

To the Graduate Council:

I am submitting herewith a dissertation written by Brandon Cole Merkl entitled “The Future of the Operating Room: Surgical Preplanning and Navigation using High Accuracy Ultra-Wideband Positioning and Advanced Bone Measurement.” I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Biomedical Engineering.

Dr. Mohamed Mahfouz, Major Professor

We have read this dissertation
and recommend its acceptance:

Dr. Aly E. Fathy

Dr. Richard D. Komistek

Dr. Lynne E. Parker

Dr. William R. Hamel

Acceptance for the Council:

Carolyn R. Hodges
Vice Provost and Dean of the Graduate School

THE FUTURE OF THE OPERATING ROOM:
SURGICAL PREPLANNING AND NAVIGATION USING
HIGH ACCURACY ULTRA-WIDEBAND POSITIONING
AND ADVANCED BONE MEASUREMENT

A Dissertation
Presented for the Doctor of Philosophy Degree
The University of Tennessee, Knoxville

Brandon Cole Merkl
December 2008

Copyright © 2008 by Brandon Merkl
All rights reserved.

DEDICATION

This dissertation is dedicated to my loving wife of three years Megan June Merkl,
without whom all of my aspirations would be merely a subtext.

ACKNOWLEDGEMENTS

I would like to acknowledge my advisor and academic mentor Dr. Mohamed Mahfouz for his tenable and diverse engineering and software development experience, which has had a fundamental contribution to my understanding of the philosophy of science. By providing an environment of innovation in research, the collective support of Dr. Mahfouz and Dr. Richard Komistek via the Center for Musculoskeletal Research and its previous incarnation, Rocky Mountain Musculoskeletal Research Laboratory, have provided me many opportunities of international travel and disparate experiences in clinical, anthropological, and engineering contexts. I would like to thank Dr. Aly Fathy, for his enduring level of difficult questions, which despite their discomfiting effect, have contributed greatly to the success of this research endeavor. Additionally, I commend my fellow doctoral students Cemin Zhang and Michael Kuhn, for their perseverance, diligent efforts, embrace of new ideas, and commitment to scientific upstanding. I thank them for their efforts on my behalf and their friendship. The author would like to thank Depeng Yang for his contribution to the DDS architecture and Adel Elsherbini for his contribution in analyzing the phase center of a single element Vivaldi antenna. To those previously mentioned as well as Anes Aref, Depeng Yang, Gary To, Emam Fatah, and David Holmes, and the many others who remain unnamed, gracious thanks is extended for all the times your personal goals were set aside for a good engineering debate, historical discussion, document assistance, or wise advice, which—despite my best efforts—has kept me a more well-rounded, better person.

ABSTRACT

This dissertation embodies the diversity and creativity of my research, of which much has been peer-reviewed, published in archival quality journals, and presented nationally and internationally. Portions of the work described herein have been published in the fields of image processing, forensic anthropology, physical anthropology, biomedical engineering, clinical orthopedics, and microwave engineering.

The problem studied is primarily that of developing the tools and technologies for a next-generation surgical navigation system. The discussion focuses on the underlying technologies of a novel microwave positioning subsystem and a bone analysis subsystem. The methodologies behind each of these technologies are presented in the context of the overall system with the salient results helping to elucidate the difficult facets of the problem.

The microwave positioning system is currently the highest accuracy wireless ultra-wideband positioning system that can be found in the literature. The challenges in producing a system with these capabilities are many, and the research and development in solving these problems should further the art of high accuracy pulse-based positioning.

TABLE OF CONTENTS

1. INTRODUCTION	1
1.1. Motivation.....	5
1.2. System Overview	10
2. SURGICAL NAVIGATION SYSTEMS: EVALUATING ELECTROMAGNETIC VERSUS OPTICAL TECHNOLOGY IN THE OR	24
2.1. Introduction.....	24
2.2. Materials and Methods.....	26
2.3. Results.....	33
2.4. Discussion	39
3. INVESTIGATION OF HIGH ACCURACY INDOOR 3-D POSITIONING USING UWB TECHNOLOGY	41
3.1. Introduction.....	42
3.2. System Architecture.....	45
3.3. Synchronization	51
3.4. Receiver-Side Peak Detection.....	57
3.5. Antenna Phase Center Variation.....	67
3.6. Localization Experiments	70
3.7. Conclusion	78
4. CHANNEL MODELING FOR UWB POSITIONING.....	80
4.1. Background	80
4.2. Statistical Channel Modeling.....	81
4.3. Ray Tracing Channel Model	83
5. PULSE AND LEADING EDGE DETECTION.....	94

5.1.	Background	94
5.2.	Sampling Requirements Simulation	94
5.3.	Iterative Peak Subtraction	98
5.4.	Max-Ratio Leading Edge Detection	101
5.5.	Genetic Programming Edge Detection	118
6.	ROBUST 3D POSITIONING IN MULTIPATH ENVIRONMENTS.....	135
6.1.	Background	135
6.2.	Kalman Filtering	146
6.3.	Protected Kalman Filtering	149
6.4.	Non-Linear Filtering	151
7.	SYSTEM CALIBRATION AND VERIFICATION	156
7.1.	System Cable Length Calibration	156
7.2.	System Time-Scale and Offset Calibration.....	159
7.3.	Base Station Calibration Algorithm using Arbitrary Path Motion	160
8.	SYSTEM CHARACTERISTICS AND ERROR ANALYSIS	173
8.1.	System testing	174
8.2.	Error Budget.....	184
8.3.	System Time Budget.....	188
9.	DISCUSSION AND FUTURE WORK	190
9.1.	System Overview	191
9.2.	Future System Improvements	192
9.3.	Future Work	192
	REFERENCES	195

APPENDICES	215
APPENDIX A – FISHER DISCRIMINANT PCA ANALYSIS ALGORITHM	216
APPENDIX B – GENETIC PROGRAMMING VERILOG MODULE GENERATION	219
APPENDIX C – C++ CODE FOR GENETIC PROGRAM SIMPLIFICATION.....	222
VITA.....	224

LIST OF FIGURES

Figure 1: a) Medtronic AXIEM™ Electromagnetic tracking [24] b) Stryker eNact™ Knee Optical Navigation System [25].....	8
Figure 2: a) Hybrid Systems: Combine tracking systems with imaging modalities, for tracked visualization and information fusion b) Positioning system used for post-operative gait lab analysis	10
Figure 3: System overview showing the interaction of the surgical preplanning subsystem and the microwave positioning subsystem with entities such as surgical navigation, gait analysis, and implant design software.....	11
Figure 4: Component diagram of the surgical preplanning subsystem	12
Figure 5: Example of a discretized surface mesh of the distal femur.....	13
Figure 6: Images highlighting the difference between male and female femora with the first principal component left out according to Algorithm A-1, red represents the largest difference, blue relatively little difference. It is clear that the largest differences between males and female occur on the medial and lateral epicondyles	14
Figure 7: Automatic landmarks and axes calculated on the distal femur	15
Figure 8: Implant placement technique showing alignment to Posterior Condylar Axis +3 degrees.	16
Figure 9: Final 3D segmentation of L5 lumbar vertebrae from a CT volume. The orange model is the result of the segmentation algorithm, while the translucent model represents the manual segmentation of an expert.....	16
Figure 10: Microwave positioning subsystem overview showing major components: control station, base station, probe, and tag	18
Figure 11: a) Power spectral density of transmitted UWB pulse after up-converted to 8 GHz carrier b) Time domain representation of normalized pulse shape, noting 300 ps pulse width (full-width-half-max)	19
Figure 12: Received signals after down-conversion, sequential analog sampling, and 8-bit AD conversion (red I, blue Q) at a rate of 50 MS/s. The the width of each pulse shown is roughly 3 μ s	20

Figure 13: Example of UWB positioning. Transmitted pulse originates from omni-directional tag and received by multiple base stations. After down conversion (not shown) analog sequential sampling reduces the bandwidth of the received signal, which is then digitally sampled, peak detection locates the sample index of the estimated peak. Computing peak, computing the difference yields time-differences which are passed to the TDOA algorithm for 3D position calculation	21
Figure 14: a) Synthetic Femur with machined landmark recesses b) Synthetic femur mounted to EM transmitter c) Femur mounted for static optical experiment	27
Figure 15: Dynamic Tracking Experiment with optical sensor mounted directly to rotating platform	28
Figure 16: Tracing acetabulum with optical probe	29
Figure 17: Sagittal view of segmented pelvis overlaid by the APP and acetabulum point cloud	30
Figure 18: Segmented pelvis overlaid by APP, acetabulum point cloud, and cup plane with APP reference frame	31
Figure 19: Patient comparison with APPs aligned. (a) Frontal view (b) Sagittal view	31
Figure 20: GPS system analogy: (a) GPS system, (b) UWB indoor positioning system.	46
Figure 21: Block diagram of indoor localization system showing one tag and three base stations which feed into the main system controller	47
Figure 22: Gaussian pulse which serves as system UWB source: (a) time domain exhibiting 300 ps pulse width, (b) frequency domain highlighting bandwidth in excess of 3 GHz.	47
Figure 23: Power spectral density of modulated pulse signal showing double sideband modulated signal with bandwidth of 6 GHz and carrier leakage at 8 GHz of -16 dBm.	48
Figure 24: Block diagram of current tag layout showing OOK digital communication and UWB transmitting architecture	49
Figure 25: Schematic of the sampling mixer highlighting the broadband balun and balanced topology [80].	49
Figure 26: ADS simulation showing the effect of Tx-Rx LO frequency offset on reconstructed sub-sampled pulse: a) original signal, b) reconstructed signal with LO offsets of 19.9 MHz, 53 MHz, 100 MHz and 118 MHz	53

Figure 27: ADS simulation of the reconstructed sub-sampled pulse with and without 3 ps of PRF clock jitter.	
.....	56
Figure 28: Experimental setup with the unsynchronized PRF clock sources to measure the effect of PRF clock jitter.....	56
Figure 29: Time variation of the pulse peak position acquired over n sample points.	56
Figure 30: Normalized received signal without multipath components (i.e. clean template).....	59
Figure 31: Actual received signal with severe multipath.	60
Figure 32: Original received multipath signal overlaid with selection function and detected peaks.	63
Figure 33: Received multipath signal $h_0(t)$ and estimated main peak signal $r(t)$	64
Figure 34: Omni-directional transmitting antenna placed in a severe multipath environment surrounded by numerous closely spaced metal objects.	65
Figure 35: Simulation of broadside Vivaldi phase center location versus frequency.	68
Figure 36: Experimental setup of Vivaldi antennas in an anechoic chamber used to measure Vivaldi antenna directivity-dependent phase center variation.	68
Figure 37: Measured Vivaldi phase center error versus angle for E-cut and H-cut.....	69
Figure 38: 3-D synchronized localization experiments: (a) 6 base stations. (b) 5 base stations. (c) 4 base stations with low position dilution of precision (PDOP). (d) 4 base stations with poor coverage in x-direction resulting in high PDOP.....	71
Figure 39: Error in x, y, and z illustrating high PDOP for base station distribution shown in Figure 38 (d). This shows the reduction in accuracy that results from poor base station spatial arrangement.	72
Figure 40: Experimental setup for 1-D unsynchronized positioning measurement.....	75
Figure 41: Measured error of 1-D unsynchronized experiment.	75
Figure 42: Power delay profile of multi cluster, S-V UWB channel model	82
Figure 43: Ray Tracing diagram showing three rays: one line-of-sight ray, and two non-line-of-sight rays	84
Figure 44: a) Simulated ray tracing results for transmitter and receiver 1 cm away from metal wall, the signal without the presence of the metal wall (green) is severely attenuated by destructive interference in the	

received signal with the wall present (red) b) Experimental results for the same configuration, noting the similarities of destructive interference in the received signal	85
Figure 45: a) Simulated ray tracing results for transmitter 1 cm and receiver 104.5 cm away from metal wall, respectively, the signal without the presence of the metal wall (green) is slightly shifted by destructive interference in the received signal with the wall present (red) b) Experimental results for the same configuration, noting the similarities of the strong dominant ray	86
Figure 46: a) Simulated ray tracing results for transmitter 1.00 m and receiver 1.0 cm away from metal wall, respectively, the signal without the presence of the metal wall (green) is slightly shifted by destructive interference in the received signal with the wall present (red) b) Experimental results for the same configuration, noting the similarities of the strong dominant ray	86
Figure 47: a) Simulated ray tracing results for transmitter 1.02 m and receiver 1.00 mm away from metal wall, respectively, the signal without the presence of the metal wall (green) is no longer shifted by destructive interference in the received signal b) Experimental results for the same configuration, noting the similarities of the strong dominant ray, and absence of significant multipath signals	86
Figure 48: Near Wall/Near Wall situation a) Frequency dependence of the channel, noting 10 dB range b) Simulated received pulse (red) showing only slight distortion referencing the transmitted pulse (blue)	87
Figure 49: Near Wall/Far Wall and Far Wall/Near Wall situations a) Frequency dependence of the channel, noting 30 dB range b) Simulated received pulse (red) showing stronger distortion referencing the transmitted pulse (blue)	88
Figure 50: Far Wall/Far Wall situation a) Frequency dependence of the channel, noting nearly 15 dB range b) Simulated received pulse (red) showing weak distortion referencing the transmitted pulse (blue) ...	88
Figure 51: InSite software simulation of ray tracing experiment. Transmitter positions (red squares) are varied in distance from the metal wall (yellow, bottom), while receiver positions are moved in two parallel tracks (green squares) relative to the transmitter.	90
Figure 52: Insite software simulation corresponding to transmitter and receiver each 1 cm away from the metal wall, corresponding to the experimental and simulation results of Figure 44.....	90

Figure 53: Insite software simulation corresponding to transmitter 1.0 cm and receiver 1.0 m away from the metal wall, corresponding to the experimental and simulation results of Figure 45. Inset is a magnified view of the dominant peak.....	91
Figure 54: Insite software simulation corresponding to transmitter 1.0 m and receiver 1.0 cm away from the metal wall, corresponding to the experimental and simulation results of Figure 46. Inset is a magnified view of the dominant peak.....	91
Figure 55: Insite software simulation corresponding to transmitter 1.0 m and receiver 1.0 m away from the metal wall, corresponding to the experimental and simulation results of Figure 47.....	92
Figure 56: Comparison of high directivity receiving antenna (blue, Vivaldi) with omni-directional receiving antenna (red, monopole), when an omni-directional monopole is used as transmitter in both cases .	93
Figure 57: TDOA system simulation on the effect of sampling error when N=4 base stations are used in the positioning calculation	96
Figure 58: TDOA system simulation on the effect of sampling error when N=6 base stations are used in the positioning calculation	96
Figure 59: TDOA system simulation on the effect of sampling error when N=8 base stations are used in the positioning calculation	97
Figure 60: Maximum percent scale factor error vs. scale factor (α) for two clock stability values, 0.5 ppm (blue) and 1.0 ppm (red)	98
Figure 61: Simulated multipath signal with a secondary peak 110% of the main peak delayed 440 ps. ...	100
Figure 62: Comparison of error (in mm) for the three peak finding algorithms versus separation of primary and secondary peaks for secondary peak amplitude of 110% Figure 61.	100
Figure 63: Interference of two 300 ps Gaussian pulses separated by 138 ps with the amplitude of the secondary pulse 110% that of the primary pulse. The large overlap between the pulses causes a significant shift in the peak of the combined pulse.....	100
Figure 64: Interference of two 300 ps Gaussian pulses separated by 262 ps with the amplitude of the secondary pulse 110% that of the primary pulse. The increased separation between the pulses reduces the lagging shift, and consequently the error in locating the peak position	101

Figure 65: Example of Max-Ratio Leading Edge Detection. The original leading-edge signal $h_o[t]$ (blue) is averaged $y[t]$ (solid red), subsequently two max-window filters $\text{max_16}[t]$ (dashed red) and $\text{max_256}[t]$ (dashed green) operate on $y[t]$, the result $r[t]$ (light blue) makes a positive transition when the ratio of $\text{max_256}[t]:\text{max_16}[t]$ is less than 4	103
Figure 66: Max-Ratio Leading Edge Simulation: Error in positioning due to multipath peaks with various power levels (1.9 dB to -40 dB) as detected by max-ratio leading edge algorithm. Note the maximum values are recorded in Table 20	104
Figure 67: Max-Ratio Leading Edge Simulation: Error in positioning due to sub-sample shifts in peak location, noting that very little error results in this case	105
Figure 68: Max-Ratio Leading Edge Simulation: Error resulting from degradation in SNR due to AWGN, 10 random trials. Note that with the exception of one outlier, error decrease is roughly linear for SNR values between 2.5 to 15	106
Figure 69: Max-Ratio Leading Edge Simulation: Instantiated signal (blue) with high SNR (30.0) and associated low error (0.56 mm), showing calculated noise threshold (black)	108
Figure 70: Max-Ratio Leading Edge Simulation: Instantiated signal (blue) with medium SNR (14.6) and associated medium error (4.30 mm), showing calculated noise threshold (black).....	108
Figure 71: Max-Ratio Leading Edge Simulation: Instantiated signal (blue) with low SNR (3.2) and associated higher error (28.7 mm), showing calculated noise threshold (black).....	109
Figure 72: Xilinx schematic describing pin inputs and outputs when Base Station Module is configured in one sample channel (I) per base station configuration as depicted in Figure 73	110
Figure 73: One sample channel (I) per base station configuration	111
Figure 74: Two sample channels (I and Q) per base station configuration	112
Figure 75: Overall diagram of FPGA implementation of DSP module	113
Figure 76: Module-level diagram showing internal structure of FIR LPF	114
Figure 77: Module-level diagram showing internal structure of Peak Detection functionality	114
Figure 78: Module-level diagram showing internal structure of Noise Detect module	115

Figure 79: ModelSim signal simulation showing Max-Ratio Algorithm as implemented in Verilog HDL targeted for FPGA	117
Figure 80: Average 1D error in mm versus sample index for prior fixed-threshold algorithm (red) and Max-Ratio algorithm (blue), noting that both the maximum error and the average error show significant improvement.....	119
Figure 81: Example of a signal processing algorithm in tree representation for use in genetic programming	120
Figure 82: Genetic Programming Algorithm Choices.....	121
Figure 83: Diagrammatic view of shifts applied to noisy signals, to avoid deterministic genetic programs	124
Figure 84: Example of how to apply the boosting meta-algorithm to improve genetic programming results	125
Figure 85: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 6, 30% mutation rate, 55% crossover, MDL scoring, 100 individuals.	127
Figure 86: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 8, 30% mutation rate, 55% crossover, MDL scoring, 100 individuals	128
Figure 87: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 10, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals.	129
Figure 88: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 8, 30% mutation rate, 25% crossover, MDL scoring, 100 individuals	129
Figure 89: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 10, 30% mutation rate, 25% crossover, MDL scoring, 100 individuals.	130

