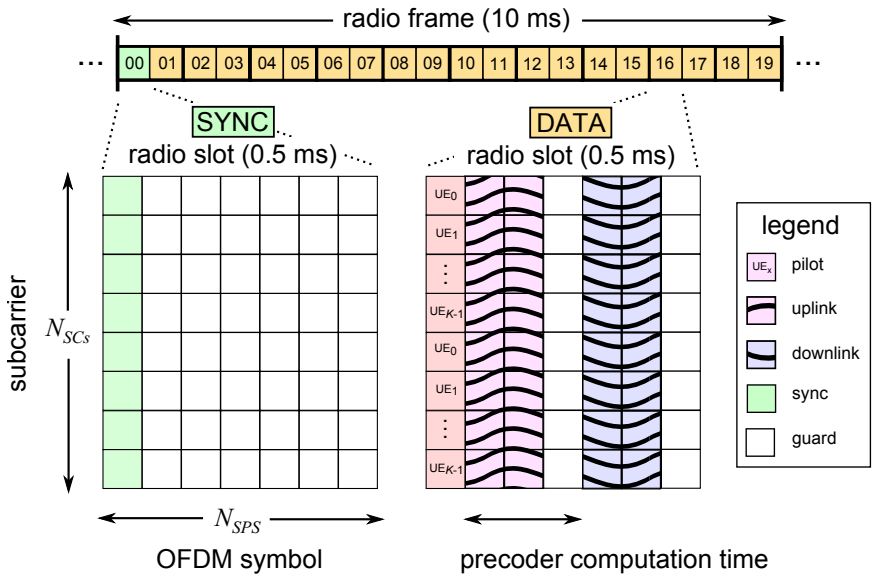

Figure 4.6: TDD radio subframe structure that supports channel reciprocity for even highly mobile channels. In a SYNC radio slot, one OFDM symbol is used for a downlink synchronization signal. The following six OFDM symbols are blank, allowing for sample shifting to align the clocks. In a DATA slot, the first OFDM symbol are uplink pilots (different subcarriers for each UE) with uplink data following shortly thereafter. Computation for the downlink precoder takes place over OFDM symbols 1-3 (indexing by zero).

backplane $N_h\Delta_h$:

$$\Delta = \Delta_f^{rx} + \Delta_o^{rx} + \Delta_e + \Delta_p + \Delta_o^{tx} + \Delta_f^{tx} + N_h\Delta_h + \phi. \tag{4.13}$$

The worst-case latency of each hop is $\Delta_h = 5$ μs for the seven hop path ($N_h = 7$), resulting in a worst-case total PCIe latency of $N_h\Delta_h = 35$ μs in the critical signal path. A high degree of pipelining is used among blocks to ensure that the frame schedule can be met. The MRT precoder can be processed point-by-point, allowing processing to be fully pipelined. Similarly, channel estimation can be performed point-by-point. The highest latency configuration is for the ZF precoder due to the matrix inverse, matrix-matrix multiplies, and its internal serial-to-parallel conversions. An implementation of high speed signal routing for pipelined data transfer is shown in Figure 4.7.

4.3.2.1 Reconfigurable 1.28 GB/s Hardware Data Router

Figure 4.7 shows a parameterizable data router used to implement nodes D, E, K, and L in Figure 4.5. This router is able to transfer one 64-bit element from one of eight possible sources to one of eight possible destination FIFOs. These FIFOs can be DMA (for peer-to-peer transfers over PCI-Express), local FIFOs, or some combination thereof. The routing logic operates at 160 MHz, providing up to 1.28 GB/s of maximum throughput and route schedules are parameterically defined using a block RAM lookup table. Using an *antenna first* data ordering convention provides lowest latency through the critical signal path. Such a convention is enforced via the routing

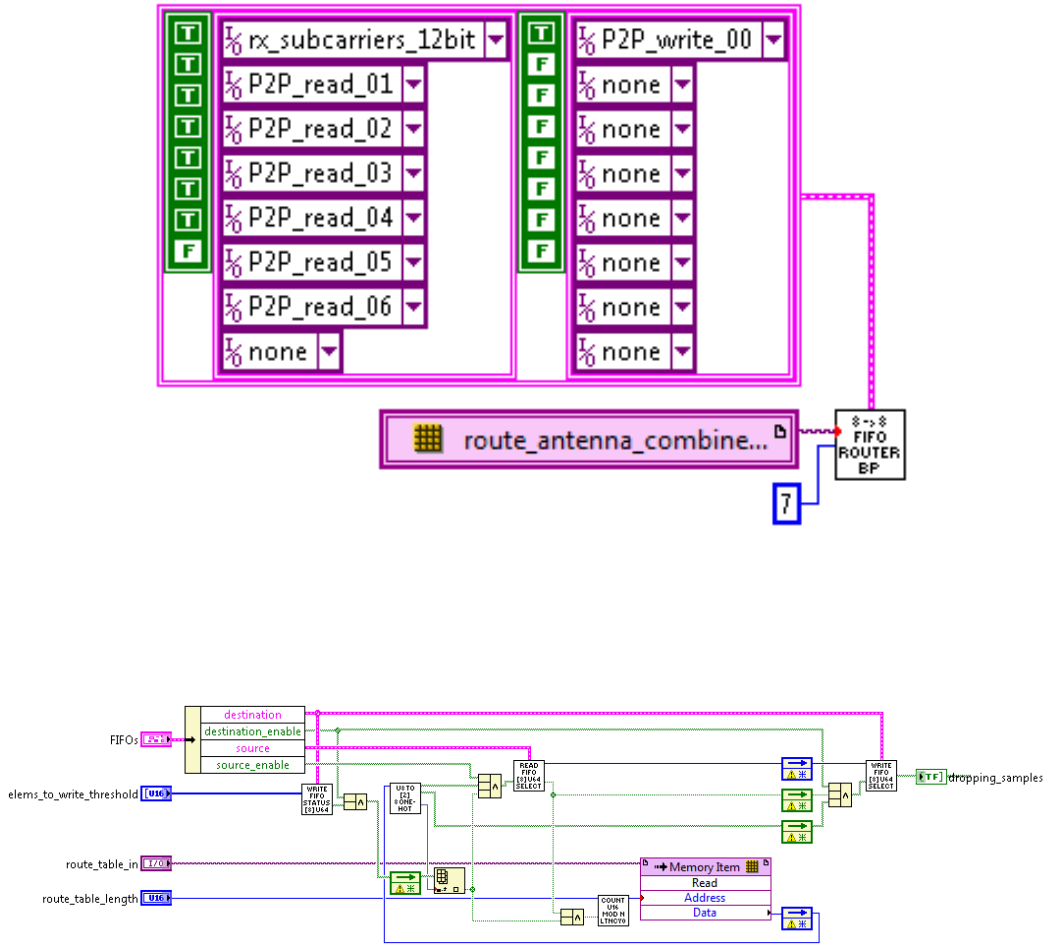

Figure 4.7: Low latency compile-time configured 1.28 GB/s signal router written in LabVIEW. FIFO read and write functions are implemented using one-hot state machines. Every node in the diagram executes in one 160 MHz clock cycle. The router can transfer up to one 64 bit element per clock cycle, providing up to 1.28 GB/s of throughput. Route schedules are parametrically configured using look-up tables implemented in block RAM. This node is configured to implement nodes D, E, K, and L in Figure 4.5. The configuration for D is shown above.

schedule. This method optimizes pipelining for ZF processing, allowing ϕ and ultimately Δ of (4.13) to be minimized.

4.3.3 Synchronization

A MIMO basestation requires time synchronization and phase coherence between each RF chain. This is achieved using a reference clock and trigger distribution network. The reference clocks are used as the source of each radio local oscillator. This provides phase coherence (yet not alignment due to phase-locked-loop uncertainty) among devices. The trigger signal is used to time synchronize initial acquisition and transmission times and synchronize state across FPGAs. State synchronization is ensured by invoking a state reset on all devices prior to the trigger event.

The hardware used to build the clock and trigger distribution network are available commercially off-the-shelf. This consists of eight OctoClock (782978-01) modules in a tree structure with a top OctoClock feeding eight lower OctoClocks (see Figure 4.8). Low skew buffering circuits and matched-length transmission cables ensure that there is low skew between the reference clock input at each USRP-RIO. The source clock for the system is an oven-controlled crystal oscillator within an NI PXIe-6674T timing module. Triggering is achieved by instigating a start pulse within the *Master* USRP-RIO via a software trigger. This trigger is then output from an output port on the master and input to the NI PXIe-6674T, which conditions and amplifies the trigger. The trigger is then propagated to the top OctoClock and distributed

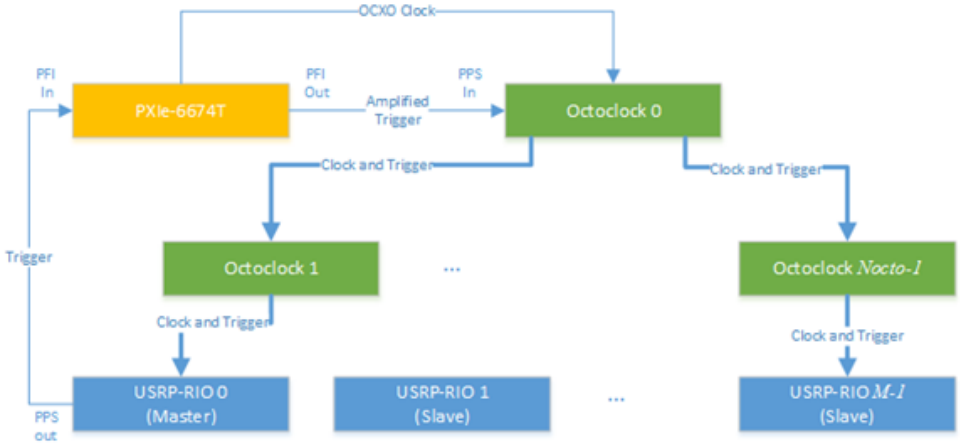

Figure 4.8: Clocks and reference signals are routed to each NI USRP-2943R device using this timing and synchronization tree. The master *USRP-RIO 0* sends a trigger signal that is fed into the PXIe-6674T module. The PXIe-6674T provides a conditioned and amplified trigger along with an OCXO reference clock. These signals are fed to the top Octoclock which fans out these signals to the lower branch of Octoclocks.

down the tree to each USRP in the system (including the master itself). This signal sets the reference clock edge to use for start of acquisition for the transmit and receive within each channel. Initial results show that reference clock skew is within 100 ps and trigger skew is within 1.5 ns. Additional signal alignment calibration is done to digitally compensate for front-end delay. This ensures the TDD signals are properly aligned over-the-air. This further allows the design to be modularly used on top of any radio hardware.

4.3.3.1 64-Antenna Wireless Synchonization Experiment

To confirm the synchonization scheme, a wireless test was conducted using 64 antennas of the LuMaMi testbed [8] and a 16×4 (horizontally long)

subset of its antenna array. The antennas chosen were a subset of the 160 dual-polarized patch antennas that form a full 2-D planar array. A uniformly-spaced planar array was build for ease of construction and compactness. Details about the antenna array are provided in [8]. A transmitter was configured to transmit a phase-coherent Gaussian pseudo noise (PN) sequence of length 8000 samples at 40 MS/s sample rate at a center frequency of 3.7 GHz. The transmitter was placed approximately 1 m from the receive array roughly in line with the normal of the plane of the array. The array consisted of shorted patch antennas but could have also have been constructed using other small antennas [82, 83]. The resulting signal occupied 40 MHz RF bandwidth at 3.7 GHz and was broadcast repeatedly at ~5 dBm transmit power. The 64 channel receiver was configured to simultaneously capture 8000 samples at 40 MS/s using the synchronization scheme outlined above. The raw received signals were then written to a file and post-processed in MATLAB. The processing performed was a circular cross-correlation with the PN sequence. A portion of the result of the magnitude of the circular cross correlation across 64 channels is shown in Figure 4.9.

As shown in the figure, a sharp, nearly planar wavefront is apparent. The leading (left) edge of the wavefront is approximately planar, as expected, with trailing delay spread. The RMS delay spread is approximately three 40 MS/s samples, or 75 ns. This is caused by the simple geometry of the physical arrangement of the test, which included a metal wall behind the transmitter and large windows behind the receiver. Receivers closest to the