Figure 90: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 8, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals.	131
Figure 91: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 10, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals.	131
Figure 92: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 12, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals. This run utilized the full set of functional nodes	132
Figure 93: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 12, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals. This run utilized a sub-set of the possible functional nodes	133
Figure 94: Average 1D error in mm versus sample index for GP algorithm (red) and Max-Ratio algorithm (blue), noting that both the maximum error and the average error are worse for the GP algorithm than the Max-Ratio algorithm or the previous fixed threshold method (Figure 80)	134
Figure 95: Calculation of tag position based on TDOA approach	138
Figure 96: For the TDOA algorithm, the degradation in PDOP due to geometric differences between the Flat 9 configuration and the Curved 9 configuration	147
Figure 97: Raw static positioning error using TDOA algorithm and a cable length calibration	148
Figure 98: A Kalman filtered result of the exact data contained in Figure 97, note that after the filter stabilizes around sample 700 error approaches 10-12 mm.....	148
Figure 99: Dynamic simulation comparing Kalman Filtering 3D position data (red) with Kalman Filtering with Mahalanobis Protection (blue). Note that the largest spikes have been eliminated.....	150
Figure 100: Dynamic simulation comparing Kalman Filtering 3D position data (red) with Kalman Filtering with D-Matrix Protection (blue), note that the RMS error has been significantly reduced	151
Figure 101: Synthesized I^2 channel data corrupted by transmitter-receiver 8 GHz LO frequency offset	152

Figure 102: Effect of several types of windowed TOA index filters on the raw data (blue), also represented in Figure 103, each of the windows are size 16, the MIQR filter (yellow) has the lowest variance of the three types.....	153
Figure 103: Cumulative distribution of pulse TOA index values showing non-Gaussian structure	154
Figure 104: Diagram showing cable length offsets and consequent need for system calibration	157
Figure 105: Experimental results showing the error trend of sampled data positioned using the cable length calibration algorithm combined with the TDOA algorithm with no filtering	158
Figure 106: Experimental results showing the error trend of sampled data positioned using the TDOA algorithm with a 1000-point averaging filter, and $\alpha = 100,000$ scale factor, noting the possibility for $\pm 11\%$ scale factor error	159
Figure 107: Measured Vivaldi phase center error versus angle for E-cut, H-cut, average of E and H cuts, and centered quadratic fit.....	161
Figure 108: Designed Vivaldi antenna, with the +z axis being the boresight direction	164
Figure 109: Instance of the 100 base station experiment	165
Figure 110: 3D Positioning Error vs. Tag x Position for 100 base station experiment illustrating the difference between Φ included (dashed) and not included (solid) in equation (4)	166
Figure 111: Base station specific results of errors shown in Figure 110, note that base stations oriented along the tag track experience the least amount of error due to phase center variation	166
Figure 112: Visual depiction of Step 3b of Algorithm 7	169
Figure 113: Comparison of N=6 base stations and N=10 base stations on the calibration algorithm convergence of RMS error. Note algorithm instability for N=6 case. Sub-mm calibrated accuracy is possible with N=10 base stations	170
Figure 114: Comparison of N=6 base stations and N=10 base stations on the calibration algorithm convergence of maximum error. Note algorithm instability for N=6 case. Sub-mm calibrated accuracy is possible with N=10 base stations	171
Figure 115: Comparison of different scale factors: 10000 (red) and 40000 (blue) on 3D RMS error, in a non-coherent, MIQR=17, 5x point averaging, using TDOA algorithm	175

Figure 116: Comparison of wireless and wired PRF (10MHz) clocks: wireless (red) and wired (blue) on 3D RMS error, static, MIQR=41, 5x point averaging, using TDOA algorithm	175
Figure 117: Raw data captured using wired PRF clocks and varying levels of LO (8GHz) coherence on each of the 4 base stations	176
Figure 118: 3D RMS error values resulting from application of various 3D Kalman filtering techniques with non-coherent base stations LOs, compare with Figure 117	177
Figure 119: 3D RMS error values resulting from index filtered data captured using wireless PRF clocks, coherent base stations, and the TDOA algorithm	178
Figure 120: 3D RMS error values resulting data captured dynamically in the presence of metal interferers using wireless PRF clocks, coherent base stations, MIQR=17, ave=5, and the TDOA algorithm	178
Figure 121: Free-form tag motion, comparing the 3D tracked points of the Optotrak (red) and UWB systems representing 6.95 mm RMS error, IQR=17, no averaging, wireless, non-coherent	179
Figure 122: Three orthographic views a) Top b) Front c) Side view of dynamic track experiment. All dimensions are in mm, where UWB data (blue) is plotted alongside Optotrak data (red). Optotrak data in the top view exhibits translation in cross-track direction due to a loose sliding bracket	181
Figure 123: Hysteresis testing showing positive (blue) and negative (red) motion down a track on two occasions. Note local nonlinearities which appear to be correlated in each direction. Hysteresis error is 4.7% and 3.3% of full-scale, respectively	181
Figure 124: 3D large range linearity testing with dynamic robot motion. Optotrak data (red) is notably smoother, however, the UWB (blue) system exhibits only a slight linearity error (~1%)	182
Figure 125: Robot repeatability experiment, where each of 20 points was visited twice. UWB data (blue) and Optotrak points are plotted for direct comparison; RMS 3D error is 7.00 mm	183
Figure 126: Best results found: 3.26 mm for static positioning and 5.14 mm for dynamic positioning	184

LIST OF TABLES

Table 1: Performance comparison of EM vs. optical tracking systems used in surgical navigation systems	9
Table 2: Comparison of commercial and research UWB positioning systems	18
Table 3: Mean native cup orientation and mean cup orientation corrected for flexion for combined optical and EM tracking experiments (six total trials per patient) in degrees	32
Table 4: Optical static landmark results in mm	34
Table 5: EM static landmark results (no metal) in mm	34
Table 6: EM static landmark results (with metal table) in mm.....	34
Table 7: Optotrak static distances vs. EM static differences in mm	Error! Bookmark not defined.
Table 8: Optotrak static distances vs. EM static distances (with metal) in mm	36
Table 9: EM static distances vs. EM static distances (with metal) in mm	36
Table 10: Optical dynamic tracking results	37
Table 11: EM dynamic tracking results.....	37
Table 12: Mean and standard deviation of corrected cup orientation for optical and EM tracking systems averaged over both patients (in degrees).....	38
Table 13: Static errors of positioning systems and the contribution to surgical axis error	39
Table 14: Comparison of the Designed Sampler with Commercial Sampling Modules.....	50
Table 15: 1-D Multipath Experiment.....	65
Table 16: PDOP Summary – Synchronized Localization Experiments.....	74
Table 17: Error Summary – Synchronized Localization Experiments.....	74
Table 18: Error Summary – Unsynchronized 1-D Experiment.....	75
Table 19: Fresnel Ellipsoid Maximum Waist Diameters for several typical Tx-Rx distances, assuming a 300 ps pulse.....	93
Table 20: Comparison of maximum error for varying secondary peak amplitudes	102
Table 21: Comparison of Max-Ratio algorithm simulation result and ModelSim result, demonstrating a consistent leading edge result	118

Table 22: Partial List of functional nodes used with genetic programming	122
Table 23: Parameters studied controlling one GP Run	123
Table 24: Synthetic Data Genetic Programming Results	126
Table 25: Comparison of various positioning algorithms, the input vector represents the measured values passed to the algorithm, while the solution vector represents the available solution upon algorithm completion	137
Table 26: Comparison of simulation error results across differing positioning algorithms when 11.0% scaling error is present using various base station configurations. *indicates the algorithm failed to converge	145
Table 27: Comparison of simulation error results across differing positioning algorithms when 1.1% scaling error is present using various base station configurations.	146
Table 28: 3D Positioning RMS errors in mm for various index filtering techniques used in conjunction with the Max-Ratio Algorithm	154
Table 29: 3D Positioning 95% C.I. in mm for various index filtering techniques used in conjunction with the Max-Ratio Algorithm.....	Error! Bookmark not defined.
Table 30: Breakdown of system errors with observed and theoretical magnitudes	185
Table 31: Maximum slew rate compared to system -10 dB band width	188
Table 32: System characteristics which determine tag positioning throughput and latency	189
Table 33: Categorized system improvements, with expectations on improvement metrics	193

NOMENCLATURE

dB	decibels
dBm	decibels referenced to 1 milliwatt
°C	degrees Celsius
cm	centimeter
fs	femtosecond
GHz	gigahertz
GS/s	giga-samples per second
Hz	hertz
Kbits/s	kilobits per second
MHz	megahertz
MS/s	mega-samples per second
μs	microsecond
mm	millimeter
ns	nanosecond
ppm	parts per million
ps	picosecond

ABBREVIATIONS

3D	Three Dimensional
ADC	Analog to Digital Converter
AWGN	Additive White Gaussian Noise
AOA	Angle-of-Arrival
CAS	Computer Aided Surgery
CM	Channel Model
COTS	Commercial off-the-shelf
CT	Computed Tomography
DC	Direct Current
DDS	Direct Digital Synthesizer
DOF	Degrees-of-freedom
DSP	Digital Signal Processing
EM	Electromagnetic
FCC	Federal Communications Commission
FDR-PCA	Fisher Discriminant – Principal Components Analysis
FIR	Finite Impulse Response
FMCW	Frequency-modulated Continuous-wave
FPGA	Field Programmable Gate Array
GA	Genetic Algorithm
GDOP	Geometric Dilution of Precision
GP	Genetic Programming
GPS	Global Positioning System

HTLS	Hankel-Total-Least-Squares
IIR	Infinite Impulse Response
IPS	Iterative Peak Subtraction
IQ	In-phase and Quadrature signals
IR	Infrared
IR	Impulse Radio
LNA	Low noise amplifier
LO	Local Oscillator (clock)
LOS	Line-of-Sight
LPF	Low pass filter
MCW	Mutual Correspondence Warping
MDL	Minimum Description Length
ML	Medio-lateral
MNV	Mean Negative Variance
MRI	Magnetic Resonance Imaging
NLOS	Non-Line-of-Sight
OOK	On-Off Keying
OR	Operating Room
PCA	Principal Components Analysis
PDOP	Position Dilution of Precision
PRF	Pulse Repetition Frequency (clock)
RF	Radio-frequency
RMSE	Root-Mean-Square Error
RSS	Received Signal Strength
Rx	Receiver
SNR	Signal-to-noise ratio
SRD	Step-recovery diode
STOA	Scaled-Time-of-Arrival
S-V	Saleh-Valenzuela
TDMA	Time-division-multiple-access
TDOA	Time-Difference-of-Arrival
THA	Total Hip Arthroplasty
TKA	Total Knee Arthroplasty
TOA	Time-of-Arrival
Tx	Transmitter
UWB	Ultra-Wideband

1. INTRODUCTION

Since the inception of human existence, mankind has sought to extend and enrich the length and quality of life by performing surgical repairs to the body. Through the use of increasing technological prowess, these types of procedures have now become a societal and medical norm. From orthotics and prostheses, mankind has a rich and varied history of replacing, repairing, or augmenting missing, deficient, or worn parts of the human anatomy. Prior to many technological breakthroughs, however, replacing bones and joints wholly within the body was a difficult matter due to the high risk of infection. At the beginning of the 20th century, the modern orthopedic industry was founded by Revra DePuy and Justin Zimmer, inventor and refiner of the metal splint, respectively [1]. As the orthopedic industry has matured over the last century, complete replacement of joints has become commonplace. As of 2008, an estimated 500,000 knee replacement surgeries per year will be performed in the U.S., with that number jumping to an estimated 3.2 million in a decade [2]. Similar surgeries exist for the hip, elbow, ankle, shoulder, and spine, with each joint having specific implants fabricated typically using a combination of metal and plastic parts designed to articulate a reproduction of anatomical motion. Depending on the complexity of the joint and the degree of weight bearing activities, the implant, in conjunction with the accuracy of implantation, can determine failure rates for the device [3]. The design methodologies of implantable devices have focused on improving the shape and size matching between the device, as well as developing new surgical navigation technologies and techniques to improve patient outcomes.

Precipitated by an ageing population, the increase in frequency of these techniques has fostered a high degree of innovation in a field traditionally dominated by a few large orthopedic device manufacturers. As the science of metrology has met with new understandings of human anatomical variation, orthopedic device manufacturers have developed implants designed for segments of the population [4]. Additionally, novel software-based tools allow surgeons new ways to plan for surgeries using pre-operative imaging modalities such as Magnetic Resonance Imaging (MRI) [6], Computed Tomography (CT) [7], ultrasound [8], video fluoroscopy [9] [10] [11], and high-resolution biplanar digital X-rays [12] [13]. By facilitating new ways to visualize and measure the human anatomy non-invasively, versus antiquated methods such as calipers, rulers, goniometers, or osteometric boards [14], surgeons can be better guided to select the correct size and placement of the orthopedic device [15]. This added benefit is realized in two ways. First, hospital managers can reduce large inventories of implantable devices when the surgeon can select one or two probable sizes prior to the surgical operation. Second, by assessing pre-operatively the position and orientation of the implant, any corrective steps can be evaluated post-operatively for closed-loop feedback of surgeon performance.

Another feedback loop necessary to the continuous improvement of implantable devices, is that of the orthopedic manufacturer's necessary adaptation to the changing needs of the market, technological advances in materials, and ultimately the body of surgical knowledge refined by surgeons. Traditionally, any device design changes undertaken by the manufacturer's exposes the company to regulatory risks, and any feedback on the efficacy of the device is muddled in the results of often contradictory clinical studies, which can require more than 10 years to complete (e.g. [16]). To complicate matters further, the results on device performance can be distributed

across multiple hospital locations with procedures performed by different surgeons each using one of many techniques [17].

To ensure that the feedback loops of surgeon performance and implant design performance offer the highest degree of reliable information, and by extension, to diminish the effects of variables and parameters that cannot be controlled, surgical pre-planning systems and surgical navigation systems deliver repeatability, accuracy, and analysis. Such systems, when endowed with some level of automation, can provide optimized solutions in cases when compromise is inevitable. Surgical pre-planning systems offer the choice of implant size and placement either by surgeon manipulation or automatic methods which seek to match patient-specific anatomy originating from pre-operative medical images [5]. In the context of a clinical study, for example, if the surgical protocol mandated a particular pre-operative templating procedure designed to constrain implant placement in some fashion, then finer-grained analyses can take place on variables or metrics of interest to the research surgeon or the device design engineer.

Pre-operative planning and design software has the potential to improve individual patient outcomes and, thus, when used as part of a broader research endeavor, empower statistical analyses [18]. On the individual level, planning software can be instrumental in reconstruction, revision, or trauma cases, when deformed or missing anatomy must be accurately reconstructed for complete rehabilitation. Surgical pre-planning systems must offer the surgeon a competing view of the surgeon's own intuition, to compel usage beyond only the rudimentary. A surgical view that defies convention or intuition might not sound optimal; however, when the system can mine statistics or infer failure rates from past performances, then adaptation to new and better techniques can be accelerated.

While pre-operative software can direct the advancement of surgical technique, surgical navigation systems working in conjunction can minimize deviations from the surgical plan. The desirable features of such a system are information visualization, overall accuracy, rapid positioning feedback, and relevant warnings or surgical suggestions. In the context of medical practice, the surgeon can exercise a greater level of confidence in the procedure by ameliorating the risk of mal-alignment. Clinical research studies using surgical navigation systems can benefit from increased levels of quantitative information regarding the procedure such as the angle of implantation. Further, surgical navigation systems that report metadata originating from pressure sensors [19] for ligament balancing, ultrasound modalities for soft or hard tissue imaging and positioning, or inflammation sensors embedded in the implant can facilitate additional dimensions of analysis. From a health management perspective, surgical navigation systems can add value by minimizing risks due to malpractice claims, albeit at the cost of longer length of time spent in the operating room (OR).

The research encompassed in this dissertation provides a broad range of prospective technologies and techniques that will dramatically shape the operating room of the future. Technologies developed will offer an alternative to positioning technologies that are currently used for computer assisted surgery (CAS), specifically in the area of surgical navigation. The positioning system described herein was built using ultra-wideband (UWB) technology operating in the microwave range of radiofrequency (RF). These frequency bands were chosen to eliminate problems of coherent interferences, multipath, and line-of-sight (LOS) requirements, which limit the usability of other technologies currently being used for positioning such as infra-red optical sensors and electromagnetic (EM) sensors. Additionally, described as part of the background to

this work is a suite of bone analysis techniques reliant on a statistical representation of bones, which, when used in conjunction with the novel UWB positioning system can aid in surgical bone preparation pre-planning.

1.1. Motivation

A surgical navigation system is desired which can encompass pre-operative planning, intraoperative tracking, and post-operative kinematic or gait analysis. A tracking system that does not impede the surgeon's work flow is another major requisite. Intelligent bone tracking and analysis—fused with a preoperative suite of analysis tools—represents the next-generation of surgical navigation systems.

Existing navigation systems use technologies that have several drawbacks [20], the most serious of which is restricting the surgeon's workable area in the OR. Optical surgical navigation systems have a line-of-sight (LOS) requirement which inhibits surgeons and their assistants from occupying certain portions of the surgical field. Additionally, electromagnetic (EM) surgical navigation systems, while not constrained to LOS operation, become degraded during dynamic tracking and in the presence of ferromagnetic interferers. The characteristics of both of these technologies are studied in the context of static and dynamic accuracy in realistic surgical scenarios in Chapter 2.

A problem afflicting major surgical navigation systems is implant placement technology, including automatic implant sizing and surgical compromise analysis. Many imageless surgical navigation systems rely heavily on landmark-based implant fitting strategies which may not be optimal when implant placement constraints and size constraints may be disjoint. Also, shape analysis software used prior to a surgery by either surgeons or teams of engineers may yield the

ability for rapid prototyping of new implant designs or the capability for patient-specific implant design.

Using UWB technology for the underlying positioning technology can solve several of the technological limitations that existing surgical navigation systems currently face. UWB technology is not limited to LOS conditions for proper operation and operates accurately in the presence of metal. Due to the extremely wide band, UWB has the ability to resolve many multipath components, removing the effects of large metal reflectors within a time-evolving, cluttered OR environment.

Also, the use of advanced shape analysis techniques applied to bones or other portions of anatomy offers automatic methods for extracting bone shape from medical images and automatic methods for surgery pre-planning steps of implant alignment

1.1.1 Surgical Procedures

Surgical procedure outcomes for certain surgeries have significant variation based on the accuracy of the surgical technique. Given the participatory role of the femur in the largest ginglymoidal and enarthrodial joints of the adult body, replacement surgeries for both the knee and hip joint, respectively, have the potential to suffer from implant mal-alignment. Total Knee Arthroplasty (TKA) involves the complete removal of all articular surfaces of the distal femur and proximal tibia. Replacing the articular surfaces are typically titanium alloy implants with a polyethylene bearing. Since the bone preparation steps of this procedure determine the final implant placement, pre-planning involves selecting which portions of the bone to remove and selecting the correct size and shape of implant according to considerations of patient-specific anatomy and pathology. Surgical techniques typically seek to either restore the patient's original

anatomy or correct a deformity. In either case, a pre-planning system that correctly measures the bone geometry and ascertains optimal implant fit can lead to improvements in sizing [5]. The longevity of the surgical procedure has been related to the implantation accuracy [21].

Total Hip Arthroplasty (THA) procedures remove the entire articular surface of the acetabulum, replacing it with a hemispherical hip cup implant most often with polyethylene, metal, or ceramic bearing material. The femoral head is completely removed and a femoral stem is inserted into the intramedullary (IM) canal. Similarly, the surgical technique seeks to restore or correct patient anatomy with the overall goal to yield the highest range of motion of the joint. Success rate of the technique is related to the accuracy of implantation (no impingement), which can be improved with pre-operative measurement of acetabular angle [22]. Also, recent work has analyzed the impact of proximal femoral morphology and its potential for correlation with indications of poor bone quality such as osteoarthritis [23].

1.1.2 Surgical Navigation Systems

Surgical navigation systems provide real-time feedback and quantitative assessment to the surgeon regarding the quality of the surgical process. This typically includes preoperative steps such as implant sizing and pre-planned placement as well as intra-operative steps such as bone cutting guide placement and kinematic assessment. Figure 1a. shows an EM surgical navigation system on a hypothetical knee surgery, while Figure 1b. shows the feedback data available to the surgeon on the user interface of a commercial optical surgical navigation system.

The taxonomy of surgical navigation systems includes both image-based systems, and imageless systems [26]. Further, imageless systems can be sub-categorized based on their reliance on either land marks to navigate the joint of interest or via dynamic model generation, known as bone




Figure 1: a) Medtronic AXIEM™ Electromagnetic tracking [24] b) Stryker eNact™ Knee Optical Navigation System [25]

morphing. Another metric of classification for surgical navigation systems is that of the system's guidance to the surgeon. When the surgical navigation system merely provides visual feedback to the surgeon, it is termed a passive surgical navigation system. This is the dominant category of surgical navigation systems with Medtronic's AXIEM™ Electromagnetic system and Stryker's eNact™ Knee Optical Navigation System (Figure 1) both being categorized as passive systems. Some systems provide either robotic positioning of cutting guides, such as Zimmer's proposed BRIGIT™ system [27], or robot assistance in bone preparation, an example being MAKO's Tactile Guidance System™ [28] which employs haptic robotic feedback to guide the surgeon and prevent gross errors. These two systems would fall under a semi-active category, due to the shared role of the navigation system and the surgeon. Another example of such a system is Intuitive Surgical's da Vinci Surgical System™ [29], where the supervising physician navigates the surgery via a co-located control station. A fully active surgical navigation system, by comparison, performs robot-controlled bone preparation under the limited supervision and guidance of a surgeon. One such system is the ROBODOC™ system [30], which automatically prepares bone surfaces with an accuracy of 0.4 mm.

With respect to operating room logistics and surgery times, not all surgical navigation systems are created equal. Image-based systems, for example, can encumber the operating field with bulky equipment, such as an intra-operative C-arm for fluoroscopy or CT / MRI apparatus. Imageless systems with large robotic components can significantly reduce the space in the operating field. Necessary components such as optical tracking cameras or computer monitors can also further constrict space in the OR. Even more importantly, optical probes themselves can block the surgeon's vision of the surgery, as they are often 10-20 cm in size.

The system-level performance requirements for surgical navigation positioning technologies are demanding, as each system must operate in the cluttered environment of the OR while providing real-time positioning feedback to the surgeon. The orthopedic industry standards for accuracy are on the order of 1-3 mm of 3D positioning accuracy, with about 1.0° orientation accuracy required [31]. Most imageless systems must track multiple bones as well as one or more tools, for a total of 3-4 simultaneously tracked objects each requiring position and orientation, or 6 degrees-of-freedom (DOF). Table 1 highlights system achievable system performance of commercially available tracking technologies.

As with any technology, new and innovative uses for surgical navigation systems are continuously developed to aid in new surgical procedures, diagnoses, and basic research. Thus, novel systems capable of eliminating shortfalls of existing technologies will begin to see

Table 1: Performance comparison of EM vs. optical tracking systems used in surgical navigation systems

	EM Tracking Systems	Optical Tracking Systems
Update Rate	300 Hz (6DOF positions/sec)	60-70 Hz
Number of Tools	2-4	15
Accuracy (RMS)	0.7-1.4 mm	0.1 mm (active), 0.25 (passive)




Figure 2: a) Hybrid Systems: Combine tracking systems with imaging modalities, for tracked visualization and information fusion b) Positioning system used for post-operative gait lab analysis

widespread use in similar systems in related medical and industrial fields of use. Shown in Figure 2a is a conceptual hybrid system fusing various medical imaging modalities with a positioning system for information fusion, and also shown in Figure 2b is the potential alternative use for tracking systems in gait labs for kinematic analysis.

1.2. System Overview

The developed system is comprised of two subsystems, each designed to address shortcomings with existing surgical navigation systems. As illustrated in Figure 3, the Surgical Preplanning Subsystem interacts with the preplanning phase of the surgical navigation system, storing preplanned data for the surgery. The Microwave Positioning Subsystem interacts primarily with the surgery phase of the surgical navigation system. Auxiliary uses for the system include a post-operative gait analysis system, as well as implant design and analysis software. The Surgical Preplanning Subsystem is retained for background of the main body of work, which is the development of the Microwave Positioning Subsystem.




Figure 3: System overview showing the interaction of the surgical preplanning subsystem and the microwave positioning subsystem with entities such as surgical navigation, gait analysis, and implant design software

1.2.1 Surgical Preplanning Subsystem Overview

Prior to surgery, the surgeon uses a host of medical imaging modalities such as CT, MRI, and/or X-ray data to plan surgical route, bone portions to resect (including osteophyte removal), as well as implant sizing and placement. The pre-planning step is important to the surgical procedure, because it provides the surgical navigation system with a basic understanding of the surgical procedure to take place. The surgical navigation's feedback intra-operatively allows for the surgeon to have confidence in the original plan or to make necessary adjustments. After the surgery, post-operative analysis can verify the surgical plan's success in accurate implant placement.

Currently, the surgical plan is typically created in a manual process using the surgeon's judgment. However, in order to move beyond this basic form of planning, robust methods must

be available to automate the process. As diagrammed in Figure 4, steps included in an automatic process would be the following: image-based reconstruction of patient anatomy, data extrapolation and cleaning, global shape analysis, bone measurement, and finally an overall calculation of implant size and fit.

1.2.1.1 Model Preparation

Prior to use in preplanning, CT, MRI, or biplanar X-ray image data is segmented or reconstructed and transformed into triangulated surface model [32]. This model is then normalized and added to a statistical bone atlas [33] using a novel algorithm [34]. This step of surface regularization can be used to smooth erroneous data that may arise from the medical images or to clear away pathological portions of the bone, as in the case of osteophytes. Figure 5 depicts a portion of the femur after surface normalization.

1.2.1.2 Statistical Bone Atlas

The fundamental characteristic of this type of statistical shape representation, termed a “bone atlas,” is a homologous set of matched points across a population of bones [32] [33]. Given such a set, Principal Components Analysis (PCA) can be performed to compress the valuable (i.e. statistically significant) portions of the shape information into a series of shape descriptors [35].




Figure 4: Component diagram of the surgical preplanning subsystem




Figure 5: Example of a discretized surface mesh of the distal femur

This compact statistical shape representation drives automatic preplanning, performs reconstruction (or extrapolation), analyzes global bone shape, automates feature generation, and aids in registration to surgical navigation system. To create the bone atlas, the matched points between sets of bones are used to calculate a model covariance matrix around the mean bone. The task of matching points is difficult considering the anatomical disparities across a population of bones. Thus, a novel non-linear iterative warping approach was developed, which is termed Mutual Correspondence Warping (MCW) [34]. This algorithm builds on previous work in the field and improves on the final matching step where a number of standardized computer vision methods are used to place the bone in a rough alignment.

1.2.1.3 Global Shape Analysis

In this portion of the pre-planning steps a single bone or sets of bones can be classified [36] [37] or manually visually inspected for shape variation or deviation from typical population statistics. For example, a single bone can be compared to various populations of bones of known pathology to classify a bone automatically. Alternatively, two distinct populations of bones can be statistically examined to highlight inter-population differences. Each of these potential

applications for global shape analysis can have future use for automated feature discovery and automated bone measurement. When the FDR-PCA method described by Algorithm A-1 is applied to the femur, the resulting magnitudes are applied to a color map, which allows the visualization of the distal femur sexual dimorphism directly (Figure 6). It is clear that when femoral scale is removed using the FDR-PCA method, the largest differences between the sexes exist on the ML aspects of the distal femur.

1.2.1.4 Surgical Landmark Assignment

Surgical landmark assignment is the localization of certain points or axes on the bone which are of surgical relevance. These landmarks, during a surgery without a surgical navigation system, allow the surgeon to approximate the correct locations to make surgical cuts and to align cutting jigs. When these landmarks are picked manually, they are subject to observer bias and measurement noise. Ultimately, the landmarks used by a surgical navigation system, whether




Figure 6: Images highlighting the difference between male and female femora with the first principal component left out according to Algorithm A-1, red represents the largest difference, blue relatively little difference. It is clear that the largest differences between males and female occur on the medial and lateral epicondyles

created internally, or by a surgeon, are used for final implant sizing and placement, and thus accuracy in landmark assignment is crucial. Surgical technique varies according to surgeon preference but can influence the necessary compromise between implant size and placement.

Certain landmarks and surgical axes are of paramount importance during the orthopedic surgeries. Often, in the absence of a surgical navigation system, the surgeon will manually identify such landmarks to appropriately place the implant. A novel method using the bone atlas was created to automatically calculate the surgical landmarks in a deterministic manner, removing ambiguity and observer bias typically present in the process [5]. A few of the landmarks and axes automatically calculated for the distal femur are shown in Figure 7.

Using the landmarks and axes previously calculated, the surgical preplanning subsystem can then take a surgeon's chosen implant technique for implant placement (i.e. "Posterior Condylar Axes plus 3 degrees") and automatically place the implant in position, using the correct size and optimizing for minimal overhang. Figure 8 shows the results of the implant placement technique (implant position relative to the bone), which are then fed into the surgical navigation system for surgeon's use during the procedure.




Figure 7: Automatic landmarks and axes calculated on the distal femur

1.2.1.5 Bone Atlas Applications

Beyond the applications previously listed, the statistical shape representation of the bone atlas is useful in other contexts. For example, segmentation is the process of labeling portions of anatomy distinct from the background of medical images. Using a bone atlas as an active shape model, 3D bone segmentation can be performed on CT or MRI medical images [34]. As shown in Figure 9, this methodology can segment highly complex shapes such as L5 lumbar vertebrae.

Again, by using the bone atlas as an active shape model, reconstruction of patient anatomy can be realized from a series of two or more X-ray images. Using Genetic Algorithms (GA) to hypothesize bone position and bone shape, this problem is posed as an 11 DOF non-linear optimization problem [38] [39]. Since the bone atlas is a statistical representation, missing bone




Figure 8: Implant placement technique showing alignment to Posterior Condylar Axis +3 degrees.




Figure 9: Final 3D segmentation of L5 lumbar vertebrae from a CT volume. The orange model is the result of the segmentation algorithm, while the translucent model represents the manual segmentation of an expert.

portions can be extrapolated using the least-squares projection of the known points onto the model. The resulting extrapolation can then be used to measure or identify characteristics of the bone which otherwise would be unavailable, such as overall bone length. Similar to the bone extrapolation and the global shape analysis applications mentioned previously, the bone atlas has applications in anthropological contexts [40] where comparison to modern population statistics is desired [41].

1.2.2 Microwave Positioning Subsystem Overview

The Microwave Positioning Subsystem utilizes a series of high-frequency UWB microwave components to provide real-time 3D positioning information to the surgical navigation system. The system is characterized by repeated transmission of narrow pulse widths (~ 300 ps), which are modulated using an 8 GHz carrier [42]. UWB as a technology is defined by the Federal Communications Commission (FCC) in terms of bandwidth (which must exceed 500 MHz, or have a fractional bandwidth greater than 20%) and typically refers to the FCC's unlicensed band of 3.1-10.6 GHz which can be used for medical as well as industrial purposes. Using the computed Time-Difference-of-Arrival (TDOA) of the pulses at various base stations placed at known locations around the OR (Figure 10), the system computes the real-time 3D position of the surgical instruments and bones [43]. A control station is responsible for processing the microwave signals and finally computing the positions of tags (3D point) or probes (6 DOF). Until recently, UWB positioning systems have been relegated to lower-accuracy applications such as asset and personnel tracking. Several commercial systems and the state-of-the-art research systems are highlighted for comparison in Table 2.



Figure 10: Microwave positioning subsystem overview showing major components: control station, base station, probe, and tag


Table 2: Comparison of commercial and research UWB positioning systems

Positioning System	Type of System	Static Accuracy	Dynamic Accuracy
Ubisense™	Commercial – TDOA	15 cm	N/A
Sapphire Dart™	Commercial – TDOA	10 cm	N/A
Low et al. [44]	Research – 1D only	1 cm	N/A
Meier et al. [45]	Research – Kalman filter, coherent system	0.1 mm	2 mm

1.2.2.1 System Characteristics

The characteristics of the microwave positioning subsystem transmitted signals are shown in Figure 11. The frequency domain depiction shows pulse after up-conversion to 8 GHz and also indicates the 3-4 GHz range bandwidth requisite of the microwave transmitter and receiver components. The time domain graph clearly shows the shape of the 300 ps (full-width half maximum FWHM) pulse.

Given the extremely large bandwidth occupied by the transmitted signals, the minimum sampling bandwidth of the system given by the Nyquist criterion exceeds that which is achievable using low-cost commercial analogue-digital converters (ADC). Consequently, the received signals first pass through a high input bandwidth analog sub-sampler (or sequential sampler), which compresses the pulse bandwidth, effectively stretching the pulse in the time-domain. Figure 12 illustrates an example of the sampled I and Q channels of the 300 ps pulse after passing through the sub-sampler. This graph highlights two of the challenges in accurately recovering the pulse shape, and thus, the correct position. First is the high frequency corruption in both signals which is caused from the use of two distinct Local-Oscillators (LOs) utilized to generate the 8 GHz carrier for up/down conversion. These clock sources cannot be synchronized



**Figure 11: a) Power spectral density of transmitted UWB pulse after up-converted to 8 GHz carrier
b) Time domain representation of normalized pulse shape, noting 300 ps pulse width (full-width-half-max)**




Figure 12: Received signals after down-conversion, sequential analog sampling, and 8-bit AD conversion (red I, blue Q) at a rate of 50 MS/s. The the width of each pulse shown is roughly 3 μ s

as one is located on the transmitter, while one is on the control station. Secondly, the use of two pulse repetition frequency (PRF) clocks—one on the transmitter and one on the receiver—causes clock drift and unintentional time scaling. This effect is analogous to trying to measure time accurately with two clocks that tick at different speeds and are out of synch.

1.2.2.2 Digital Signal Processing

Given the unique characteristics of the system, there are challenges that must be overcome by the onboard digital signal processing (DSP). First, UWB Signals have higher bandwidth than can be adequately sampled directly. Consequently, an analog sub-sampler is used prior to the ADC which time-extends the received pulse. This leads to the following signal processing challenges which will be discussed further [42]: unintentional time scaling at the receiver due to frequency bias between the two asynchronous PRF clocks, jitter in sample point locations, and incoherent LOs. The incoherent LOs cause frequency leakage and misalignment in down converted IQ (Figure 12).

The main steps for the signal processing consist of off-the-shelf components which provide for the digital conversion and digital signal processing (Figure 13). First, each IQ channel is sampled at each base station. The resultant digital signal is passed to a Field Programmable Gate Array (FPGA), which has the primary function of de-noising and ascertaining the time-of-arrival (TOA) of the peak. The FPGA processing first reconstructs the envelope of the original pulse, then performs digital filtering, and then performs peak detection or leading edge detection. The peak position is transformed into a time position index which is transmitted to a secondary FPGA, common to all base stations. The control station FPGA subtracts the time position indices, creating differenced time-of-arrival values. Finally, the control station FPGA serially transmits the values to a PC for final 3D position calculation.




Figure 13: Example of UWB positioning. Transmitted pulse originates from omni-directional tag and received by multiple base stations. After down conversion (not shown) analog sequential sampling reduces the bandwidth of the received signal, which is then digitally sampled, peak detection locates the sample index of the estimated peak. Computing the difference yields time-differences which are passed to the TDOA algorithm for 3D position calculation

1.2.2.3 Communication Protocols

Since the finalized system must be able to provide positioning information on a number of tags or probes, each tag will communicate with the system during a specified time-slice in a Time-division-multiple-access (TDMA) scheme. Tag switching, wake up, and sleep commands are issued by a microcontroller on the control station and communication will occur over a separate wireless technology such as Bluetooth™. Each tag will maintain a unique tag id, the system software running on the PC will maintain tag-to-probe associations and related geometry information.

Ultra-wideband systems can also be used for implementing a communication channel and a large body of UWB literature exposes principles of communication using the RF band. One such is the use of orthogonal-frequency division multiplexing (OFDM) [46], as well as transmitted-reference (TR) [47], and time-reversal [48]. Recent work by K. Liu et al. suggests that concurrent UWB communication schemes are better suited to channel capacity than TDMA methods [49]. While, Zhao and Liu examined rake and prerake systems for UWB communication systems that can handle pulse overlapping and narrow-band interference [50]. Subsequent work by others working on this project is to implement tag switching protocols using the UWB as a communications layer.

1.2.2.4 Final 3D Positioning

The final positioning calculation is completed on a PC, which maintains a list of active tags and tools currently in use for the surgical procedure. The positioning algorithm is one of many multilateration algorithms, Time-Difference-of-Arrival (TDOA), Time-of-Arrival (TOA), or variants of these two, which will be discussed in depth in Chapter 6. The system design

expectations for position fixes is a total of 85 tags at a rate of 24 Hz, or roughly 14 tools with full 6 DOF. The algorithms that perform the peak detection, the 3D positioning calculation, and data filtering have been designed using a pipelined approach. Using this approach, a time delay will certainly occur between the time the signals are all received at the base stations and when the final positioning value is rendered on the computer screen. However, the pipelined approach seeks to maximize throughput, and thus the maximum possible rate of position fixes is optimal given the digital hardware. Restated, for each point received at the system control station a point is calculated and recorded by the system.

2. SURGICAL NAVIGATION SYSTEMS: EVALUATING ELECTROMAGNETIC VERSUS OPTICAL TECHNOLOGY IN THE OR

This chapter is a slightly revised version of a paper that was recently presented at the Annual Meeting of American Academy of Orthopaedic Surgeons, San Francisco, USA, 2008, by Brandon C. Merkl, Michael J. Kuhn, Emam Fatah, Mohamed Mahfouz, David DeBoer [20]. The use of “we” refers to the original authors of the paper as listed above. This work was a preliminary study of static and dynamic accuracy testing for comparison to the developed UWB system.

2.1. Introduction

The use of optical and electromagnetic (EM) tracking for computer assisted surgery (CAS) is a topic which has been widely discussed [51] [52]. Hassan et al. mounted both EM and passive optical sensors to a mechanical robot which simulated upper limb motion. The robot provided a sub-mm reference frame to measure dynamic error of both systems while performing a range of arm motions (e.g. flexion-extension, varus-valgus, etc.). Reported results show that an RMS difference of less than one degree is possible between the passive optical and EM trackers when moving the robot through a realistic compound arm motion, although this required smoothing and least squares fitting of the data [52]. Maletsky et al. reported errors of less than one degree and less than one millimeter in measuring the relative transformation between two rigid bodies with an active optical tracking system [53]. Shuler et al. report errors upwards of 9 mm in testing the ability of an EM tracker to track knee motion under gait analysis [54], although post-

processing of the data should yield better results, similar to [52]. Poulin et al. measured the errors in an EM tracking system due to normal metal interference that may occur in OR during surgery (e.g. lamps, instrument table, arthroscope, etc.) [55]. Finally, Khadem et al. did an extensive analysis of both passive and active optical tracking systems with focus on how system jitter limits accuracy [56].