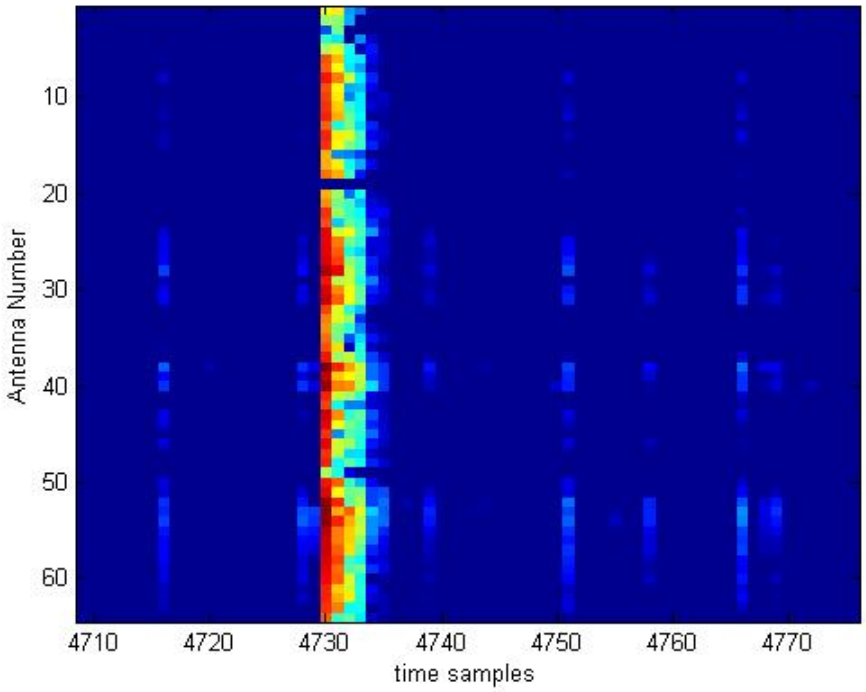

Figure 4.9: Channel magnitude response across 16×4=64 planar array over 40 MHz of bandwidth. Time synchronization is achieved within 25 ns across the 64 antennas.

center of the array horizontally receive the highest signal power with decaying power toward the edges. Antenna 19 has low receive power and was later found to be faulty. Additional details about the test are provided in [8]. This experiment confirmed sample acquisition across the basestation to be synchronized to within 25 ns (the resolution of the test signal).

4.3.3.2 Phase Stability Experiment

To validate the stability of the synchronization over time, a cabled test was conducted using for receivers spanning four lower branches of the clock distribution tree. An input sinusoid was generated by a transmit device and fed to three 1-to-2 RF splitters in a tree configuration to provide four quarter power outputs. The receivers were configured to log measured data every minute over the course of 68 minutes. The logged samples were then post-processed in MATLAB to evaluate their phase stability as a function of time. Figure 4.10 shows the measured phase at each of the four receive channels. As expected, the phases are not aligned due to the phase-coherent RF architecture of the NI USRP-2943R. However, the phases do remain very stable over one hour, showing variation of, at most, five degrees. A majority of the phase drift is caused by temperature variations in the transmit and receive circuits (note correlated drifts at start of measurements). These drifts are independent between transmit and receive circuits. This drift over time must be compensated for using reciprocity calibration, though these results demonstrate that retraining intervals can be relatively long (minutes or hours).

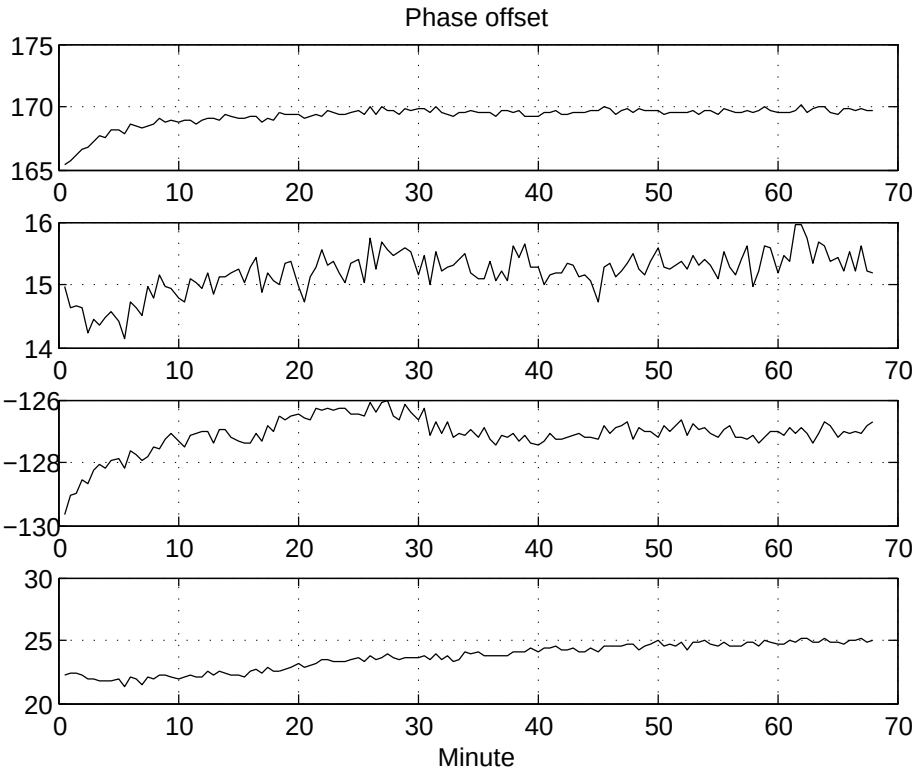

Figure 4.10: The phase stability between transmit and receive as a function of time in degrees. Phase measurements were taken using a cabled signal split four ways to span four branches of the clock and trigger distribution network. Phase drift is within 5 degrees over 1 hour and is primarily caused by variations in temperature (note variations caused at start of measurements).

The measured results for time (see Figure 4.9) and frequency (Figure 4.10) demonstrate that synchronization is within the system is working properly and sample continuity (no loss of samples) is confirmed.

4.4 Conclusion

In order to make Massive MIMO prototyping more realizable, I have developed a scalable platform for large antenna count MIMO using commercial off-the-shelf hardware. The key contributions of this design are (*i*) scaling processing and interfaces to 20 MHz bandwidth over 128 antennas, (*ii*) the development of low-latency routing and frame structure to support reciprocity in high-mobility channels, and (*iii*) the time and frequency synchronization of up to 128 antennas at the basestation. The design is capable of processing up to 409.6 Gbps of raw baseband samples (128 antennas $\times$ 120 MS/s $\times$ 32 bits/S) in both the uplink and downlink directions with a channel estimate to precoder turnaround time of less than 500 μs. The design is reconfigurable and flexible to accommodate various signal structures and system requirements for current and future Massive MIMO prototyping efforts. This platform is used in LuMaMi, the world's first 100-antenna MIMO testbed.

Chapter 5

Conclusions and Future Work

In this dissertation, I apply multi-dimensional signal processing methods to improve communication system performance in underwater acoustic, powerline, and multi-antenna cellular applications. I defend the following thesis statement:

Multi-dimensional signal processing methods can be applied to underwater acoustic, powerline, and multi-antenna cellular to dramatically enhance communications performance without sacrificing real-time requirements.