As shown in [51]- [56], extensive analysis has been done in measuring the accuracy of optical and EM tracking systems in the OR. Beyond just base system accuracy, some research has also focused on how the accuracy of these systems can be used to measure relative transformations, which is ultimately used in kinematics (e.g. [53]). In this work the focus is on how accuracy of the tracking system can affect measuring relevant surgical axes in the knee (e.g. transepicondylar) and the hip (e.g. anterior pelvic plane). First, the basic static and dynamic accuracy of optical and EM systems is discussed, including how metal interference affects EM dynamic error. This is followed by experimental results from simulated hip and knee surgeries, with the static and dynamic error of the tracking systems being translated into errors in the automatic calculation of surgical axes. From these results a quantitative measure is obtained of how the performance of these two systems will affect hip and knee surgeries. When analyzing the effects of these systems' dynamic and static tracking errors in actual surgeries, it is useful to delineate their different roles during surgery in terms of landmarking, bone tracking, bone morphing (picking point clouds), and use during image guidance surgery. Landmarking is largely a static activity while picking point clouds is much more dynamic. This analysis provides a more useful metric for surgeons in terms of how these systems can affect the overall surgical outcome.

2.2. Materials and Methods

2.2.1 Static and Dynamic Tracking Experiment

In order to achieve a baseline measurement for both the EM and optical systems tested, a battery of tests was run to determine both static and dynamic tracking error. The tasks were intended to be realistic to typical OR use of surgical navigation systems. Environmentally, the lab was similar to expected OR conditions; no ferromagnetic shielding was present as this had little effect on the outcome of the accuracy testing [55]. These series of tests were also designed to test the three paradigms of use in modern surgical navigation systems, namely: landmarking, bone tracking, and bone morphing. For a system to yield highly accurate landmarks, static accuracy is of the highest regard. To perform well for bone tracking, both dynamic tracking and static accuracy is requisite. For bone morphing algorithms, dynamic accuracy and the refresh rate are the largest determiners of success.

The first experiment involved landmarking on a synthetic distal femur to test the static accuracy of both systems. Small reference holes were bored in 8 locations on the bone and the bone was mounted in a fixed position relative to the tracking equipment. During this experiment, all metal objects were removed at a distance greater than 1.5m from the bone, which was made entirely of ABS plastic. The landmarks were placed at the surgically relevant locations of the medial epicondyle (TM), lateral epicondyle (TL), antero-medial condyle (AM), antero-lateral condyle (AL), distal-medial condyle (DM), distal-lateral condyle (DL), postero-medial condyle (PM), and posterior-lateral condyle (PL), as shown in Figure 14. At each location, a series of points was digitized and recorded for each system. Care was taken to ensure a good geometric location for the synthetic bone, so that the optical probe had a clear line-of-sight (LOS) with each



Figure 14: a) Synthetic Femur with machined landmark recesses b) Synthetic femur mounted to EM transmitter c) Femur mounted for static optical experiment

landmark position. Using the mean of each digitized point as the true value, the standard deviation of measurement was recorded for each system. Next, pairwise distances between all landmarks were calculated for each system, similar to how a landmark-based navigation system would ascertain relevant surgical axes.

As a second stage to this experiment, a metal platform (~2kg) was brought next to the transmitter and EM sensor, and the landmarking experiment was repeated in order to see the degradation of the system with respect to metal in the work area. Again, each landmark was digitized repeatedly with the mean value representing the true value. The pairwise distances were again measured using this mean value and the distances were compared directly to the static results from the optical and EM systems under the metal free conditions.

The second measurement experiment involved dynamic tracking of the probe on a rotating platform shown in Figure 15.

In this study, the aim was to see how accurately the system could track the sensors while the rotating platform revolved at a constant angular velocity. The points were tracked using both systems and recorded for the platform to perform one full revolution. Since the EM system




Figure 15: Dynamic Tracking Experiment with optical sensor mounted directly to rotating platform

would be affected by the metal present at the rotating platform, a non-ferromagnetic mounting structure was used to separate the EM probe from the rotating platform at a distance of roughly 30 cm. At the onset and end of the revolution, there is an acceleration and deceleration period of the platform, respectively, that is disregarded in this analysis of angular tracking error. In this experiment, the 3D points were acquired and then fit to a plane using orthogonal distance regression. With the plane of rotation known, the center is calculated using a circular regression algorithm, using 2D points that have been projected onto the plane. After the plane and the center of the 2D circle are known, the points are transformed into cylindrical coordinates using a rotating inertial reference frame with constant angular velocity. In this frame, 3 types of errors were measured: radial error, out-of-plane error, and angular tracking error. Since the radius of

the circle is known, the tracking error is transformed into a linear distance at the radius and subsequently recorded.

2.2.2 Hip Cup Tracking Experiment

The alignment of the acetabular component is critical in primary Total Hip Arthroplasty (THA) [57] [58] [59]. Clinical results have shown that anteversion of the cup of $<40^\circ$ and $>60^\circ$ can significantly increase the odds of hip dislocation postoperatively [58] (although the reported anteversion is affected by the flexion angle, so it is better to report an adjusted anteversion angle as done in [59]). In [57], the performance of a mechanical guide for proper cup placement is tested experimentally. By comparing the results of the mechanical guide to a hip navigation system, it is shown that a mechanical guide is insufficient for proper cup alignment.

In this experiment, an optical tracking system (Optotrak 3020, Northern Digital Inc., Toronto [60]) and an EM tracking system (Aurora, Northern Digital Inc., Toronto) were compared in hip navigation by performing a standard postero-lateral THA on cadaveric specimens with both systems, illustrated in Figure 16.



Figure 16: Tracing acetabulum with optical probe

This experiment was undertaken by 6 total graduate level biomedical engineering and anthropology students. Using a posterior approach mock surgery, a cadaveric specimen was opened and the acetabulum was exposed by dislocating the femoral head. In order to obtain the Anterior Pelvic Plane (APP), each observer acquired points on the right Anterior-Superior Iliac Spine (ASIS) by touching the probe tip to a manually palpated position on the unexposed pelvis. The left ASIS point and also the pubis symphysis were recovered by manually picking points on the segmented pelvis with a computer. Subsequently, the observer collected points by tracing the rim of the acetabulum and making an 'X' pattern in the cup. Before performing this mock surgery, each patient was CT scanned and the whole pelvis manually segmented. Figure 17 shows the APP, which was constructed from the two ASIS points and the pubis symphysis, as well as the points traced along the acetabulum overlaid on the segmented pelvis model. Next, the cup plane was approximated using auto regression to fit a plane to the point cloud taken around the circumference of the acetabulum, as shown in Figure 18. Once the APP and cup planes were obtained, the abduction, anteversion, and flexion angles of the cup—which define




Figure 17: Sagittal view of segmented pelvis overlaid by the APP and acetabulum point cloud

the orientation of the acetabulum and are critical for a positive outcome of THA—were calculated. This calculation was done by defining a coordinate frame where the anterior normal of the APP is oriented in positive z , while positive x is defined as the vector pointing from the left ASIS point to the right ASIS point, making positive y orthogonal to x and z and in the superior direction, as shown in Figure 18. Next, a rotation matrix describing the alignment of the cup plane normal (in the anterior inferior direction) relative to the positive z axis was obtained. Three rotations were extracted from the rotation matrix, corresponding to the abduction, anteversion, and flexion of the acetabulum. This was done using a Z-Y-X Euler rotation, corresponding to abduction, anteversion, and flexion. The procedure outlined in [59] for adjusting the abduction and anteversion angles was used since these angles are affected by the amount of flexion native to the acetabulum. This effect was extreme in these results, due to the inherent anatomical differences between the two cadavers studied. Figure 19 shows the




Figure 18: Segmented pelvis overlaid by APP, acetabulum point cloud, and cup plane with APP reference frame




Figure 19: Patient comparison with APPs aligned. (a) Frontal view (b) Sagittal view

segmented pelvises of the two patients analyzed in this study with their anterior pelvic planes aligned. The patient with the red pelvis and red cup axis has roughly 11° of flexion while the patient with the blue pelvis and blue cup axis has approximately 39° of flexion. The large difference in flexion angles significantly shifts the native abduction and anteversion angles, as shown in Table 3 contains the mean cup orientation angles averaged over three trials of EM combined with three trials of Optotrak for each patient. Correcting for the flexion angle requires negating the flexion angle by post-multiplying the overall rotation matrix by a negative rotation of the extracted flexion angle about the fixed x-axis. This removes the flexion angle from the rotation matrix, allowing the correctly adjusted abduction and anteversion angles to be extracted from the rotation matrix. This method was obtained from [59] and the corrected results in agree well with [59].

2.2.3 Knee Axes Determination Experiment

In order to determine the effect of errors in the tracking technology, an evaluation was performed on surgically relevant axes calculated from noisy points. Using a representative bone, with the axes calculated using an automated procedure previously developed [5], synthetic Gaussian noise was added to each of the landmarks from which the axes are calculated. This effectively measures the error due to the positioning system in surgical axes angle calculations. Such measurements serve as the inputs to calculate resection plane orientations; consequently, if the

Table 3: Mean native cup orientation and mean cup orientation corrected for flexion for combined optical and EM tracking experiments (six total trials per patient) in degrees

Patient	Abduction	Anteversion	Flexion	Abduct - Corrected	Antev - Corrected
1	37.3	32.2	11.2	40.4	25.2
2	76.8	48.8	39.6	41.1	24.6

axes exhibit a significant amount of error, the feasibility of using a positioning system for axes calculation may be called into question.

In this experiment, the worst-case standard deviations of measurement from the static landmarking experiment were recorded for the optical system, the EM system in the absence of metal, and the EM system with metal nearby. The standard deviation was used to create zero mean Gaussian noise around each of the landmarks of the representative bone. Subsequently, the axes orientations were calculated and compared to the robust automatic measurements. The data was collected for the following axes: Mechanical Axis (MA), Posterior Condylar Line (PCL), Distal Bearing Line (DBL), Transepicondylar Axis (TEA), Anterior Capsule Axis (ACA), and Distal Anatomic Axis (DAA).

2.3. Results

2.3.1 Positioning Results - Static

Comparing the results from the static landmarking experiment (Table 4, Table 5, and Table 6), the salient results indicate that the optical system yields the least landmark uncertainty. Following the optical system, the EM system without metal present proceeded to provide similar results, though the standard deviations are roughly an order of magnitude larger. When metal is introduced to the EM system setup, the performance degrades significantly, increasing the standard deviation by another order of magnitude to 0.25 mm on average.

The pairwise landmark distances were calculated using each of the positioning scenarios. Some of these distances (e.g. medial transepicondyle to lateral transepicondyle) represent a surgical

Table 4: Optical static landmark results in mm

	Mean X	Mean Y	Mean Z	Std. X	Std. Y	Std. Z
AL	-545.37	84.32	-1846.61	0.002	0.001	0.003
AM	-536.64	52.11	-1839.69	0.017	0.014	0.033
DL	-536.64	89.12	-1807.93	0.017	0.009	0.031
DM	-525.70	89.12	-1801.64	0.023	0.020	0.048
PL	-545.32	88.43	-1789.97	0.011	0.006	0.021
PM	-547.22	52.06	-1784.11	0.013	0.007	0.018
TL	-551.21	103.36	-1813.45	0.008	0.007	0.013
TM	-548.54	30.29	-1810.68	0.012	0.005	0.011

Table 5: EM static landmark results (no metal) in mm

	Mean X	Mean Y	Mean Z	Std. X	Std. Y	Std. Z
AL	72.02	-28.48	-113.60	0.0628	0.0214	0.0684
AM	58.56	1.020	-121.87	0.0198	0.0000	0.0684
DL	36.61	-43.04	-137.86	0.0857	0.0569	0.0575
DM	20.58	3.700	-138.58	0.1110	0.1084	0.0858
PL	16.51	-50.82	-119.33	0.0122	0.0611	0.0532
PM	4.470	-11.27	-120.18	0.0167	0.0291	0.0539
TL	45.00	-55.41	-113.95	0.0369	0.0503	0.0505
TM	22.28	13.74	-112.21	0.0509	0.0000	0.0414

Table 6: EM static landmark results (with metal table) in mm

	Mean X	Mean Y	Mean Z	Std. X	Std. Y	Std. Z
AL	358.64	165.05	-303.70	0.185	0.246	0.200
AM	352.62	138.04	-317.10	0.211	0.257	0.253
DL	390.82	169.36	-336.70	0.266	0.447	0.263
DM	377.19	122.98	-344.89	0.301	0.189	0.285
PL	417.94	160.83	-328.15	0.231	0.264	0.169
PM	408.83	120.52	-333.95	0.248	0.218	0.206
TL	391.05	180.66	-319.72	0.220	0.336	0.398
TM	375.07	113.20	-319.20	0.194	0.354	0.237

axis while others do not. As indicated in Table 7, Table 8, and Table 9, the pairwise distance differences show that the optical measurements agree best with the EM system when no metal is present. Some of the pair-wise errors between the systems can be explained when considering that the probe tips, being slightly different diameters, fit differently into the machined holes. The EM system in the presence of metal does not agree well with either the optical system or the EM system without metal. The standard deviations listed in Table 6 cannot be thought of as accuracy, but rather system precision under static conditions.

Table 7: Optotrak static distances vs. EM static differences in mm

	AL	AM	DL	DM	PL	PM	TL
AM							-0.64
DL						2.78	2.30
DM					-2.28	-0.62	-2.74
PL				3.33	4.64	3.66	2.06
PM			-0.34	-1.05	0.57	-2.71	4.46
TL		-0.52	-0.81	0.66	-3.16	0.92	1.02
TM	0.30	1.48	1.66	2.70	3.41	-2.65	-3.85

Table 8: Optotrak static distances vs. EM static distances (with metal) in mm

	AL	AM	DL	DM	PL	PM	TL
AM							-3.33
DL						3.75	3.67
DM					-6.16	-2.20	-3.13
PL				7.50	7.88	4.92	2.18
PM			3.25	4.66	4.17	2.18	4.84
TL		0.70	-1.95	-7.14	-7.25	6.02	5.07
TM	-8.40	-4.66	-1.38	1.99	2.90	3.19	-3.85

Table 9: EM static distances vs. EM static distances (with metal) in mm

	AL	AM	DL	DM	PL	PM	TL
AM							-2.70
DM						0.97	1.38
DL					-3.88	-1.58	-0.38
PL				4.16	3.23	1.27	0.12
PM			3.59	5.71	3.60	4.89	0.38
TL		1.22	-1.14	-7.80	-4.10	5.10	4.05
TM	-8.70	-6.14	-3.04	-0.71	-0.51	5.84	-3.49

2.3.2 Positioning Results - Dynamic Results

The results in Table 10 and Table 11 clearly indicate a large difference in dynamic errors between optical and EM technologies. Comparing the standard deviations of the radii, out-of-plane, and tracking errors, the optical probe yields higher accuracy and consistency than the EM probe. The optical results for the measures of radii and out-of-plane errors are less than 0.1 mm, while the tracking error exhibits the largest magnitude. The EM experiments indicate that the tracking error exhibits the opposite trend, as the standard deviation of radii and out-of-plane errors are slightly larger. Overall, the EM system had roughly 1 mm of error in each of the cylindrical coordinate directions. The tracking error corresponds well with the larger uncertainty in angular velocity as measured by the EM system relative to the optical system.

2.3.3 THA - Acetabulum Orientation Results

The results in Table 12 show that the error between using EM and optical tracking systems in

Table 10: Optical dynamic tracking results

Exp #	Std. Radii (mm)	Mean Radii (mm)	Std. Out-of-plane (mm)	Mean Ang. Vel. (radians/sec)	Std. Ang. Vel. (radians/sec)	Std. Track Error (mm)
1	0.030	97.6	0.0721	1.12	0.0219	0.104
2	0.028	97.6	0.0733	1.12	0.0238	0.106
3	0.029	97.6	0.0744	1.12	0.0228	0.105

Table 11: EM dynamic tracking results

Exp #	Std. Radii (mm)	Mean Radii (mm)	Std. Out-of-plane (mm)	Mean Ang. Vel. (radians/sec)	Std. Ang. Vel. (radians/sec)	Std. Track Error (mm)
1	0.839	99.9	0.990	1.13	0.234	0.918
2	1.040	100.	1.01	1.13	0.248	0.893
3	1.050	100.	1.03	1.12	0.297	0.835
4	0.964	99.8	1.03	1.13	0.224	0.836

calculating the corrected abduction and anteversion angles in hip navigation is negligible. The corrected abduction and anteversion angles in Table 12 agree well with the results obtained in [59]. It should be noted that the sampling rate for the EM system was three to five times larger than that of the optical system, resulting in 3-5x as many points obtained in picking points and tracing the acetabulum for the EM system when compared to the optical system. This means that the error in picking points with the EM system is higher than optical, but since this can be solved with a higher sampling rate for the EM tracking system, it does not pose a problem in practice.

2.3.4 TKA - Axes Results

The results in Table 13 show that for all positioning scenarios, shorter axes (ACA, PCL, and DBL) experience greater error in angular orientation than longer axes due to uncertainties. The longest axis was the MA. This axis experienced the least amount of orientation error relative to the other axes. In general, the optical system had the greatest performance with respect to all axes when compared to the EM system. The presence of metal near the EM system increased the maximum error by a factor of about four. For the longer axes, this amounted to negligible error; however, the shorter axes had maximum errors of 2.43-4.76 degrees. Considering TKA placement protocols that call for PCL alignment plus 3 degrees, this magnitude is deemed significant.

Table 12: Mean and standard deviation of corrected cup orientation for optical and EM tracking systems averaged over both patients (in degrees)

Tracking System	Abduct - Corrected	Antev - Corrected	Flexion - Patient 1	Flexion - Patient 2
Optical	39.3 +/- 1.91	24.0 +/- 2.24	11.2 +/- 1.36	37.6 +/- 1.11
EM	40.8 +/- 0.92	24.9 +/- 2.64	11.4 +/- 0.66	39.2 +/- 2.99

Table 13: Static errors of positioning systems and the contribution to surgical axis error

Exp Type		MA (deg)	ACA (deg)	PCL (deg)	DBL (deg)	TEA (deg)	DAA (deg)
Optotrack	Std. Error	0.00	0.08	0.10	0.05	0.03	0.02
	Max Error	0.02	0.34	0.51	0.21	0.14	0.09
EM no metal	Std. Error	0.01	0.18	0.28	0.11	0.06	0.04
	Max Error	0.05	0.89	1.21	0.53	0.29	0.22
EM w/ metal	Std. Error	0.04	0.66	1.11	0.42	0.24	0.18
	Max Error	0.20	2.76	4.76	2.43	1.19	0.93

2.4. Discussion

Tested were two competing tracking technologies used in surgical navigation systems through measuring static and dynamic tracking error and quantitatively studying how these errors affect automatic axes calculation during TKAs and THAs. Both systems underwent static, dynamic, and cadaveric testing. The results from the static error analysis also lead to a simulated testing of surgical axes calculations. Both systems proved adequate in the static results, although the performance of the EM system seriously degraded when metal was placed in proximity to either the transmitter or the EM sensor. In dynamic testing, the optical system provided an order of magnitude higher accuracy over the EM system. Under these conditions, both systems would be adequate for range-of-motion measurement, especially considering the refresh rate of the EM system is typically kept three to five times higher than that of the optical. The increased number of samples effectively reduces the observed tracking error of the EM system down to levels comparable to optical tracking when the EM data is down-sampled to the refresh rate of the optical system. This effect, as mentioned in the hip cup experiment, caused the EM and optical systems to perform similarly in hip navigation as there was little difference between the systems

in the final cup axis calculation. However, when both systems were used to calculate surgical axes based on known landmarks, the optical system's higher accuracy created a negligible difference in angular orientation with respect to automatically calculated axes. Metal placed in proximity to the EM system caused a noticeable degradation in surgical axes calculation. Of the three primary uses of surgical navigation systems, namely landmarking, bone morphing, and dynamic bone and instrument tracking, the optical system was superior or comparable to the EM system. The EM system has the distinct advantage of not requiring line-of-sight between the tracking sensors and the system. As a result, it is the finding of this study that optical systems can be used in large incision total joint surgeries or in instances where the highest levels of accuracy are needed. EM systems have lower performance when used in highly dynamic or ferromagnetic conditions but have advantages with point acquisition and thus remain viable for bone morphing and static landmarking. Finally, both systems have certain disadvantages, leaving the door open for new tracking technologies, such as wireless tracking, assuming new systems can reduce error levels to the accuracy of EM and optical systems.