5.1 Summary of Contributions

In this dissertation, I have presented contributions to underwater acoustic, powerline, and multi-antenna cellular communications. The multi-dimensional signal processing methods I develop are summarized below.

5.1.1 Space-Time Methods for Underwater Acoustic Communications

In Chapter 2, I develop two methods for increased underwater acoustic communication performance in the presence of strong reverberation (self-interference) typical of shallow water channels:

1. a windowed Doppler compensation technique coupled with a multi-stage time-adaptive decision-feedback equalizer.

2. a space-time filtering method coupled with a decision-feedback structure to spatially filter the communication channel.

The first method is used to process more complicated motion effects observed in experimental data. The second method is used to enhance coherent (*desirable*) paths while suppressing incoherent (*undesirable*) paths. These methods are combined to form a robust underwater acoustic communication receiver for shallow water channels that is scalable to large arrays.

Results from a shallow water test environment are presented. Experimental and simulated results agree, demonstrating reduction in interference in excess of 10 dB at the equalizer output. In addition to performance gains, computational complexity is reduced by using the smaller number of pre-processed array outputs of a many-element receiver. Finally, this method is used in conjunction with a multi-user scenario to demonstrate gains in sum spectral efficiencies an order of magnitude above prior state-of-the-art.

5.1.2 Time-Frequency Methods for Powerline Communications

In Chapter 3, I characterize cyclic statistics of powerline noise. I confirm that the noise is predominantly cyclostationary with dominant period equal to half the AC cycle. Other noise such as asynchronous impulsive noise and narrowband noise are observed. I use these results to develop two meth-

ods:

1. an adaptive cyclic modulation and coding scheme method over time and frequency

2. an implementation of an asynchronous impulsive noise mitigation technique.

First, I demonstrate the utility of this adaptive cyclic MCS adaptation using measured noise samples and the G3-PLC MCS set. My results show that this scheme is capable of offering up to a 2× increase in throughput using the current G3-PLC modulation set. Further operating point improvements are observed up to 25 dB for simulated P1901.2 channel models. Second, I design and implement a real-time OFDM receiver with AMP impulsive noise mitigation on FPGAs. My design flow was to (*i*) model the receiver in synchronous dataflow (SDF), (*ii*) convert the receiver to fixed-point data and arithmetic, and (*iii*) synthesize the fixed-point SDF model onto two Xilinx Vertex-5 FPGAs. In the third step, I generalize the SDF receiver model to a globally asynchronous, locally synchronous computational model in LabVIEW DSP Design Module. My OFDM transceiver testbed achieves real-time streaming throughput at G3-PLC rates and mitigates up to 8 dB SNR of impulsive noise.

5.1.3 Space-Time-Frequency Methods for Multi-Antenna Cellular Communications

In Chapter 4, I develop a scalable platform for large antenna count MIMO using commercial off-the-shelf hardware. This platform is used to make

Massive MIMO prototyping more realizable. The key contributions of this design are the development of:

1. scaling processing and interfaces to 20 MHz bandwidth with up to 128 antennas
2. a low-latency routing and frame structure to support reciprocity in high-mobility channels
3. time and frequency synchronization of up to 128 antennas at the base station

The design is capable of processing up to 491.5 Gbps of raw baseband samples (128 antennas $\times$ 120 MS/s $\times$ 32 bits/S) in both the uplink and downlink directions with a channel estimate to precoder turnaround time of less than 500 μs. The design is reconfigurable and flexible to accommodate various signal structures and system requirements for current and future Massive MIMO prototyping efforts. This platform is used in LuMaMi, the world's first 100-antenna MIMO testbed.

5.2 Future Work

I conclude this dissertation with a list of possible topics for follow-on work. These proposed works are extensions of work presented in this dissertation.

- **Develop spatial capacity model for underwater acoustic communications**

  In Chapter 2, I develop an SDMA approach coupled with spatial filtering and a large array receiver to achieve sum spectral efficiencies an order of magnitude above prior state-of-the-art. Eight simultaneous users spread somewhat evenly over 45° in azimuth. Potentially even higher sum rates could be achieved with more users and wider azimuthal spread. This follow-on work could analyze contributions in incoherent reverberation from each user and its effects on the asymtotic behavior of spatial capacity in shallow water underwater acoustic channels.

- **Extend space-time methods to OFDM underwater acoustic communication systems**

  In Chapter 2, I develop space-time methods to process received signal in shallow water channels. I use these methods to increase the time-frequency coherence of received signals and reject incoherent reverberation. These methods enhance SINR of the received signal. In this future work, the appropriate extensions could be made to apply these methods to OFDM-based underwater acoustic communication systems. Scalar weights could be applied to each subcarrier for each monopulse beam output in the frequency domain. This could allow independent nulls to be steeered over each narrowband channel, affording additional flexibility for interference conditions.

- **Develop robust uplink power control mechanism for ACOMMS**

  In Chapter 2, I use spatial division multiple access to increase sum rate spectral efficiencies 10× above prior state-of-the-art. This technique is sensitive to variations in uplink power level from each user, a problem typical in MU-MIMO systems. Future work could develop a robust uplink power control mechanism for underwater acoustic channels.

- **Quantify cyclostationary coherence time of powerline noise**

  In Chapter 3, I develop a method to combat cyclostationary noise using a cyclic MCS map agreed upon by PLC transmitter and receiver. Cyclic coherence is measured and presumed to be constant over an OFDM frame. This map could be used for subsequent frames depending on frame length and the cyclostationary coherence period. In this work, long duration measurements can be used to quantify the cyclostationary coherence times. This can be used to determine re-training intervals needed in PLC protocols adopting this cyclic mapping.

- **Joint impulsive and cyclostationary noise mitigation**

  In Chapter 3, I develop apply two separable methods to address cyclostationary and impulsive noise. More extensive study could be performed on applying these methods together in a practical system. Estimation methods may need to be constructed to estimate background and impulsive noise components over small sub-intervals of the AC cycle. This could allow potentially higher gains to be achieved.

- **Improved time-frequency noise power estimators**

  In Chapter 3, I develop a link adaptation method to address cyclostationary noise. In equation 3.9, the time-frequency noise estimate $\mathbb{S}$ is used. This estimate is obtained in the simulations presented in pgs 76-79 using a time average of noise samples that are measured directly during quiet times in the channel. The estimate could be improved using sideband information in guard bands or by improving convergence of estimates.

- **Apply baseband compression methods to Massive MIMO subcarrier data to increase proposed system beyond 128 antennas and/or 20 MHz bandwidth**

  In Chapter 4, I develop a processing framework for up to 128 antenna MIMO using commercial off-the-shelf components. The primary bottleneck for future scalability is the overall subcarrier bandwidth, particularly the inter-switch rates within the PXI chassis. Further word size reduction could be applied to subcarriers based on methods outlined in [84] that allow SQNR to be maintained while further reducing baseband data rates. Reducing data rates could allow further scaling of the system beyond 128 antennas and/or 20 MHz of bandwidth. Further compression could be achieved by applying similar methods along the spatial dimension.

  Results in the uplink dimension will highly depend on channel

(higher rank = less compressibility) and results in the downlink direction will highly depend on the precoder used. High coherence bandwidth with respect to subcarrier spacing may make this data more compressible.

- **Develop online channel reciprocity calibration method for using 100-antenna testbed**

  Using the processing framework of Chapter 4, a channel reciprocity calibration method could be applied to demonstrate Massive MIMO in real-time. This online calibration method could use coupling between transmitter and receiver, which is typically undesirable to allow estimation of phase differences between transmit and receive. This could be a low-overhead method calibration. The measured phase differences can be applied to form the equivalent of a phase-aligned system even though the hardware currently only supports phase coherency.

- **Evaluate different precoding and decoding methods using 100-antenna testbed**

  The 100-antenna testbed of Chapter 4 can be used to evaluate different MIMO precoders and decoders proposed for Massive MIMO. These algorithms can be mapped to the real-time framework and evaluated over real channels. Results from experimentation can be used to refine the system architecture and tradeoff design parameters for a more optimal system design.

- **Evaluate performance of different antenna array geometries**

The 100-antenna testbed of Chapter 4 is built using a 2-D planar uniformly-spaced array. In practice, a non-uniformly (randomly) spaced antenna array might do better than uniformly-spaced ones, e.g. distributed arrays, due to additional randomness that helps ensure good singular values in the channel matrix. This also helps protect against local shadowing effects on portions of the array. Future work could explore the impacts of array geometries in efforts to improve spatial channel models for Massive MIMO.

Bibliography

[1] G3-PLC, "Open standard for smart grid implementation." http://www.maximintegrated.com/products/powerline/g3-plc/, 2013. xvi, 51, 61, 62, 63, 64, 69

[2] ITUG.9901-9904, "Narrowband orthogonal frequency division multiplexing power line communication transceivers." http://www.itu.int/rec/T-REC-G, 2013. xvi, 51, 61, 62

[3] 3^{rd} Generation Partnership Project, "3GPP Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Access Network (E-UTRA); Base Station Radio Transmission and Reception (Release 12)," 2013. xvi, 14, 105, 116

[4] 3^{rd} Generation Partnership Project, "3GPP Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Access Network (E-UTRA); Base Station Radio Transmission and Reception (Release 10)," 2011. xvi, 95

[5] P1901.2, "IEEE Draft Standard for Low Frequency (less than 500 kHz) Narrow Band Power Line Communications for Smart Grid Applications." http://standards.ieee.org/develop/project/1901.2.html, 2013. xvii, 51, 57, 60, 61, 62, 69, 72, 73

[6] PRIME Alliance, "Interoperable standard for advanced meter management and smart grid." `http://www.prime-alliance.org/`, 2013. xvii, 51, 61, 62

[7] C. R. Johnson and W. A. Sethares, *Telecommunication Breakdown*. Prentice Hall, 2004. 4

[8] J. Vieira, S. Malkowsky, K. F. Nieman, Z. Miers, N. Kundargi, L. Liu, I. Wong, O. Edfors, and F. Tufvesson, "A flexible 100-base station antenna massive MIMO testbed," *Proc. IEEE Global Communications Conference (GlobeCom)*, 2014. accepted for publication. 7, 94, 104, 107, 121, 122, 124

[9] R. E. Francois and G. R. Garrison, "Sound absorption based on ocean measurements 2. boric acid contribution and equation for total absorption," *J. Acoust. Soc. Am.*, vol. 72, no. 6, pp. 1879–1890, 1982. 13

[10] D. B. Kilfoyle and A. B. Baggeroer, "The state of the art in underwater acoustic telemetry," *Oceanic Engineering, IEEE Journal of*, vol. 25, no. 1, pp. 4–27, 2000. 14, 21, 22

[11] A. Ghosh, J. Zhang, J. G. Andrews, and R. Muhamed, *Fundamentals of LTE*. Prentice Hall, 2011. 15

[12] A. P. Kannu and P. Schniter, "Design and analysis of MMSE pilot-aided cyclic-prefixed block transmissions for doubly selective channels," *Signal*

Processing, IEEE Transactions on, vol. 56, no. 3, pp. 1148–1160, 2008. 15, 21

[13] M. J. Lopez, A. C. Singer, S. L. Whitney, and G. S. Edelson, "A DFE coefficient placement algorithm for underwater digital acoustic communications," in *OCEANS '99 MTS/IEEE. Riding the Crest into the 21st Century*, vol. 2, pp. 996–1001 vol.2, 1999. 15, 21

[14] Y. M. Aval and M. Stojanovic, "A method for differentially coherent multichannel processing of acoustic OFDM signals," in *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2012 IEEE 7th*, pp. 73–76, 2012. 15

[15] D. E. Lucani, M. Medard, and M. Stojanovic, "Capacity scaling laws for underwater networks," in *Proc. 42nd Asilomar Conference on Signals, Systems and Computers*, pp. 2125–2129, Oct. 2008. 15, 22

[16] D. W. Ricker, *Echo Signal Processing*. Kluwer Academic Publishers Group, 2003. 16, 26

[17] K. A. Perrine, K. F. Nieman, T. L. Henderson, K. H. Lent, T. J. Brudner, and B. L. Evans, "Doppler estimation and correction for shallow underwater acoustic communications," in *Signals, Systems and Computers (ASILOMAR), 2010 Conference Record of the Forty Fourth Asilomar Conference on*, pp. 746–750, 2010. 17, 19, 26, 30, 47

[18] K. F. Nieman, K. A. Perrine, K. H. Lent, T. L. Henderson, T. J. Brudner, and B. L. Evans, "Multi-stage and sparse equalizer design for communication systems in reverberant underwater channels," in *Signal Processing Systems (SIPS), 2010 IEEE Workshop on*, pp. 374–379, 2010. 17, 19, 24, 31, 47

[19] K. F. Nieman, K. A. Perrine, T. L. Henderson, K. H. Lent, T. J. Brudner, and B. L. Evans, "Wideband monopulse spatial filtering for large receiver arrays for reverberant underwater communication channels," in *OCEANS 2010*, pp. 1–8, 2010. 17, 19, 40, 47