3. INVESTIGATION OF HIGH ACCURACY INDOOR 3-D POSITIONING USING UWB TECHNOLOGY

This chapter is a slightly revised version of a paper that was recently published in the journal *IEEE Transactions on Microwave Theory and Techniques* by Mohamed Mahfouz, Cemin Zhang, Brandon Merkl, Michael Kuhn, and Aly Fathy [42]:

Investigation of High Accuracy Indoor 3-D Positioning Using UWB Technology

Mohamed R. Mahfouz, *Senior Member, IEEE*, Cemin Zhang, *Student Member, IEEE*, Brandon C. Merkl, *Student Member, IEEE*, Michael J. Kuhn, *Student Member, IEEE*, Aly E. Fathy, *Fellow, IEEE*

The use of “we” refers to the original authors of the paper as listed above. My specific contributions to this work, in order of importance, include: (1) development and description of the Iterative Peak Detection algorithm, (2) study and literature review of existing channel modeling techniques in conjunction with receiver-side peak detection, (3) experimental testing of various peak detection techniques, (4) development and description of PDOP calculation and resulting analysis of system error as a result of PDOP, (5) analysis of unintentional time scaling, (6) testing of modified TDOA algorithm to mitigate time scaling, (7) theoretical time scaling calculations.

Abstract — There are many challenges in building an ultra wideband (UWB) indoor local positioning system for high accuracy applications. These challenges include reduced accuracy

due to multipath interference, sampling rate limitations, tag synchronization, and antenna phase center variation. Each of these factors must be addressed to achieve mm or sub-mm accuracy. The developed system architecture is presented where a 300 ps Gaussian pulse modulates an 8 GHz carrier signal and is transmitted through an omni-directional UWB antenna. Receiver-side peak detection, a low cost sub/sequential-sampling mixer utilizing a direct digital synthesizer, high fidelity 10 MHz crystals, and Vivaldi phase center calibration are utilized to mitigate these challenging problems. Synchronized and unsynchronized experimental results validated with a sub-mm accurate optical tracking system are presented with a detailed discussion of various system errors.

3.1. Introduction

Wireless local positioning has many diverse applications and has been extensively studied [60] [61]. While Global Positioning Systems (GPS) use ultra high precision atomic clocks to measure time-of-flight [62], a more standard method for indoor localization systems is Time Difference of Arrival (TDOA), where all of the base stations or receivers are synchronized, and the difference in time is measured between each pair of receivers to triangulate the position of an unsynchronized tag [60]. Two main technologies have emerged as possible solutions for TDOA systems: frequency-modulated continuous-wave (FMCW) and UWB. FMCW systems can be found both in the literature [63] [64] [65] and as commercial products [66] with various levels of accuracy. For example, Stelzer et al. achieved accuracy of greater than 10 cm for an outdoor application tracking a car around a 500 m² racecar track [63] while Wiebking et al. achieved accuracy around 20 cm for an indoor application covering a 15x25 m² 2-D area [67]. Finally, Roehr et al. achieved accuracy of 1 cm in a line-of-sight (LOS), multipath free environment

using a novel chirp technique centered at 5.8 GHz [68]. Meanwhile, interest in UWB for radar applications has increased greatly following the Federal Communications Commission's (FCC) decision to open up the bands from 3.1 - 10.6 GHz and 22 – 29 GHz for UWB use in 2002 [43]. UWB technology has inherent advantages for indoor applications in terms of robustness to multipath interference and potential for high ranging accuracy [69].

Commercial systems are already available which utilize UWB sensors for indoor asset tracking. For example, Multispectral Solutions, Inc. has the Sapphire DART™ system which can operate in indoor environments of greater than 50 m ranges while using TDOA in conjunction with UWB sensors to achieve accuracy of 10 cm [70]. Ubisense has a comparable indoor system which uses UWB in conjunction with TDOA and Angle of Arrival (AOA). It has an operating range of up to 50 – 100 m and can detect sensors with accuracy of 15 cm [71]. Higher accuracy has been reported in the literature for indoor UWB positioning systems. For example, Low et al. achieved centimeter-range accuracy in a 1-D short range indoor LOS environment utilizing UWB pulse signals [44]. Zetik et al. reported sub-mm 1-D accuracy but with only extremely short displacements while accuracy decreased to 1.5 cm for 2-D localization over a $2 \times 2 \text{ m}^2$ area [72]. Recently, Meier et al. designed a 24 GHz system which uses a Kalman filter combined with correlation and phase information to reduce the uncertainty of a static point to 0.1 mm, although the uncertainty increases to 2 mm when the tag is in motion [45]. These experimental results show that UWB technology has the potential for high precision indoor localization even in harsh environments with significant multipath effects.

Commercial UWB systems have localization accuracy of 10 - 15 cm. Certain short-range industrial and medical applications such as dynamic part tracking, structural testing, and

computer-assisted therapy require significantly higher accuracy than commercial UWB systems. Current technologies used for these applications include infrared (IR), electromagnetic (EM), and ultrasound tracking, which have mm or even sub-mm accuracy. However, IR has a short transmission range and can be easily disturbed by a fluorescent lamp or other light sources in a room. EM has reduced performance near metal. Finally, ultrasound has limitations due to multipath interference because of its limited bandwidth compared to UWB. UWB tracking systems have inherent advantages over these existing technologies since UWB does not suffer from the drawbacks mentioned.

However, there are many challenges in developing such a real-time indoor UWB localization system which has accuracy orders of magnitude greater than existing commercial systems. These challenges include multipath interference, sampling rate limitations, synchronization errors, and antenna phase center error. Errors associated with these design issues must be accounted for to develop a system for these high accuracy applications. In this chapter, a system is proposed [73] [74] and various system-level design issues in the context of prototyping a high accuracy UWB localization system are discussed. Through advanced sub-sampling techniques, receiver-side signal processing, a modified TDOA algorithm, and antenna phase center calibration, mm-range accuracy in a real-time system is possible as will be demonstrated here.

This chapter is organized as follows. In Section 3.2 an overview of the positioning system is presented including a block diagram of major system components, the mobile tag architecture, and a detailed discussion of the sub-sampling mixer. In Section 3.3, synchronization issues are addressed with focus given to how they affect overall system accuracy. Section 3.4 details the peak detection algorithm used to identify pulse peak position in dense multipath environments.

Experimental results are included to show how multipath interference contributes to system error. The effects of phase center variation of the base station antennas on system accuracy are discussed in Section 3.5. Synchronized and unsynchronized localization experiments including results are outlined in Section 3.6. Finally, Section 3.7 concludes.

3.2. System Architecture

3.2.1 Overview

A GPS-like scheme is utilized along with TDOA to locate 2-D and 3-D transmitting tag positions, as shown in Figure 20. The average output power spectral density for indoor systems has an upper bound of -41.3 dBm/MHz [43] as specified by the FCC. One typical UWB localization method is the use of Impulse Radio (IR) UWB, where a baseband UWB pulse is filtered by a UWB antenna to conform to FCC regulations [44]. However, with IR-UWB the received signals are noisy due to the complex transmitted waveform and added multipath signals, which makes it difficult to accurately locate the position of the received LOS signal. In the system, we modulate an UWB pulse with an 8 GHz carrier signal which resides at the upper end of the 3.1 – 10.6 GHz band. The use of this band reduces the size of the wideband RF components in the transmitter and receiver and also bypasses many of the interfering frequency bands that exist at the lower end of the 3.1 – 10.6 GHz band.

The complete experimental setup of the developed system is shown in Figure 21. In the developed system, we transmit a modulated narrow Gaussian pulse with a carrier frequency and demodulate it at the receiver side. The source of the UWB positioning system is a step-recovery diode (SRD) based pulse generator with a pulse width of 300 ps and bandwidth of greater than 3




Figure 20: GPS system analogy: (a) GPS system, (b) UWB indoor positioning system.

GHz, shown in Figure 22. A detailed discussion can be found in [74]. At the input of the traditional delay-line type pulse generator [75], a novel input matching network has been introduced to prevent pulse echoing, minimize pulse width broadening, and suppress any significant pulse distortion. The modulated Gaussian pulse is then transmitted through an omnidirectional UWB antenna. Multiple base stations are located at distinct positions in an indoor environment to receive the modulated pulse signal. The received modulated Gaussian pulse at each base station first goes through a directional Vivaldi receiving antenna and then is amplified through a low noise amplifier (LNA) and demodulated to obtain the I/Q signals. After going through a low pass filter (LPF) with a passband of DC-5 GHz to suppress the 8 GHz carrier signal, the I/Q signals are sub-sampled using an UWB sub-sampling mixer, extending them to a larger time scale (i.e. μs range) while maintaining the same pulse shape. The sub-sampling mixer uses extended time techniques to achieve equivalent sampling rates in excess of 100 GS/s, which yields mm-range sample spacing and provides the peak detection algorithm with ample data. Finally, the extended I/Q signals are processed by a conventional analog to digital converter (ADC) and standard Field Programmable Gate Array (FPGA) unit.




Figure 21: Block diagram of indoor localization system showing one tag and three base stations which feed into the main system controller.




Figure 22: Gaussian pulse which serves as system UWB source: (a) time domain exhibiting 300 ps pulse width, (b) frequency domain highlighting bandwidth in excess of 3 GHz.

The transmitting tag uses an elliptical monopole structure combined with a modified ground plane which provides an omni-directional radiation pattern [76]. The base stations use a single element Vivaldi antenna with a flared-antipodal design which has demonstrated almost a constant gain and beamwidth over a wide band [77] [78].

The output power spectral density of the transmitted signal in the system has been measured and plotted in Figure 23. The modulated pulse signal has a 10 dB bandwidth of approximately 6 GHz, exceeding the 500 MHz minimum bandwidth required under the FCC rules governing UWB. Its average output power spectral density satisfies the FCC indoor limit by a margin of more than 3 dB for a majority of the useable bandwidth. The 8 GHz carrier signal leaks through the mixer and is shown as a peak at -16 dBm in Figure 23. This leakage could be suppressed by adding a band-notched filter or utilizing a band-notched monopole [79].

A clear Gaussian pulse signal with a signal-to-noise ratio (SNR) of about 20 dB was detected at a range of up to 3.8 m. The sampler, described in Section 3.2.2, has a sensitivity of -45 dBm [80]. Figure 21 shows the receiver chain. The two amplification stages increase the signal by 25 dB while the I/Q down-converter has 8 dB conversion loss. This makes the overall sensitivity of each base station around -62 dBm. The dynamic range of the system is limited by the sampler and is over 50 dB. The proposed system can be extended to more than 5 m considering the transmitted signal spectrum has a margin of 3 dB below FCC limitations. Figure 24 shows the current layout of the tag. A microcontroller is currently used to implement a Time Division Multiple Access (TDMA) scheme for communicating with multiple tags. TDMA is the preferred multiple access scheme in current UWB commercial systems [70] [71]. The modulation scheme used is On-Off Keying (OOK) with the 8 GHz carrier signal, although pulse modulation could be used to increase system dynamic range. A unique ID is stored on each tag. Each tag is in a low power or sleep state until activated by the main control station. The control station calls each tag in a round robin fashion.




Figure 23: Power spectral density of modulated pulse signal showing double sideband modulated signal with bandwidth of 6 GHz and carrier leakage at 8 GHz of -16 dBm.

3.2.2 UWB Sub-Sampling Mixer

To detect narrow pulses on the order of a few hundred picoseconds (i.e. 300 ps or 3 GHz bandwidth in the system), analog to digital converters with at least 6 GS/s are needed to satisfy the Nyquist criterion. However, such high performance ADC units are currently either not commercially available or too expensive for most applications. A realistic alternative approach to real-time sampling is to sub-sample the UWB pulses while maintaining the initial pulse shape through extended time techniques. The extended UWB signals can then be handled by readily available commercial ADCs, reducing overall system cost [81] [82]. We have developed a compact UWB sub-sampling mixer integrated with a fast strobe generator for this purpose [80]. The sampler utilizes a simple broadband balun structure [83] and a balanced topology. Figure 25 shows the schematic of the designed sampler.




Figure 24: Block diagram of current tag layout showing OOK digital communication and UWB transmitting architecture.




Figure 25: Schematic of the sampling mixer highlighting the broadband balun and balanced topology [80].

The designed sampler based on the sub-sampling method uses the same technique as the commercial off-the-shelf (COTS) high speed sampling modules, such as the Tektronix 80E01 [84] and Picosecond 7040 [85]. The comparisons between the designed sampler and the commercial sampling modules are listed in Table 14. The Picosecond 7040 sampling module, for example, utilized the nonlinear transmission line (NLTL) technique to achieve a fast rise time of 14 ps, and the associated bandwidth of 25 GHz. The 90 ps rise time of the designed sampler is limited by the SRD transition time and the utilized microwave integrated circuit (MIC) fabrication. However, the designed sampler with 4 GHz bandwidth is adequate to sample the 300 ps pulse signal shown in Figure 22a. The compact size of the sampler compared to the commercial samplers in Table 14, combined with the low fabrication cost, make it an attractive and necessary alternative to current COTS products. An even wider bandwidth could be potentially achieved by setting a proper gating duration through adjusting the strobe impulse amplitude [80].

For measurement, a direct digital synthesizer (DDS) circuit is used to trigger the strobe-step generator with a pulse repetition frequency (PRF) of $PRF_2 = f_o \pm \Delta f$. The original transmitted Gaussian pulse has a repetition frequency of $PRF_1 = f_o = 10$ MHz. If the added offset frequency

Table 14: Comparison of the Designed Sampler with Commercial Sampling Modules

	Tektronix 80E01	Picosecond 7040	Designed sampler
Rise time	7 ps	14 ps	90 ps
Bandwidth	50 GHz	25 GHz	4 GHz
Sampling rate	N/A	Up to 10 MHz	Up to 35 MHz
Fabrication	N/A	MMIC	MIC
Size (cm ³)	13.5x7.9x2.5	5.1x3.8x1.3	4.3x3.0x2.0
Cost	Very High	High	Very Low

(Δf) is set to be 100 Hz, the corresponding extending ratio is $\alpha = f_o / \Delta f = 100,000$. Noticeable ranging errors can occur during the time extension process due to the stability of PRF_1 and PRF_2 . This is due to drifting between the two clocks during the 10 ms interval needed to time extend the signal. Section 3.3 discusses how the clock stability affects system accuracy.

3.3. Synchronization

There are three synchronization issues for the proposed localization system in Figure 21: 1) LO and PRF synchronization between receivers, 2) carrier frequency synchronization between the tag and receiver, and 3) PRF_1 and PRF_2 synchronization between the tag and receiver.

Synchronization between different receivers is solved by wiring the N receivers together to use a single crystal clock (PRF_2) for transmission of clock data output from DDS. Also, a single carrier (LO_2) is used between all the base stations.

3.3.1 Tx-Rx Carrier Synchronization

Carrier frequency synchronization between transmitting tag and receiver is solved by direct I/Q down-conversion. The transmitted signal $s(t)$ from the tag is given by

$$s(t) = \sin(\omega_c t)(p(t) + K) \quad (1)$$

where $p(t)$ is the Gaussian pulse signal, K is the carrier signal leakage factor, and ω_c is the carrier frequency generated by LO_1 . The received I and Q signals before passing through the low pass filter are given by

$$\begin{aligned} I &= [\sin(\omega_c t) \cdot (p(t) + K)] \cdot \sin(\omega_c t + \Delta\omega t) \\ &= \frac{1}{2} [p(t) + K] \{ \sin(2\omega_c t) \sin(\Delta\omega t) + [1 - \cos(2\omega_c t)] \cos(\Delta\omega t) \} \end{aligned} \quad (2)$$

$$\begin{aligned}
Q &= [\sin(\omega_c t)(p(t) + K)] \cdot \cos(\omega_c t + \Delta\omega t) \\
&= \frac{1}{2} [p(t) + K] \{ \sin(2\omega_c t) \cos(\Delta\omega t) - [1 - \cos(2\omega_c t)] \sin(\Delta\omega t) \}
\end{aligned} \tag{3}$$

where $\Delta\omega$ is a small offset frequency of LO_2 relative to the carrier ω_c generated by LO_1 from the tag. After passing through the LPF with a passband of DC-5 GHz which suppresses the 8 GHz carrier signal, the I and Q signals become

$$I = \frac{1}{2} [p(t) + K] \cos(\Delta\omega t) \tag{4}$$

$$Q = \frac{1}{2} [p(t) + K] \sin(\Delta\omega t) \tag{5}$$

Finally, the filtered I and Q data are sampled and processed by the FPGA, and the reconstructed received signal is given by

$$\bar{p}(t) = \sqrt{I^2 + Q^2} = \frac{1}{2} [p(t) + K] \tag{6}$$

From (6), the recovered signal $\bar{p}(t)$ is not affected by the offset carrier frequency ω_c and contains the same information as the transmitted Gaussian pulse signal $p(t)$ with a DC bias $K/2$ which can be blocked using a DC block after the LPF. To validate the above analysis and study how the offset frequency $\Delta\omega$ between the tag and base stations would affect the signal performance, a system level simulation using Agilent ADS2006A has been carried out. The simulation in Figure 26 was conducted with four randomly chosen offset frequencies, 19.9 MHz, 53 MHz, 100 MHz and 118 MHz. As shown in Figure 26, the original pulse has been successfully reconstructed with four different LO offset frequency conditions.




Figure 26: ADS simulation showing the effect of Tx-Rx LO frequency offset on reconstructed sub-sampled pulse: a) original signal, b) reconstructed signal with LO offsets of 19.9 MHz, 53 MHz, 100 MHz and 118 MHz.

3.3.2 Tx-Rx PRF Synchronization

The PRF_1 and PRF_2 clock signals between the tag and receiver in this system are not synchronized. This results in an interesting synchronization problem when incorporating the sub-sampling mixer (discussed in Section 3.2.2) since 100,000 pulses are needed to extend the pulse from 300 ps to 30 μ s, which corresponds to a 1000 GS/s sampling rate. In the prototyped system, the extended time signal can then be adequately sampled with a 50 MS/s ADC, given the bandwidth of the sub-sampled signal is 0.03 MHz, although in the actual system the 50 MS/s ADC can be replaced by a higher performance ADC (e.g. 24 bit, 250 MS/s). This can potentially increase system performance by reducing quantization error through extending the dynamic range and increasing the sampling rate with little difference in chip and manufacturing costs.

Techniques exist to calculate clock jitter in the time domain given the phase noise of the crystal [86], and crystal manufacturers provide crystal stability specs in terms of parts-per-million (ppm). On the surface, the stability factor would appear to have less of an effect than actual clock jitter. For instance, the stability factor of ± 0.5 ppm for the 10.0 MHz Vectron VTC4-A0AA10M000 crystal yields clock jitter of 50 fs [87]. However, the effects of the stability factor are amplified during the sub-sampling process and cause time-scaling to occur.

Any frequency offset occurring between PRF_1 and PRF_2 would cause apparent time-scaling in the sub-sampled signal. The ± 0.5 ppm stability of the crystal equates to 5 Hz of variation in the clock signal. This causes the nominal frequency of PRF_1 to be $10.000000 \text{ MHz} \pm 5 \text{ Hz}$ and the nominal frequency of PRF_2 to be $9.999900 \text{ MHz} \pm 5 \text{ Hz}$, with the worst case scenario being a difference of 10 Hz between the two signals. If the offset frequency (Δf) is set to be 100 Hz, the corresponding extending ratio is $\alpha = f_o / \Delta f = 100,000$. However, with a potential offset of ± 10 Hz, Δf has the potential to be $100 \pm 10 \text{ Hz}$, thus $\alpha = 100,000 \pm 11,111$. This time scaling may be time-varying depending on the thermal stability of the clocks, although this is a slow variation with a drift rate of less than 10 Hz, which allows the TDOA algorithm to calibrate such systematic error. Since all the base stations acquire synchronous samples clocked by PRF_2 , the time-scaling effect will be unknown, but consistent across receivers. Consequently, the 1-D ranging errors will be roughly $\pm 11\%$. The TDOA algorithm using time differences will likewise be affected. A simulation was enacted to test the ability to calibrate out the scale factor in a post-processing step. In this simulation, 4-8 base stations in realistic constellations were used to localize a simulated tag trajectory that varied within a volume defined by half of the constellation radius. Using an iterative approach, the TDOA algorithm was allowed to calculate a position using an estimate of the scaling factor, initially set at 100,000. The calculated location was then used to determine a new least squares estimate of the scaling factor, based on the original ranges. Iterating these steps between 3-5 times yielded accuracy consistent with the system's 1-D accuracy of 1.49 mm of error, as discussed in Table 17. This problem has been examined in the field of geolocation [88], and has been likened to multidimensional scaling [89]. In these contexts, this time-scaling problem has been successfully solved even in situations when each

receiver undergoes unique scaling, which may indicate that base stations synchronized by PRF_2 are not necessary.

Besides stability, the jitter of these 10 MHz clocks must be examined. When the phase noise is integrated from 1 Hz – 5 MHz using an Agilent E5052A Signal Source Analyzer, RMS jitter for the crystal is found to be between 1 – 1.5 ps [87] [90]. This technique is the most accurate way to measure clock jitter of highly stable crystals. If 1.5 ps of RMS jitter is assumed for both 10 MHz crystals, total system RMS jitter σ_{sys} due to the two unsynchronized clocks of 2.12 ps is obtained through

$$\sigma_{sys} = \sqrt{\sigma_{clk1}^2 + \sigma_{clk2}^2} \quad (7)$$

where σ_{clk1} and σ_{clk2} are assumed to be uncorrelated normal random variables [86] [91] of mean $\mu = 0$ and standard deviation $\sigma = 1.5$ ps. The jitter described in (7) will cause normally distributed noise of $\mu = 0$, $\sigma = 2.12$ ps (corresponds to 0.64 mm error) to be added to each sampled point. A simulation using Agilent ADS2006A has been carried out to study how such jitter affects the sampler performance during the sub-sampling process. Based on the ADS simulation results shown in Figure 27, we notice that such jitter would cause tiny signal distortion, and could be significantly reduced after simple digital processing such as a low pass filter.

Figure 28 outlines the experimental setup used to test the jitter effect caused by the unsynchronized PRF clock signals. A coaxial cable was utilized to connect the pulse generator and the sampler so that no channel noise is included. Figure 29 shows the measured time variation of the pulse peak position at consecutive measurement cycles. The peak to peak variation is below 10 ps while the RMS jitter is 3.48 ps. The measured RMS jitter, of 3.48 ps, is 1.36 ps larger than the theoretical result of 2.12 ps, which can be interpreted as the added jitter




Figure 27: ADS simulation of the reconstructed sub-sampled pulse with and without 3 ps of PRF clock jitter.




Figure 28: Experimental setup with the unsynchronized PRF clock sources to measure the effect of PRF clock jitter.




Figure 29: Time variation of the pulse peak position acquired over n sample points.

from the sub-sampling mixer, DDS, and ADC circuitry. The measured system clock jitter of 3.48 ps corresponds to 1.05 mm error.

3.3.3 Temperature Effects on Synchronization

The variation in phase noise (and hence jitter) is negligible for both the 10 MHz crystals and also the 8 GHz LOs over the anticipated operating temperature range of 10°C to 30°C. The 10 MHz crystal has stability of ± 0.5 ppm over this temperature range. Although this introduces time-scaling, this effect can be removed with a revised TDOA algorithm, as mentioned in Section 3.3.2. The 8 GHz LOs in the current setup do vary in frequency from 7.983-8.017 GHz over the anticipated operating temperatures. These effects can be mitigated by using I/Q down-conversion (outlined in Section 3.3.1) given that the temperature has a slow rate of variation or by using the phase locked dielectric resonator oscillator (PLDRO) with extremely low phase noise and minimal temperature sensitivity.

3.4. Receiver-Side Peak Detection

Much study and devoted work has been accomplished in the field of indoor UWB positioning. The foundations of this work lie with the exploration of UWB channel models as a description of the actual propagation expected in indoor environments. To develop a highly accurate positioning system it is necessary to construct an algorithm with knowledge of expected propagation effects. Beginning with the Saleh-Valenzuela (S-V) indoor propagation model [92], many adjustments and refinements have been added. Cramer, Sholtz and Win [93] used arrays of antenna measurements to better model angle-of-arrival and waveform shape. Kunisch and Pamp [94] used a battery of measurements to model the frequency dependence of the channel impulse response. Building on this concept, Irahauten, et al. [95] modeled the frequency-domain

impulse response with an autoregressive (AR) model. Wong and Lee [96] examined a minimum description length (MDL) method to extract the best fitting model while minimizing the number of parameters. Recently Sibille et al. [97] looked into incorporating antenna effects by looking at the joint antenna-channel problem concluding that channel angular dependence must be taken into account. Also, recently the IEEE 802.15.4a task group approved an extended and greatly parameterized version of the S-V channel model [98].

Further study has been undertaken to translate channel models into achievable ranging systems. The core problem to this task is achieving high-accuracy pulse peak detection or leading edge detection. Relevant methods to UWB positioning in multipath environments can be decomposed into three major categories: frequency-domain spectrum fitting [99] [100], covariance methods [101] [102] such as Hankel-Total-Least-Squares (HTLS) [103], and pulse detection and subtraction in the time domain [104] [105] [106]. The peak detection and subtraction method has also been used in frequency domain analysis to extract spectral content [107]. The computational complexity of the frequency-domain spectrum fitting involves *a priori* knowledge of the number of multipath components and a potentially expensive optimization step. The covariance-based methods employ matrix inversion and eigenvalue decomposition which become costly when more points are taken to form the covariance matrix.

Alternatively, simplified methods using only matched filters employed by Low et al. [44] cannot offer high accuracy in cases where multipath components experience delay of less than half of the pulse width. Consequently, the proposed algorithm, Iterative Peak Subtraction, follows the track of the “search and subtract” paradigm, with additional preprocessing (Algorithm 1, steps 3-4, following in Section 3.4.1) and a peak selection evaluation step (Algorithm 1, Step 7). We

evaluated the additions against two popular basic estimates for time-of-arrival: Received Signal Strength (RSS) and first detected peak. In Section 3.4.1, all three algorithms were tested using experimental data, and 1-D ranging errors are presented for each.

3.4.1 Iterative Peak Subtraction Algorithm

This algorithm was developed to subtract multipath peaks from the received multipath signal using a clean template to obtain an accurate estimation of the main peak. This algorithm was developed first as a post-processing step using Matlab. The aim was first high accuracy, secondarily speed. The critical portions of this algorithm are subsequently being implemented on a digital hardware device for additional speedup. Primitive FPGA implementations of this algorithm have already been constructed, and it should be feasible to run this algorithm in real-time on time extended signals. The algorithm begins using a normalized received signal envelope without any multipath components to form a clean template used for peak subtraction, as shown in Figure 30. The received pulse in Figure 30 contains added dispersion compared to Figure 22a due to the indoor channel as well as the receiver hardware.

The received UWB signal $h_0(t)$ with multipath components is represented by a summation of




Figure 30: Normalized received signal without multipath components (i.e. clean template).

pulses of varying amplitude and time delays as given by equation

$$h_0(t) = \sum_n a_n(t) s(t - \tau_n(t)) \quad (8)$$

where $a_n(t)$ represents the amplitude and $\tau_n(t)$ represents the time delay of each multipath pulse. Here only one cluster is considered in this simplified version of the S-V channel model [92] which has been used by many others [93] [94] [95] [96] [99] [100] [101] [102] [104] [105] [106]. For the purpose of localization geared towards high accuracy LOS conditions, this assumption remains valid. For conditions where non-line-of-sight (NLOS) is unavoidable, slight modifications can be made to the selection function presented in the following algorithm (which models the intra-cluster delay time of the first cluster) for increased performance. A treatment of associated NLOS errors is given by Denis et al. [100] suggesting that first detected peak is optimal. The phases of received rays within this signal are considered to be matched to sample points constituting the minimum resolvable delay bins. To accurately measure the main peak location, the effects of the multipath signal must be detected and removed, where in general $a_n(t)$ and $\tau_n(t)$ are not known *a priori*. The actual received multipath signal used as illustration is shown in Figure 31.




Figure 31: Actual received signal with severe multipath.

Algorithm 1: Iterative Peak Subtraction

1. The received signal is passed through an averaging filter to remove noise and is denoted by

$$y(t) = \text{ave}[h_i(t)] \quad (9)$$

where $h_i(t)$ represents a windowed moving average of the received signal after the i^{th} iteration (discussed in step 6), initially $h_0(t)$. A zero-phase filter $\text{ave}[\cdot]$ is needed to insure no shifting of the peak occurs.

2. The averaged signal is then subtracted from the initial received signal to get a measure of the channel noise as shown

$$n(t) = h_0(t) - y(t) \quad (10)$$

The noise is assumed to be Gaussian and the standard deviation is calculated and denoted by σ_n .

3. A large window (~ 200 ps) parabolic filter is used to detect large regions of positive and negative concavity and is denoted by

$$w(t) = b[y(t)] \quad (11)$$

where $b[\cdot]$ represents application of the parabolic filter to the averaged signal $y(t)$.

4. The concavity information held within $w(t)$ is used to localize regions of negative concavity within the averaged signal $y(t)$. The lower and upper time limits of the first region of negative concavity (which represents the first temporal peak) are obtained from $w(t)$ as shown

$$(l, u) = g[w(t)] \quad (12)$$

where $g[\cdot]$ is the function used to obtain the lower and upper limits l and u .

5. The peak in the region (l, u) is then located. If it has an amplitude greater than $3\sigma_n$, it is recorded

$$(p_t^i, p_a^i) = \max[y(l, u)] \quad (13)$$

where p_t^i is the time position and p_a^i is the amplitude of the i^{th} peak. If the amplitude is smaller than $3\sigma_n$, steps 4 and 5 are repeated on subsequent regions of negative concavity until a peak with amplitude larger than $3\sigma_n$ (6 dB above the noise floor) is found. If no relevant peaks are located, the algorithm terminates.

6. The first detected peak $p_a^0 s(t - p_t^0)$ is subtracted from the original received signal $h_0(t)$ and is denoted by $h_1(t)$. Each subsequent detected peak is subtracted in an iterative process as shown

$$h_{i+1}(t) = h_i(t) - p_a^i s(t - p_t^i) \quad (14)$$

This process is repeated (steps 1, 3-6) until all relevant peaks ($> 3\sigma_n$) have been detected and removed.

7. After all relevant peaks have been detected and subtracted, the main peak position P_t and amplitude P_a are detected by locating the peak with maximum amplitude after application of a weighting function (decaying exponential) as shown

$$(P_t, P_a) = \max \left[p_a^i \exp \left(\frac{-(p_t^i - P_t^0)}{\tau} \right) \right], 0 \leq i \leq M \quad (15)$$

where τ is a time constant, set to $\tau = 651$ ps, and M is the total number of relevant detected peaks.

The time constant τ is set to the intra-cluster decay time constant for the first cluster (%) used in the 802.15.4a LOS industrial channel model (CM 8) [98]. Figure 32 shows the original received multipath signal (also shown in Figure 31) with the detected peak positions and amplitudes marked by vertical bars. The decaying exponential is also overlayed to illustrate the behavior of the weighting function.




Figure 32: Original received multipath signal overlaid with selection function and detected peaks.

8. All peaks that occur after the main peak P_t are assumed to be multipath peaks. These peaks are located using

$$(MP_t^j, MP_a^j) = z(p_t^i, p_a^i, P_t), 0 \leq j \leq L \quad (16)$$

where $z(\cdot)$ returns the multipath peak positions MP_t^j and amplitudes MP_a^j for the L multipath peaks, respectively.

9. The new received signal is obtained by subtracting the detected multipath peaks from the original received multipath signal $h_0(t)$ as seen

$$r(t) = h_0(t) - \sum_{j=0}^L MP_a^j s(t - MP_t^j) \quad (17)$$

where $r(t)$ is the estimated received signal without any multipath components.

10. The last step is cross correlation of the estimated received signal $r(t)$ with the multipath-free received signal $s(t)$ as indicated

$$FP_t = \max[r(t) \otimes s(t)] \quad (18)$$

where FP_t is the temporal position corresponding to the maximum cross correlation value of the signals. Before performing cross correlation, all negative values in the estimated signal $r(t)$ are set to zero. FP_t represents the final main peak position of the received signal with all multipath components subtracted out.

It should be noted that the signals used in the cross correlation of step 10 have not been placed through any type of averaging filter. The large scale in which cross correlation is performed inherently filters out the uncorrelated Gaussian noise contained in $r(t)$ and $s(t)$. Also, as seen in Figure 32, there is a tendency for a small artifact found in front of the main peak in multipath conditions; this artifact is not subtracted. Figure 33 shows the estimated received signal $r(t)$ along with the original received multipath signal $h_o(t)$. All of the multipath peaks present in the original signal have been subtracted out.

3.4.2 Experimental Results

Testing of the three algorithms was performed against actual signals acquired in various multipath conditions. These experiments were performed by holding the transmitter and receiver in fixed positions, while progressively placing metal objects near the transmitter, as shown in Figure 34, with the purpose of evaluating the performance of the algorithms in worsening multipath conditions.

Initially no metal objects were placed near the transmitter and a signal was measured with no multipath interference. Each algorithm was allowed to determine the peak for this interference-free signal, the position for each algorithm was recorded. As additional pieces of metal were




Figure 33: Received multipath signal $h_o(t)$ and estimated main peak signal $r(t)$.




Figure 34: Omni-directional transmitting antenna placed in a severe multipath environment surrounded by numerous closely spaced metal objects.

Table 15: 1-D Multipath Experiment

Method	Ave. Bias (mm)	RMSE (mm)
RSS	10.4	160
First Peak (6dB threshold)	-3.26	58.6
First Peak (15dB threshold)	3.00	4.10
Iterative Peak Subtraction	0.0938	2.77

placed near the transmitter, causing multipath interference, each algorithm again determined a new peak position. This new peak position was compared with the original peak position for the particular algorithm, which allows for direct comparison of the error metrics across the various algorithms. In these experiments, the dominant ray originating from a LOS path was always preserved. A total of two multipath-free signals and six multipath signals were recorded for a fixed position. As shown in Figure 31, one of the received severe multipath signals was obtained from this experiment. Table 15 records the evaluation of the spatial Root Mean-Squared Error (RMSE) and average bias error given by each algorithm over the course of this experiment.

The iterative peak subtraction algorithm performs better than the other basic algorithms, but at higher processing cost. Only pulses that lag by less than one pulse width influence the peak position of the first pulse. The algorithm can be simplified to only search for pulses that are less

than one pulse width from the leading pulse. This will reduce the computational cost associated with finding all multipath peaks. In the situations encountered in this experiment the multipath peaks are spaced at intervals of roughly equal or larger size to the pulse width (Figure 31). This allows highly accurate results to be obtained from the iterative peak subtraction algorithm while the RSS algorithm results in relatively larger error, due to a small number of gross mispredictions. The large positive bias in the RSS algorithm is due to several instances where the maximum signal is significantly later than the original peak. The iterative peak subtraction algorithm performs marginally better than the first peak algorithm, when the first peak algorithm assumes a 15dB signal-to-noise ratio (SNR) to calculate an appropriate threshold. The performance of the first peak algorithm degraded significantly when the threshold value was set under a 6dB SNR assumption.

The effect of coherent interference (e.g. 802.11a) on receiver-side peak detection deserves consideration. Coherent narrowband and wideband interference can affect TOA or TDOA UWB systems, as outlined in [108]. Intelligent thresholding when detecting the peak of the incoming signal can significantly mitigate the effects of narrowband and wideband interferers in UWB positioning systems [108]. 802.11a signals (5.0 – 5.83 GHz) can also be filtered out either by the antenna acting as a bandpass filter or by the addition of a bandreject filter since this frequency band only includes a small portion of one of the sidebands of the UWB signal. Through filtering of the coherent interference or through intelligent thresholding of the incoming signal, or through a combination of these two techniques, we anticipate mitigating the effects of this type of interference. We do not expect this effect to be magnified by the extra time needed to do the sub-sampling since the same interference will be present in each frame needed to reconstruct the time

extended signal, which implies that, even if the narrowband interference is included in the signal, the signal itself can be assumed to be constant in shape over the duration of cycles needed to create the time extended reconstruction. Lastly, if a leading edge detection scheme is used at the base stations instead of peak subtraction, this will make the system even more robust to any time shifts which could be introduced via 802.11a interference.

3.5. Antenna Phase Center Variation

Accounting for antenna phase center variation at the transmitters and receivers is critical for performance in high accuracy localization systems. Ideally all frequencies contained in the pulse are radiated from the same point of the UWB antenna and thus would have a fixed phase center [109]. In this case, all frequencies travel the same distance within the same time, and the pulse can be received undistorted.

In practice, however, the phase center varies with both frequency and direction. For localization systems that require high accuracy, this can result in significant localization errors. For example, to compensate for phase center variation in GPS antennas, automated high precision robots are used in a calibration procedure to move a GPS antenna into 6000 – 8000 distinct orientations [110]. In the case of the transmitting antenna, which is an UWB monopole, phase center variation is less than 1 mm and is considered negligible (both across the frequency band from 6 – 10 GHz and as the azimuth angle is varied). Phase center variation along the broadside direction was simulated to estimate the axial position of the Vivaldi phase center. Figure 35 shows the simulated phase center variation over the desired frequency band at broadside using CST software, with the original point set at the input of the Vivaldi antenna, as shown in Figure 36.




Figure 35: Simulation of broadside Vivaldi phase center location versus frequency.




Figure 36: Experimental setup of Vivaldi antennas in an anechoic chamber used to measure Vivaldi antenna directivity-dependent phase center variation.

The variation in phase center over frequency can be further reduced by placing a dielectric rod over the end-fire portion of the Vivaldi antenna [111]. The average phase center position across the frequency band of 6 – 10 GHz is obtained at 39.5 mm which is later used as the “apparent phase center” in directivity-dependent phase center measurements.