[20] K. F. Nieman, K. A. Perrine, T. L. Henderson, K. H. Lent, and T. J. Brudner, "Sonar array-based acoustic communication receivers with wideband monopulse processing," *USN Journal of Underwater Acoustics*, vol. 61, no. 2, 2011. 17, 19, 47, 49, 50

[21] P.-P. J. Beaujean and L. R. LeBlanc, "Adaptive array processing for high-speed acoustic communication in shallow water," *Oceanic Engineering, IEEE Journal of*, vol. 29, no. 3, pp. 807–823, 2004. 21, 22

[22] L. Freitag, M. Grund, J. Catipovic, D. Nagle, B. Pazol, and J. Glynn, "Acoustic communication with small UUVs using a hull-mounted conformal array," in *OCEANS, 2001. MTS/IEEE Conference and Exhibition*, vol. 4, pp. 2270–2275 vol.4, 2001. 22

[23] T. C. Yang, "A study of spatial processing gain in underwater acoustic

communications," *Oceanic Engineering, IEEE Journal of*, vol. 32, no. 3, pp. 689–709, 2007. 22

[24] J. Tao, Y. Zheng, C. Xiao, T. C. Yang, and W.-B. Yang, "Channel equalization for single carrier MIMO underwater acoustic communications," *EURASIP Journal on Advances in Signal Processing*, 2010. 22, 50

[25] P. Moose, "A technique for orthogonal frequency division multiplexing frequency offset correction," *IEEE Transactions on Communications*, vol. 42, pp. 2908–2914, Oct. 1994. 26, 30

[26] Lower Colorado River Authority, "Historical lake levels: Highland lakes." `http://www.lcra.org/water/conditions/historical.html`, 2009. 28

[27] Measurement Computing Corporation, "Iotech personal daq/3000 series users manual." `http://www.mccdaq.com/manuals.aspx`, 2010. 28, 30

[28] K. A. Perrine, K. F. Nieman, T. L. Henderson, K. H. Lent, and T. J. Brudner, "University of texas applied research laboratory nov. 2009 five-element acoustic underwater dataset." `http://users.ece.utexas.edu/~bevans/projects/underwater/datasets/ARLUT_01_doc_01.pdf`, 2009. 30

[29] D. Lucani, M. Medard, and M. Stojanovic, "Capacity scaling laws for underwater networks," in *Proc. 42nd Asilomar Conference on Signals, Systems and Computers*, pp. 2125–2129, Oct. 2008. 30

[30] T. L. Henderson, "Matched beam theory for unambiguous broadband direction finding," *Acoustical Society of America Journal*, vol. 78, pp. 563–574, Aug. 1985. 34

[31] T. L. Henderson, "Wide-band monopulse sonar: Processor performance in the remote profiling application," *Oceanic Engineering, IEEE Journal of*, vol. 12, pp. 182–197, Jan 1987. 34

[32] T. L. Henderson and S. G. Lacker, "Seafloor profiling by a wideband sonar: simulation, frequency-response optimization, and results of a brief sea test," *Oceanic Engineering, IEEE Journal of*, vol. 14, pp. 94–107, Jan 1989. 34

[33] T. L. Henderson and T. J. Brudner, "A space-time filtered gradient method for detecting directions of echoes and transient sounds," *Acoustical Society of America Journal*, vol. 118, pp. 679–695, Aug. 2005. 34

[34] C. A. Balanis, *Antenna Theory: Analysis and Design, 3rd Edition*. Wiley-Interscience, 2005. 35

[35] J. W. S. Rayleigh, *The Theory of Sound*. Dover Publications, 1945. 41, 42

[36] M. Nassar, J. Lin, Y. Mortazavi, A. Dabak, I. H. Kim, and B. L. Evans, "Local utility powerline communications in the 3-500 khz band: Channel impairments, noise, and standards," *IEEE Signal Processing Magazine*, vol. 29, no. 5, pp. 116–127, 2012. 52, 56

[37] D. Shaver, "Low Frequency, Narrowband PLC Standards for Smart Grid The PLC Standards Gap!." `http://cms.comsoc.org/SiteGen/Uploads/Public/Docs_Globecom_2009/6_-_12-03-09_shaver_smart_grid_panel_final.pdf`, 2009. 52

[38] S. Güzelgöz, H. B. Celebi, T. Guzel, H. Arslan, and M. C. Mihcak, "Time frequency analysis of noise generated by loads in PLC," *Proc. IEEE International Conference on Telecommunications*, 2010. 53

[39] M. Katayama, T. Yamazato, and H. Okada, "A mathematical model of noise in narrowband power line communication systems," *IEEE Journal on Selected Areas Communications*, vol. 24, no. 7, 2006. 53, 57

[40] K. F. Nieman, J. Lin, M. Nassar, K. Waheed, and B. L. Evans, "Cyclic spectral analysis of power line noise in the 3—200 kHz band," in *Power Line Communications and Its Applications (ISPLC), 2013 17th IEEE International Symposium on*, pp. 315–320, 2013. 54, 67, 70

[41] M. Zimmermann and K. Dostert, "Analysis and modeling of impulsive noise in broad-band powerline communications," *Electromagnetic Compatibility, IEEE Transactions on*, vol. 44, no. 1, pp. 249 –258, 2002. 53

[42] S. Barmada, A. Musolino, and M. Tucci, "Response bounds of indoor power-line communication systems with cyclostationary loads," *Power Delivery, IEEE Transactions on*, vol. 24, pp. 596 –603, april 2009. 55

[43] M. H. L. Chan and R. W. Donaldson, "Amplitude, width, and interarrival distributions for noise impulses on intrabuilding power line communication networks," *Electromagnetic Compatibility, IEEE Transactions on*, vol. 31, no. 3, pp. 320 – 323, 1989. 56

[44] M. Nassar, A. Dabak, I. H. Kim, T. Pande, and B. L. Evans, "Cyclostationary noise modeling in narrowband powerline communication for smart grid applications," *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2012. 56

[45] A. D. Spaulding and D. Middleton, "Optimum reception in an impulsive interference environment–part I: Coherent detection," *Communications, IEEE Transactions on*, vol. 25, no. 9, pp. 910–923, 1977. 56, 58

[46] K. Blackard, T. Rappaport, and C. Bostian, "Measurements and models of radio frequency impulsive noise for indoor wireless communications," *IEEE Journal on Selected Areas Communications*, vol. 11, no. 7, p. 991?1001, 1993. 56, 58