Since the UWB pulse contains broadband frequency information, a more accurate method for defining the phase center variation of the Vivaldi antenna is to employ time domain techniques. As shown in Figure 36, an experiment was setup in an anechoic chamber to quantify how the phase center is affected by the directivity based on time domain measurement. Both transmitting and receiving Vivaldi antennas were put face to face and separated by a distance of 1.5 m. The

receiving antenna was rotated around the calculated “apparent phase center” (at 39.5 mm, shown in Figure 36) from -45° to 45° at 5° per step. The apparent phase center was tracked on the receiving Vivaldi antenna as it was rotated from -45° to 45° with an optically tracked probe. These reference points from the optical system were used to calculate the actual center of rotation during the experiment. This allowed changes in the actual phase center as the receiving antenna was rotated to be separated from physical movement of the apparent phase center, shown in Figure 36.

Figure 37 shows the measured phase center displacement for both the E and H cuts. As shown in Figure 37, the measured phase center variation versus rotating angle indicates a small phase center variation of less than 2 mm within $\pm 20^{\circ}$ while the variation degrades dramatically with an angle greater than 30° . TDOA-related error due to this phase center error can be minimized by varying the number of base stations and their orientation in a standard size room. Additionally, other techniques to calibrate out the phase center error in a practical system by either assuming or measuring orientation for each base station is required to achieve sub-mm accuracy.




Figure 37: Measured Vivaldi phase center error versus angle for E-cut and H-cut.

3.6. Localization Experiments

3.6.1 Synchronized Experiments

In the early experimental trials, 1-D, 2-D, and 3-D localization experiments were performed to test the system. In order to test the feasibility of a high accuracy short range system, the equivalent time Tektronix TDS8200 oscilloscope was used instead of the receiver architecture outlined in Figure 21. Also, for these experiments the tag clock PRF_1 was synchronized with the base station clock PRF_2 . This allows errors due to the system architecture and TDOA algorithm to be isolated by accounting for clock synchronization, receiver-side sampling, etc. These experiments provide a benchmark on achievable system accuracy. When testing a 3-D localization system for mm-range accuracy, a highly accurate reference positioning system is required for calibration and validation. For this purpose the Optotrak 3020 [112] is used to obtain accurate reference data. It is a 3-D IR tracking system that provides 3-D positioning data with accuracy of less than 0.3 mm, which is needed to validate the developed UWB positioning system. The indoor environment for the 1-D, 2-D, and 3-D experiments is in a laboratory with dense multipath effects including reflection from side walls, floor, furniture, ceiling, test equipment, and human bodies. UWB signals that pass through or refract off the human body (e.g. high permittivity materials) will exhibit a large time delay [113]. This introduces significant error into the TDOA system. Consequently, in these experiments only LOS cases were studied. The monopole transmitting antenna was constantly within $\pm 20^\circ$ of broadside direction from the Vivaldi antennas to minimize the phase center variation at the base stations.

1-D and 2-D measurements similar to [76] were conducted by moving the UWB monopole antenna along a precision optical rail. Results for the 1-D and 2-D experiments are included in

Table 17. Figure 38(a) shows the experimental setup with 6 base stations put at fixed locations around the transmitting tag under test. Both the tag and base stations were put on the supporting rails. The tag was then moved along the rail over 8 discrete positions.

The placement of the base stations is critical for the accuracy of a TDOA localization system. Ideally, the base stations would be spherically positioned around the volume of interest. If only four base stations are used, placing them at the nodes of a tetrahedron gives optimal results, although this may not be possible in actual indoor environments. Therefore, the number of base stations used for TDOA is varied, and the effect this has on overall system accuracy is analyzed. Figure 38(b)-(d) outline the other three configurations considered where only 4 or 5 base stations were used in TDOA localization. It should be noted that the data from the experiment outlined in Figure 38(a) is used for all configurations by leaving out specific base stations to realize the other three configurations.




Figure 38: 3-D synchronized localization experiments: (a) 6 base stations. (b) 5 base stations. (c) 4 base stations with low position dilution of precision (PDOP). (d) 4 base stations with poor coverage in x-direction resulting in high PDOP.

Table 17 summarizes the mean, standard deviation and worst case error of overall distance for the 1-D, 2-D, and 3-D experiments. The 3-D experiment includes two 4 base station scenarios, one with 5 base stations, and a 6 base station configuration, as shown in Figure 38. As shown in Table 17, using all 6 base stations yields the highest accuracy at 2.45 mm. The system accuracy gets progressively better as the number of base stations is increased from 4 to 6. Also, the largest error is seen using the 4 base station configuration shown in Figure 38(d). All of the base stations in this configuration reside on the +x side of the tag. As shown in Figure 34, this results in large error in the x position, which can be attributed to poor base station placement relative to the tag, also known as position dilution of precision (PDOP). PDOP is discussed in detail in the following section.

3.6.2 Position Dilution of Precision

PDOP can be computed upon convergence of the TDOA algorithm as in [114] and is dependent solely on the base station geometry relative to the tag position. The TDOA algorithm involves linearizing the relative range measurements of each of the base stations about a position estimate




Figure 39: Error in x, y, and z illustrating high PDOP for base station distribution shown in Figure 38 (d). This shows the reduction in accuracy that results from poor base station spatial arrangement.

using matrix H [115]

$$H D_x = R \quad (19)$$

where R represents a matrix with the range difference elements, and D_x is the position update vector. The least squares solution of (19) is shown in (20)

$$D_x = (H^T H)^{-1} H^T R \quad (20)$$

where D_x can be used to update the current position estimate P_i

$$P_{i+1} = P_i + D_x \quad (21)$$

The dilution of precision parameters can be calculated using (22)

$$\begin{bmatrix} \sigma_x^2 & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_y^2 & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_z^2 \end{bmatrix} = (H^T H)^{-1} \quad (22)$$

Combining the diagonal elements the overall PDOP is found

$$PDOP = \sqrt{\sigma_x^2 + \sigma_y^2 + \sigma_z^2} \quad (23)$$

Alternatively, the PDOP can be explicitly defined for each coordinate direction as

$$PDOP_x = |\sigma_x| \quad (24)$$

The 3-D error can be estimated by combining the PDOP of the geometric configuration with the 1-D uncertainty of the system [62] as shown in (25)

$$\sigma_{3D} = PDOP \sigma_{1D} \quad (25)$$

For each of the base station configurations the mean PDOP is reported across the eight tag positions for each of the coordinate axes as well as the overall PDOP. The results are reported in Table 16.

Table 16: PDOP Summary – Synchronized Localization Experiments

	Mean PDOP _x	Mean PDOP _y	Mean PDOP _z	Mean Overall PDOP
6 BS	0.87	0.44	1.61	1.88
5 BS	0.90	0.59	1.70	2.02
4 BS - Fig. 19(c)	1.40	0.60	1.89	2.43
4 BS - Fig. 19(d)	20.2	1.71	3.33	20.6

Table 17: Error Summary – Synchronized Localization Experiments

	Mean Error (mm)	Std. Dev. Error (mm)	Worst Case (mm)	Mean PDOP
1-D	1.49	0.69	3.0	-
2-D	2.61	0.69	3.6	-
6 BS	2.45	0.93	3.3	1.88
5 BS	3.13	1.20	4.2	2.02
3-D 4 BS - Fig. 19(c)	3.62	1.53	5.6	2.43
4 BS - Fig. 19(d)	43.3	32.4	96.1	20.6

3.6.3 Unsynchronized 1-D Experiment

A 1-D experiment with unsynchronized LOs and PRF clock sources was also carried out to validate the theory in Section 3.3 and to test the robustness of the system. The experimental setup is shown in Figure 40. Two base stations are needed for 1-D measurements. The low bandwidth output of the sub-sampler is fed through an ADC. Next, the signal is fed to the FPGA which uses RSS rather than the peak subtraction algorithm to locate peak position. As seen in Figure 41, mm-range accuracy was consistently achieved for 1-D unsynchronized measurements at 8 separate locations with a 5 cm distance between two successive measurements, although system jitter does cause noticeable short term variation in the error at each static point of roughly ± 10 mm. This short term variation was mitigated by averaging 16 pulses at each static point. As shown in Table 17 and Table 18, the mean error in measuring 1-D static points increases from 1.49 mm to 3.07 mm. The increase in error of 1.58 mm is comparable to the measured error of




Figure 40: Experimental setup for 1-D unsynchronized positioning measurement.




Figure 41: Measured error of 1-D unsynchronized experiment.

Table 18: Error Summary – Unsynchronized 1-D Experiment

	Mean Error (mm)	Std. Dev. Error (mm)	Worst Case (mm)	Mean PDOP
1-D	3.07	2.39	6.4	-

1.05 mm due to PRF clock jitter shown in Figure 29. The error also includes time scaling effects, which were not removed in post processing of the data.

3.6.4 Error Discussion

The data in Table 17 and Table 18 can be understood by considering the following points.

The PDOP values for each coordinate direction can be understood by examining the 6 base station experiment as shown in Figure 38(a). The spatial spread of the base stations along the y -axis is the largest (1.86 m), thus the $PDOP_y$ value listed in Table 16 is the lowest among the coordinate axes. This is a common trait to the experimental set up in Figure 38(b-c), and consequently the PDOP values are consistent across these experiments. In these experiments the $PDOP_z$ values are consistently the highest, due to the maximum range of base station positions along the z -axis being relatively the smallest (0.57 m). Figure 38(d), which represents a poor geometric configuration, has $PDOP_x$ as the worst due to the low spatial diversity of base stations along the x -axis.

The phase center of the Vivaldi receiving antenna varies with directivity, which can have large effects on system accuracy especially when the tag is in locations on the border or outside of the target volume. We insured in this experiment that all tag positions remained within $\pm 20^\circ$ of the broadside direction of each single element Vivaldi in order to minimize phase center effects. The error originating from this effect is estimated to be less than 1 mm, although if the angle from broadside increases beyond $\pm 20^\circ$, the phase center error increases dramatically.

Multipath interference from extremely close metal (e.g. metal bar supporting transmitting tag) causes pulse peak shifting. The developed algorithm can handle dense multipath situations but still has substantial uncertainty of around 3 mm under severe multipath conditions, as shown in

Table 15. These localization experiments were done with strong LOS signals and only minimal amounts of multipath interference.

As discussed in Section 3.3, LO and PRF clock jitter and LO frequency offset from the tag relative to the base stations can affect the sub-sampling process and introduce error into the TDOA measurements. For the synchronized experiments, the LOs and PRF clock sources for both the tag and base stations were synchronized to evaluate overall system accuracy in the absence of clock jitter and frequency offset. As shown in Table 17, mean error of 2.45 mm is possible in localizing 3-D points with 6 base stations.

A 1-D experiment was conducted where the LOs and PRF clock sources for the tag and base stations were unsynchronized. With 16x averaging, the mean error was 3.07 mm, roughly twice that observed in the 1-D synchronized case. Using the PDOP from the 6 base station experiment and (25), the expected 3-D mean error for the unsynchronized system is 5.77 mm. This error could be reduced by applying the peak subtraction algorithm to the received signals. More investigation of how the 3-D real-time, unsynchronized system performs in terms of overall error will be investigated in future work.

The overall system error can be significantly reduced by increasing the number of base stations. This is highlighted by considering the different sources of error outlined previously. As shown in Table 17, the PDOP decreases with a higher number of base stations. Also, system error due to multipath interference is reduced with increased number of base stations since the received signals at each base station will traverse through a different UWB channel realization. Next, by calibrating the system to determine base station orientations, phase center effects can be drastically reduced. This calibration step is more accurate with increased number of base

stations. Also, the error due to LO frequency offset, which affects the I/Q down conversion, is reduced with increased number of base stations since LO_2 is mixed with the received pulse at different phase offsets for each base station depending on the distance between each base station and the tag. Additional base stations result in additional measurements at random phases of LO_2 which reduces this effect. Finally, the 1-D ranging error of 1.49 mm represents inherent 1-D system error and can be considered an unbiased Gaussian random variable. Filtering the 1-D ranging values or increasing the number of base stations can minimize the impact of this 1-D ranging error on overall 3-D system accuracy.

3.7. Conclusion

A short range UWB indoor localization system with mm-range accuracy has many promising applications. Conversely, many challenging problems must be overcome in order to realize such a system in practice. We are proposing a system architecture similar to available commercial systems but equipped with more advanced receiver-side hardware for sub-sampling the incoming pulse train and detecting the main LOS peak. We have addressed the main challenges being faced in building this system, including sampling limitations, multipath interference, phase center error, and timing errors due to clock jitter, LO offsets, temperature effects, etc. Sub-sampling techniques, sophisticated receiver-side leading edge detection, increased number of base stations, phase center calibration, and high fidelity PRF crystals combined with a TDOA approach have been proposed as solutions to these problems while the experimental results show the feasibility of achieving mm-range accuracy with an UWB indoor positioning system in highly reflective indoor environments. The sub-sampling mixer is a viable option for sampling the incoming pulse train with a high sample rate. The effects of clock jitter can be mitigated by

using PRF clocks with low phase noise and temperature compensation, although the total effect of PRF clock jitter on the sub-sampling process can cause substantial error on the order of 1 mm. Time scaling originating from PRF clock stability has been shown through simulation to not impact final system accuracy by using a modified version of the TDOA algorithm. The designed peak subtraction algorithm is robust to multipath interference. System performance in dense indoor environments can be further increased by greater number of base stations and optimal base station placement. Base station phase center error has been shown to be significant. In the finalized system, calibration of base station orientation will allow this error to be removed. We have shown that, by using a system level approach, various errors encountered in designing a mm-accuracy UWB positioning system can be eliminated or reduced.

4. CHANNEL MODELING FOR UWB POSITIONING

4.1. Background

In the field of UWB channel modeling, two approaches have dominated the literature over the recent years. The first approach seeks to classify the environments in which the channel exists (e.g. indoor office space, industrial, outdoor, etc.) and to what degree the environment affects or determines a number of statistical parameters on which the model depends. Essentially, by categorizing the various possible environments, the model quantifies on a macro scale the density of metal reflectors and the presence or absence of other dielectrics. This type of channel modeling further describes whether the propagating signals have a LOS component or are received purely as either reflected or refracted signals in a NLOS condition.

To develop and subsequently quantify the accuracy and efficacy of such models, researchers make time-domain or often frequency-domain measurements in the various LOS or NLOS conditions across a variety of environments [116]. The measurements seek to approximate the widely-varying situations in which a UWB communication link or a positioning system might be placed and ultimately attempt to statistically describe the channel impulse response. The drawbacks of this paradigm center on the inability of a finite set of measurements to be entirely representative of an entire class of environments, even supposing that the set of environments is complete. For example, the statistical model may fail when transmissions that originate in an office space terminate in an urban canyon. A second criticism for this type of channel modeling involves the resolution of the measurements and their potential suitability for other purposes. If a battery of tests is performed, for example, with low-resolution positioning in mind, to what

degree can information contained in the model be utilized for the development of a high-speed wireless communication channel? Despite these limitations, the most broadly studied aspects of the UWB channel have been modeled using this statistical description technique.

Another promising, yet under-developed, channel modeling method is that of a deterministic description of the channel using computational or analytical ray-tracing. In this channel modeling technique, the geometry of the environment is explicitly defined, including but not limited to, the position and orientation of the transmitter and receiver, the location and orientation of metal surfaces, and the relative permittivity, size, and location of dielectric materials. Thus, this model only considers a particular volume of interest, and as such, signals that would propagate from reflections and electromagnetic interactions outside of this region are intrinsically discarded. Models described using ray-tracing techniques cannot possibly account for any and all possible transmission paths within the volume of interest, and thus are limited to the first several rays, which are expected to be of the highest power. Alternatively, analytical simulation software can search for all possible rays with only a few reflections or refractions. The advantages of the ray tracing methodology are clear when the geometry is well known or is unchanging, such as transmission paths from base stations in a large city [117]. However, in time-evolving channels or cases when the transmission geometry cannot be known *a priori*, the ray tracing channel modeling method is less successful in accurately predicting the channel realization.

4.2. Statistical Channel Modeling

A large amount of study and devoted work have been performed in regard to channel modeling in the field of indoor UWB communications [116]. The foundations of this work lie with the

exploration of UWB channel models as a description of the actual propagation expected in indoor environments, fitting under the statistical parameter-based technique. Beginning with the

$$h(t) = \sum_l \sum_k \beta_{kl} e^{j\theta_{kl}} \delta(t - T_l - \tau_{kl}) \quad (26)$$

Saleh-Valenzuela (S-V) indoor propagation model [92], many adjustments and refinements have been added. The S-V model defines the UWB channel impulse response as:

where $\beta_{kl} e^{j\theta_{kl}}$ is the complex amplitude, T_l is the time of the l^{th} cluster, and τ_{kl} is the arrival of the k^{th} ray within the l^{th} cluster (Figure 42).

By modeling the inter-cluster arrival times and inter-ray arrival times as Poisson processes, the S-V model can adequately describe mean excess delay and RMS delay spread metrics, which have important consequences in terms of communication applications. In positioning applications, however, the received signals coming after the LOS, or direct ray, carry relatively inaccurate positioning information. In the absence of a direct ray, it has been shown [100] that selecting the first ray out of the first detected NLOS cluster is considered optimal. Due to measurements conducted in realistic environments, the S-V model was extended with additional statistical parameters [98].




Figure 42: Power delay profile of multi cluster, S-V UWB channel model

Writrisal and Pausini derived the first and second order statistics of the autocorrelation and cross-correlation functions of UWB channel modeling [118]. Cheng et al. modeled the underwater acoustic channel as a modified version of the S-V UWB channel model, which uses the TOA algorithm for final positioning [119]. Gaertner and Nuallain showed the relatively smaller degree of small scale signal power fading of wideband signals versus narrowband signals across a distance of 5 m [120]. Research has also attempted to unify deterministic channel models, such as ray tracing and purely statistical models, by computing joint TOA/AOA statistics for particular scatterer geometries [121].

4.3. Ray Tracing Channel Model

Ray tracing methodologies take a different approach than typical statistical modeling methods. For example, ray tracing seeks to model the complex channel impulse response against specific transmitter-receiver geometries by deterministically finding the strongest set of rays. Each ray originates from the transmitter and terminates at the receiver, undergoing multiple reflections, refractions, and/or attenuations based on the various mediums encountered. While the goals of a statistical model involve more generalized metrics as the previously mentioned mean excess delay and RMS delay spread, ray tracing involves close inspection of received signals and purports to explain such effects as fast fading, slow fading, and time-dependent fading, as well as provide an estimation of the frequency dependence of the channel. As shown in Figure 43, rays originating at the transmitter each travel different path lengths based on the whether the path is LOS or NLOS, noting that the phase response of the reflection or refraction can change the received signal phase angle. In the ray tracing paradigm, the rays are collectively summed at the




Figure 43: Ray Tracing diagram showing three rays: one line-of-sight ray, and two non-line-of-sight rays

receiver in the complex domain to measure constructive or destructive interference present in the channel.

As an alternative to the S-V model, a simple simulation was developed to test the presence of multipath interferers when the UWB transmitter and receiver were placed at varying distances away from a highly reflective metal wall. A total of 34 measured locations were used in this simulation and corresponding experiment. The measurements were classified into four categories based on the proximity of the transmitter and receiver to a metal wall: Near Wall/Near Wall, Near Wall/Far Wall, Far Wall/Near Wall, and Far Wall/Far Wall. The simulation portion involved a ray tracing model of 6 primary rays: LOS, metal wall reflection, ceiling reflection (concrete, $\epsilon_r = 15$), a wall-ceiling reflection, ground reflection (concrete, $\epsilon_r = 15$), and wall-ground reflection. Using parameters of the system (300 ps pulse, with a carrier of 8 GHz), the E-field of the transmitting and receiving plane was oriented vertically. One assumption used in this experiment was that all rays experience $1/R^2$ path loss, which is a relatively optimistic path

loss value, considering the constrained indoor environment in which the experiment took place. Additionally, the reflection coefficient was not considered to be frequency dependent.