[47] M. Sanchez, L. de Haro, M. Ramon, A. Mansilla, C. Ortega, and D. Oliver, "Impulsive noise measurements and characterization in a UHF digital TV channel," *Electromagnetic Compatability, IEEE Tranactions on*, vol. 41, no. 2, pp. 124–136, 1999. 56, 58

[48] W. A. Gardner and L. Franks, "Characterization of cyclostationary random signal processes," *IEEE Transactions on Information Theory*, vol. 21, no. 1, pp. 4–14, 1975. 57

[49] K. Gulati, B. L. Evans, and S. Srikanteswara, "Joint temporal statistics of interference in decentralized wireless networks," *IEEE Transactions on Signal Processing*, vol. 60, no. 12, pp. 6713–6718, 2012. 58

[50] M. Nassar, P. Schniter, and B. L. Evans, "A factor graph approach to joint OFDM channel estimation and decoding in impulsive noise environments," *Signal Processing, IEEE Transactions on*, vol. 62, pp. 1576–1589, March 2014. 60, 61, 62, 80, 81, 91

[51] D.-F. Tseng, Y. S. Han, W. H. Mow, L.-C. Chang, and A. J. H. Vinck, "Robust clipping for OFDM transmissions over memoryless impulsive noise channels," *IEEE Communications Letters*, vol. 16, no. 7, 2012. 61

[52] J. Hring and A. J. Han Vinck, "OFDM transmission corrupted by impulsive noise," *Proc. Int. Symp. Powerline Communications (ISPLC)*, 2000. 61

[53] G. Caire, T. Y. Al-Naffouri, and A. K. Narayanan, "Impulse noise cancellation in OFDM: an application of compressed sensing," in *Proc. IEEE International Symposium on Info. Theory*, 2008. 61

[54] J. Lin, M. Nassar, and B. L. Evans, "Impulsive noise mitigation in powerline communications using sparse bayesian learning," *Selected Areas in Communications, IEEE Journal on*, vol. 31, pp. 1172–1183, July 2013. 61

[55] K. F. Nieman, M. Nassar, J. Lin, and B. L. Evans, "Fpga implementation of a message-passing ofdm receivers for impulsive noise channels," in *Proc. Asilomar Conf. on Signals, Systems, and Computers*, 2013. 61

[56] A. Mehboob, L. Zhang, J. Khangosstar, and K. Suwunnapuk, "Joint channel and impulsive noise estimation for ofdm based power line communication systems using compressed sensing," *Proc. IEEE Int. Symposium on Power Line Communications and Its Applications*, 2013. 61

[57] D. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," *Proc. Natl. Acad. Sci.*, vol. 106, pp. 18914–18919, 2009. 62, 80

[58] L. Cavigelli and P. A. Hager, "Implementing the approximate message passing (AMP) algorithm on a GPU," *ETH Report*, 2012. 62, 80

[59] P. Maechler, C. Studer, D. Bellasi, A. Maleki, A. Burg, N. Felber, H. Kaeslin, and R. Baraniuk, "VLSI design of approximate message passing for signal restoration and compressive sensing," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 2, no. 3, pp. 579–590, 2012. 62, 80

[60] J. Lin and B. L. Evans, "Cyclostationary noise mitigation in narrowband powerline communications," in *Signal Information Processing Association Annual Summit and Conference (APSIPA ASC), 2012 Asia-Pacific*, pp. 1–4, Dec 2012. 62

[61] J. Antoni, "Cyclic spectral analysis in practice," *Mechanical Systems and Signal Processing*, no. 21, pp. 597–630, 2007. 65, 66

[62] W. Gardner and L. Franks, "Characterization of cyclostationary random signal processes," *Information Theory, IEEE Transactions on*, vol. 21, no. 1, pp. 4 – 14, 1975. 65

[63] N. Sawada, T. Yamazato, and M. Katayama, "Bit and power allocation for power-line communications under non-white and cyclostationary noise environments," *Proc. IEEE International Symposium on Power Line Communication and Its Applications*, 2009. 69

[64] K. Waheed and K. F. Nieman, "Adaptive cyclic channel coding for orthogonal frequency division multiplexed (OFDM) systems," 2014. US patent pending. 69

[65] J. G. Proakis, *Digital Communications*, pp. 638–642, 660–668. McGraw-Hill, 2 ed., 1989. 71

[66] M. Sandell, J. van de Beck, and P. Borjesson, "Time and frequency synchronization in OFDM systems using the cyclic prefix," *IEEE International Symposium on Synchronization*, pp. 16–19, 1995. 87

[67] G. J. Foschini, "Layered space-time architecture for wireless communication in a fading environment when using multi-element antennas," *Bell Laboratories Technical Journal*, pp. 41–59, 1996. 95

[68] J. S. Belrose, "A radioscientist's reaction to marconi's first transatlantic wireless experiment-revisited," in *Antennas and Propagation Society International Symposium, 2001. IEEE*, vol. 1, pp. 22–25 vol.1, July 2001. 95

[69] IEEE 802.11 Working Group, "Wireless local area networks." `http://www.ieee802.org/11/`, 2013. 95

[70] 3^{rd} Generation Partnership Project, "3GPP Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Access Network (E-UTRA); Base Station Radio Transmission and Reception (Release 8)," 2009. 95, 96

[71] A. J. Paulraj, D. A. Gore, R. U. Nabar, and H. Bolcskei, "An overview of MIMO communications - a key to gigabit wireless," *Proceedings of the IEEE*, vol. 92, no. 2, pp. 198–218, 2004. 95

[72] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *Wireless Communications, IEEE Transactions on*, vol. 9, pp. 3590–3600, November 2010. 96, 99

[73] F. Rusek, D. Persson, B. K. Lau, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, "Scaling up MIMO: Opportunities and challenges with very large arrays," *Signal Processing Magazine, IEEE*, vol. 30, no. 1, pp. 40–60, 2013. 96, 99, 102, 103

[74] E. G. Larsson, F. Tufvesson, O. Edfors, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *CoRR*, vol. abs/1304.6690, 2013. 96, 99

[75] K. T. Truong and R. W. Heath, "Effects of channel aging in massive MIMO systems," *Journal of Communications and Networks*, Aug 2013. 96

[76] H. Q. Ngo, E. Larsson, and T. Marzetta, "Massive mu-mimo downlink tdd systems with linear precoding and downlink pilots," in *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*, pp. 293–298, Oct 2013. 99

[77] C. Shepard, H. Yu, N. Anand, E. Li, T. L. Marzetta, R. Yang, and L. Zhong, "Argos: Practical many-antenna base stations," *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2013. 102, 103

[78] Samsung, "Samsung takes first 5G steps with advanced antenna." `http://www.pcworld.idg.com.au/article/461656/samsung_takes_first_5g_steps_advanced_antenna/`, 2013. 103

[79] Nutaq, "Nutaq Announces 5G Massive MIMO Partnership with Université de Sherbrooke." `http://nutaq.com/en/library/news/nutaq-announces-5g-massive-mimo-partnership-universit`, 2014. 103

[80] Xilinx, "LTE Small Cell Baseband Solutions." `http://www.xilinx.com/publications/prod_mktg/lte_small_cell_baseband.pdf`, 2013. 106

[81] K. F. Nieman, I. C. Wong, N. U. Kundargi, and B. L. Evans, "High speed processing framework for high channel count MIMO," in *IEEE International Symposium on Circuits and Systems*, 2015. , to be submitted. 109

[82] H. Huang, K. F. Nieman, Y. Hu, and D. Akinwande, "Electrically small folded ellipsoidal helix antenna for medical implant applications," in *Antennas and Propagation (APSURSI), 2011 IEEE International Symposium on*, pp. 769–771, July 2011. 122

[83] H. Huang, K. F. Nieman, P.-Y. Chen, M. Ferrari, Y. Hu, and D. Akinwande, "Properties and applications of electrically small folded ellipsoidal helix antenna," *Antennas and Wireless Propagation Letters, IEEE*, vol. 11, pp. 678–681, 2012. 122

[84] K. F. Nieman and B. L. Evans, "Time-domain compression of complex-baseband lte signals for cloud radio access networks," in *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*, pp. 1198–1201, Dec 2013. 133

[85] D. Middleton, "Non-Gaussian noise models in signal processing for telecommunications: new methods an results for class a and class b noise mod-

els," *IEEE Journal of Information Theory*, vol. 45, no. 4, pp. 1129–1149, 1999.

[86] J. Zhang, R. Chen, J. G. Andrews, A. Ghosh, and R. W. Heath, "Networked MIMO with clustered linear precoding," *IEEE Transactions on Wireless Communications*, vol. 8, no. 4, pp. 1910–1921, 2009.

[87] K. F. Nieman, N. U. Kundargi, I. C. Wong, and B. C. Prumo, "Synchronization of large antenna count systems," 2014. US patent pending.

[88] I. C. Wong, K. F. Nieman, and N. U. Kundargi, "Signaling and frame structure for Massive MIMO cellular telecommunication systems," 2014. US patent pending.

[89] N. U. Kundargi, I. C. Wong, and K. F. Nieman, "Distributed low latency Massive MIMO telecommunication transceiver processing framework and use," 2014. US patent pending.

[90] D. De Koning and W. Verhelst, "On psychoacoustic noise shaping for audio requantization," in *Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03). 2003 IEEE International Conference on*, vol. 5, pp. V–453–6 vol.5, 2003.

[91] M. Stojanovic, J. Catipovic, and J. Proakis, "Phase-coherent digital communications for underwater acoustic channels," *IEEE Journal of Oceanic Engineering*, vol. 19, pp. 100–111, Jan 1994.

[92] R. Urick, *Principles of Underwater Sound*. Peninsula, 3 ed., 1996.

[93] B. Li, S. Zhou, M. Stojanovic, L. Freitag, and P. Willett, "Non-uniform Doppler compensation for zero-padded OFDM over fast-varying underwater acoustic channels," in *Proc. IEEE Oceans Europe*, pp. 1–6, June 2007.

[94] M. Johnson, L. Freitag, and M. Stojanovic, "Improved Doppler tracking and correction for underwater acoustic communications," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1, pp. 575–578 vol.1, Apr 1997.

[95] E. C. Farnett and G. H. Stevens, *Radar Handbook*, pp. 10.1–10.39. McGraw-Hill, 1990.

[96] M. R. Ducoff and B. W. Tietjen, *Radar Handbook*, ch. Pulse Compression Radar, pp. 8.3–8.11. McGraw-Hill, 3 ed., 2008.

[97] Lower Colorado River Authority, "Historical lake levels: Highland lakes." http://www.lcra.org/water/conditions/historical.html, 2007.

[98] T. Kang and R. Iltis, "Fast-varying doppler compensation for underwater acoustic OFDM systems," in *Proc. IEEE Asilomar Conf. on Signals, Systems and Computers*, pp. 933–937, Oct. 2008.

[99] M. Chitre, S. Shahabudeen, L. Freitag, and M. Stojanovic, "Recent advances in underwater acoustic communications and networking," in *OCEANS 2008*, pp. 1–10, Sept. 2008.

[100] L. Freitag, M. Johnson, M. Stojanovic, D. Nagle, and J. Catipovic, "Survey and analysis of underwater acoustic channels for coherent communication in the medium-frequency band," in *OCEANS 2000 MTS/IEEE Conference and Exhibition*, vol. 1, pp. 131–138, 2000.

[101] M. Stojanovic, "Recent advances in high-speed underwater acoustic communications," *IEEE Journal of Oceanic Engineering*, vol. 21, pp. 125–136, Apr. 1996.

[102] T. Oberg, B. Nilsson, N. Olofsson, M. Nordenvaad, and E. Sangfelt, "Underwater communication link with iterative equalization," in *OCEANS*, pp. 1–6, Sept. 2006.

[103] V. Torres, A. Perez-Pscual, T. Sansaloni, and J. Valls, "Design of high performance timing recovery loops for communication applications," in *IEEE Workshop on Signal Processing Systems Design and Implementation*, pp. 1–4, Oct. 2006.

[104] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Springer, 1992.

Vita

Karl Frazier Nieman received the Bachelor of Science in Electrical Engineering (*Summa Cum Laude*) from the New Mexico Institute of Mining and Technology in 2009. He was appointed at Applied Research Laboratories upon joining The University of Texas at Austin and later became part of the Wireless Networking and Communication Group (WNCG) in Fall 2010. He obtained his Masters of Science in 2011 and has since been pursuing his PhD in Electrical Engineering at the same university. He has held several internships, most recently at Freescale Semiconductor and National Instruments where he has developed optimized embedded signal processing algorithms for use in wireless and power line communication systems. His research interests include multi-dimensional signal processing for high-speed, multi-antenna communication systems and physical modeling of array antenna systems. He won the best paper award at the 2013 International Symposium on Powerline Communications and the best student paper in the Architecture and Implementation Track at the 2013 Asilomar Conference on Signals, Systems, and Computers.

Email: karl.nieman@utexas.edu

This dissertation was typeset with LaTeX[†] by the author.

[†]LaTeX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's TeX Program.