The purpose of this experiment and simulation was three-fold. First, the purpose was to examine the pulse distortion and attenuation under several realistic conditions. Secondly, the time delay of the multipath components can be measured directly. Finally, this ray tracing model serves as the benchmark generative model for comparison of peak-detection/leading edge algorithms discussed in the remainder of this chapter.

Categorically, the configurations where the transmitter and receiver were both close to the metal wall (Figure 44), suffered the greatest distortion and attenuation when compared to the other studied configurations. Another configuration involved when the transmitter was near the wall and the receiver was placed away from the wall (Figure 45), showing a slight delay (~ 35 ps). Similarly, when the transmitter was placed farther away from the wall and the receiver was placed near the wall, results were very similar due to symmetry (Figure 46). When both the transmitter and receiver were placed far away from the metal wall, the power received by the dominant ray was the highest out of the four categories of experiments (Figure 47).




Figure 44: a) Simulated ray tracing results for transmitter and receiver 1 cm away from metal wall, the signal without the presence of the metal wall (green) is severely attenuated by destructive interference in the received signal with the wall present (red) b) Experimental results for the same configuration, noting the similarities of destructive interference in the received signal




Figure 45: a) Simulated ray tracing results for transmitter 1 cm and receiver 104.5 cm away from metal wall, respectively, the signal without the presence of the metal wall (green) is slightly shifted by destructive interference in the received signal with the wall present (red) b) Experimental results for the same configuration, noting the similarities of the strong dominant ray




Figure 46: a) Simulated ray tracing results for transmitter 1.00 m and receiver 1.0 cm away from metal wall, respectively, the signal without the presence of the metal wall (green) is slightly shifted by destructive interference in the received signal with the wall present (red) b) Experimental results for the same configuration, noting the similarities of the strong dominant ray




Figure 47: a) Simulated ray tracing results for transmitter 1.02 m and receiver 1.00 mm away from metal wall, respectively, the signal without the presence of the metal wall (green) is no longer shifted by destructive interference in the received signal b) Experimental results for the same configuration, noting the similarities of the strong dominant ray, and absence of significant multipath signals

A second type of analysis that resulted from this 6-ray model was that of the frequency dependence of the channel, also known as frequency fading [117]. In the Near Wall/Near Wall case, shown in Figure 48, the frequency fading occupies a 10 dB range across the frequency bandwidth of the pulse (5-11GHz), yielding an only slightly distorted pulse in the time domain. While in the Near Wall/Far Wall case (and by symmetry, the Far Wall/Near Wall case), shown in Figure 49, the frequency fading has a range of nearly 30 dB, leading to significant distortion to the received pulse. Also, in this case note that the highest relative power is 17 dB below that of the Near Wall/Near Wall case, indicating significant power loss due to destructive interference between the LOS and the secondary wave. Lastly, in the Far Wall/Far Wall case, depicted in Figure 50, the frequency fading trend is similar to the Near Wall/Near Wall case, except with roughly a 4 dB lower overall power. Consequently, in this case the pulse undergoes very slight distortion (Figure 50b).

In addition to the developed 6-ray model, a commercial software package, InSiteTM (Remcom Inc.) [122], was also used to compare the previous simulation and experimental results. The




Figure 48: Near Wall/Near Wall situation a) Frequency dependence of the channel, noting 10 dB range b) Simulated received pulse (red) showing only slight distortion referencing the transmitted pulse (blue)




Figure 49: Near Wall/Far Wall and Far Wall/Near Wall situations a) Frequency dependence of the channel, noting 30 dB range b) Simulated received pulse (red) showing stronger distortion referencing the transmitted pulse (blue)




Figure 50: Far Wall/Far Wall situation a) Frequency dependence of the channel, noting nearly 15 dB range b) Simulated received pulse (red) showing weak distortion referencing the transmitted pulse (blue)

geometry used in the InSite software simulation included the entire lab as shown in Figure 51. Consequently, the received rays found by the InSite simulator include rays not explicitly simulated in the 6-ray model previously mentioned.

The results of these simulations compare favorably with the prior 6-ray simulation and experimental results, but with a few important differences. The InSite results for the Near Wall/Near Wall case (Figure 52) did not compare well with either the experimental data, or the 6-ray model. It would seem that the InSite results did not take destructive interference effects into account. Both the Near Wall/Far Wall (Figure 53) and the Far Wall/Near Wall (Figure 54) InSite results are consistent with the 6-ray model and the experimental results, with the characteristic pulse delay shift, and the slight attenuation. In the simpler case of Far Wall/Far Wall, the InSite results shown Figure 55 agree well with the 6-ray model and the experimental results.

Another way to consider the problem of interacting rays is to view the ray tracing problem in terms of more generalized geometries. A particularly simple geometry to consider is that of the Fresnel ellipse [117]. The Fresnel ellipse, or ellipsoid when considering three dimensions, has foci at the transmitter and receiver locations, respectively. The waist or widest point of the ellipsoid is centered at the midpoint of the line segment with endpoints at the transmitter and receiver given by

$$w_F = \sqrt{\lambda r} \quad (27)$$

where r is the distance between the transmitter and receiver and λ is the wavelength of interest. Scatterers shown to be interior to the Fresnel ellipsoid will have a significantly stronger effect on the channel impulse response than those exterior. The wavelength corresponding to the center




Figure 51: InSite software simulation of ray tracing experiment. Transmitter positions (red squares) are varied in distance from the metal wall (yellow, bottom), while receiver positions are moved in two parallel tracks (green squares) relative to the transmitter.




Figure 52: Insite software simulation corresponding to transmitter and receiver each 1 cm away from the metal wall, corresponding to the experimental and simulation results of Figure 44




Figure 53: Insite software simulation corresponding to transmitter 1.0 cm and receiver 1.0 m away from the metal wall, corresponding to the experimental and simulation results of Figure 45. Inset is a magnified view of the dominant peak.




Figure 54: Insite software simulation corresponding to transmitter 1.0 m and receiver 1.0 cm away from the metal wall, corresponding to the experimental and simulation results of Figure 46. Inset is a magnified view of the dominant peak.




Figure 55: Insite software simulation corresponding to transmitter 1.0 m and receiver 1.0 m away from the metal wall, corresponding to the experimental and simulation results of Figure 47

frequency of the 8 GHz system is nominally 3.74 cm. Consequently, a 1.0 meter transmitter-receiver distance will have a Fresnel ellipsoid with a maximum waist diameter of 19.3 cm. As seen in the Near Wall/Near Wall and Near Wall/Far Wall cases highlighted in the simulations and experiments above, the system undergoes dramatic changes until all reflectors are outside of this zone. Also, to a lesser degree, a similar ellipsoid can be drawn by substituting the wavelength for the spatial width of the 300 ps pulse, specifically 9.0 cm. The interior of this ellipsoid will describe spatial points where reflected pulses will overlap with the LOS pulse (Table 19).

A final experiment was enacted to demonstrate antenna effects on received signals. As shown in the InSite geometry of Figure 51, multipath signals can impinge on the receiving antenna from a wide set of possible angles. If the correct direction of the signal is roughly known, then proper antenna selection can be used to remove signals received from unwanted directions. Using an omni-directional monopole antenna as both a transmitter and receiver, allows for the fullest capture of all multipath signals, while using an antenna with higher directivity can eliminate

multipath signals. As shown in Figure 56, the received signal using a Vivaldi antenna with higher directivity eliminates multipath signals that can be clearly discerned from the received signal when an omni-directional monopole is used as the receiving antenna. It is noted in this application, the transmitting antenna must always be omni-directional as it is not possible for the transmitting tag to communicate with only one base station using the TDOA approach. This is a consequence of needing to measure simultaneous time-difference-of-arrival values between base stations.

Table 19: Fresnel Ellipsoid Maximum Waist Diameters for several typical Tx-Rx distances, assuming a 300 ps pulse

Tx-Rx Distance	Maximum Waist Diameter
1 m	43.3 cm
2 m	60.6 cm
5 m	95.2 cm




Figure 56: Comparison of high directivity receiving antenna (blue, Vivaldi) with omni-directional receiving antenna (red, monopole), when an omni-directional monopole is used as transmitter in both cases

5. PULSE AND LEADING EDGE DETECTION

5.1. Background

As has been mentioned in Chapter 3, a large body of work has been presented in the literature discussing methodology and challenges in pulse and leading edge detection, often referred to as TOA estimation. R. Fontana states, “Among the most significant problems encountered in the measurement of the time-of-flight for a short pulse emission are the deleterious effects of multipath cancellation and reverberation [123].” The field of TOA estimation typically approaches the problem from a macroscopic viewpoint with less stringent accuracy constraints than those of this project (~ 1 mm). For example, Kamil examines TOA estimation techniques, such as the CLEAN and generalized maximum likelihood (GML) algorithms, and measures error in nanoseconds [124], which corresponds to a spatial error on the order of 30 cm. Also, Saeed et al. adds a synchronization scheme to TOA estimation, using OFDM and Direct Sequence (DS) UWB positioning and found that the lower bound of errors ranged from 3 cm to 12 cm depending on the UWB standard used [125]. Overall, TOA estimation has focused primarily on large-range, low-accuracy applications, such as tracking users via cellular communication networks or local area networks (LAN), where even meter-range errors are acceptable [126].

5.2. Sampling Requirements Simulation

In considering TOA estimation, there exists a direct relationship to sampling rate and the potential for temporal error. Without using any super-sampling techniques, the sampling error is dictated by the spatial width of the samples. The worst case occurs when the point to be detected occurs midway between two samples and is given by:

$$E_{max}^{samp} = 0.5 \frac{c}{F_s} = \max \left\{ \frac{pc}{2F_s} \right\} \quad (28)$$

where F_s is the sampling frequency in Hz and c is the speed of light, and p is a uniformly distributed random variable on $[0,1]$. This model supposes that in the case that the true peak or leading edge surpasses the midpoint between two samples it will always be resolved to the next sample. In the presence of uniformly distributed distances, the expectation of error due to this phenomenon is

$$E_{RMS}^{samp} = \sqrt{\frac{1}{12} \frac{c}{F_s}} = E \left[\frac{pc}{2F_s} \right] \quad (29)$$

where $E[\]$ represents the statistical expectation operator. A small study was enacted to test the system sampling requirements when considering this effect alone in conjunction with 4 or more base stations used for positioning. In this simulation, 4, 6, or 8 base stations were randomly distributed over a 3 m radius sphere, representing the various geometries possibly encountered in real environments. A simulated tag trajectory involved 81 positions in a helix pattern generated by Equations (79)-(80), between x-axis values of $\pm R/2$. A nominal sampling frequency was selected between 1 GS/s and 512 GS/s in powers of two. For each tag position-base station pair, the true ranging value was rounded to the nearest integer multiple of the spatial sample width. Finally, the TDOA algorithm was run on these degraded range values, and the maximum error and RMS error were averaged over 10 trials. Figure 57 shows the case of the 4-base station experiment, where it becomes apparent that 1 cm mean RMS accuracy is possible with a 128 GS/s sampling rate (or 2.34 mm spatial sample width). Figure 58 shows the case of the 6-base station experiment, where it becomes possible to achieve 1 cm mean RMS accuracy at about 70 GS/s sampling rate (or 4.3 mm spatial sample width).




Figure 57: TDOA system simulation on the effect of sampling error when N=4 base stations are used in the positioning calculation




Figure 58: TDOA system simulation on the effect of sampling error when N=6 base stations are used in the positioning calculation




Figure 59: TDOA system simulation on the effect of sampling error when N=8 base stations are used in the positioning calculation

Using 8 base stations, 1 cm mean RMS accuracy is possible at about 50 GS/s (6 mm spatial sample width) (Figure 59). Further, sub-mm mean RMS accuracy is achievable when a 512 GS/s sampling frequency is used, which corresponds to 0.58 mm spatial sample width.

Given the target accuracy of 1 mm, the nominal value of the sample spacing used for much of the finalized system testing was 0.6 mm. This corresponds to $\alpha = 10,000$ scaling factor on the sampling mixer discussed in Chapter 3 and a 50 MS/s ADC, which gives an equivalent resolution of a 500 GS/s sampling rate. This offers an attractive trade-off: setting $\alpha = 10,000$ will only cause $\pm 1.1\%$ unintentional scaling error (Figure 60), while spatial sampling error will only be responsible for a minimal portion of the total system error. When α is set at 100,000 the system accuracy is 11% with a 0.5 ppm clock stability figure.




Figure 60: Maximum percent scale factor error vs. scale factor (α) for two clock stability values, 0.5 ppm (blue) and 1.0 ppm (red)

5.3. Iterative Peak Subtraction

5.3.1 Iterative Peak Subtraction Algorithm Validation - Simulation

The described peak subtraction algorithm (Chapter 3) was compared against both the RSS and first peak-finding algorithms in extremely severe multipath conditions. The first comparison was done using simulated multipath data. These signals were constructed by placing the clean received pattern in a known temporal location and adding a time delayed version of the same pattern with various amplitudes and time delays. The secondary peak was delayed over a range of 0 ps to 440 ps and the secondary peak amplitude was varied from 1% to 125% (-40dB to 1.9dB) of the primary peak amplitude. A synthetic multipath signal representing the primary peak combined with a secondary peak at a lag time of 440 ps with relative amplitude of 110% is shown in Figure 61.

Three principal algorithms are compared: RSS, first peak-finding similar to [44], and the developed IPS method. In Figure 62 the errors (in mm) of the three methods in finding the true peak are compared for the simulated multipath signal shown in Figure 61. The errors for first peak-finding and iterative peak subtraction algorithms are highest when the first two multipath peaks begin to exhibit extensive overlap (< 300 ps). With large separation, RSS has large errors due to the algorithm exclusively finding the secondary peak. It should be noted that the RSS and first peak-finding algorithms will return identical results for all situations where the primary peak has the largest amplitude as well as situations where any additional multipath peak(s) add to the main peak to form one large, combined peak. Figure 63 and Figure 64 explain the cause of this error.

Initially, the primary and secondary peaks are nearly aligned. Even though the two peaks combine to form one combined peak, the center of this combined peak is positioned in close proximity to the center of the true peak. This results in low errors for separations in the range of 0-50 ps. As the separation between the two peaks increases, the center of the combined peak shifts. This results in increased error between the center of the combined peak and the center of the true peak, as shown in Figure 63. Finally, as the two peaks continue to become more separated, the combined peak splits into two distinct peaks. This once again reduces the error in locating the center of the true peak, as is shown in Figure 64. This behavior indicates that the error due to this effect is initially small, increases to a maximum near 150 ps (half the pulse width), and decreases back towards zero as the peaks become separated from each other.

The maximum errors found across various time delays over a range of amplitudes are summarized for each of the primary peak detection methods in Table 20. The iterative peak




Figure 61: Simulated multipath signal with a secondary peak 110% of the main peak delayed 440 ps.




Figure 62: Comparison of error (in mm) for the three peak finding algorithms versus separation of primary and secondary peaks for secondary peak amplitude of 110% Figure 61.




Figure 63: Interference of two 300 ps Gaussian pulses separated by 138 ps with the amplitude of the secondary pulse 110% that of the primary pulse. The large overlap between the pulses causes a significant shift in the peak of the combined pulse.




Figure 64: Interference of two 300 ps Gaussian pulses separated by 262 ps with the amplitude of the secondary pulse 110% that of the primary pulse. The increased separation between the pulses reduces the lagging shift, and consequently the error in locating the peak position

subtraction algorithm yields considerably better results than either RSS or first peak-finding, especially for multipath conditions where the secondary peak has an amplitude as large or larger than the primary peak. It is necessary to note that in these simulations, the secondary peak was extremely close to the primary peak. In actual received multipath signals this is not often the case, indicating that the iterative peak subtraction algorithm will accurately detect the correct peak with a high degree of accuracy. Also, the Max-Ratio Algorithm has improved results over the Iterative Peak Subtraction algorithm and will be discussed in depth in the next section.

5.4. Max-Ratio Leading Edge Detection

In order to solve some of the inconsistencies of the IPS algorithm, a secondary algorithm was developed that was based on the LOS Ray tracing results discussed previously. The Max-Ratio algorithm seeks to reduce the overall algorithm complexity, minimize the use of floating point operations, and eliminate iteration steps. The main rationale for the algorithm is the simplified channel model described by Equation (8). In this model the most uncorrupted portion of the

Table 20: Comparison of maximum error for varying secondary peak amplitudes

Relative Height of Second Peak normalized amplitude - dB	RSS (mm)	First Peak (mm)	Iterative Peak Subtraction (mm)	Max-Ratio Algorithm (mm)
0.01 -40 dB	1.50	1.50	0.37	0.05
0.05 -26 dB	2.25	2.25	0.37	0.25
0.1 -20 dB	3.00	3.00	2.62	0.53
0.25 -12 dB	6.00	6.00	4.50	1.47
0.5 -6.0 dB	10.9	10.9	8.24	2.54
0.75 -2.5 dB	16.1	16.1	12.4	3.79
1.0 0.0 dB	133	22.1	16.5	5.86
1.1 0.8 dB	133	25.1	18.7	6.54
1.25 1.9 dB	133	58.5	37.1	7.65

signal that can be used for positioning is the rising edge of the first (LOS) ray. Additionally, in NLOS conditions selecting the first ray is considered optimal [100].

Algorithm 2: Max-Ratio Leading Edge Detection

1. The received signal is described in the discrete time domain as

$$h_0[t] = \sum_n a_n s[t - \tau_n] \quad (30)$$

where a_n is the ray amplitude, and τ_n is the ray arrival time. The signal $s[t]$ represents the clean signal envelope.

2. The received signal is passed through a sliding window average filter given as

$$y[t] = \text{avewindow}(h_0[t], 16) \quad (31)$$

where the function $\text{avewindow}(\cdot)$ maps the original signal to a filtered version of the signal using a window size given by the second parameter.

3. The averaged signal is then passed through two sliding window max filters

$$\text{max_16}[t] = \text{maxwindow}(y[t], 16) \quad (32)$$

$$\text{max_256}[t] = \text{maxwindow}(\text{max_16}[t], 256) \quad (33)$$

where the function `maxwindow(.)` returns the value of the maximum valued sample within the window.

4. By taking a ratio of these two maximum filtered signals and comparing the ratio to a fixed threshold

$$r[t] = (\min_16[t] * 4 \geq \min_256[t]) \quad (34)$$

where the value 4 is taken as the threshold, a resulting binary valued signal $r[t]$ is created which can be used to indicate the time of arrival of the leading edge.

5. Alternatively, the time of the leading edge can be solved using a least-squares approach to calculate sub-sample resolution for the leading edge arrival time.

$$FP_t = \arg \min_t ((\min_16[t] * 4 - \min_256[t])^2) \quad (35)$$

Figure 65 shows an example of the Max-Ratio pulse detection algorithm displayed graphically.




Figure 65: Example of Max-Ratio Leading Edge Detection. The original leading-edge signal $h_0[t]$ (blue) is averaged $y[t]$ (solid red), subsequently two max-window filters $\max_16[t]$ (dashed red) and $\max_256[t]$ (dashed green) operate on $y[t]$, the result $r[t]$ (light blue) makes a positive transition when the ratio of $\max_256[t] : \max_16[t]$ is less than 4

5.4.1 Max-Ratio Leading Edge Simulation Results

In order to simulate the performance of this algorithm under multipath interference conditions, a simulation was performed to examine the effect of 1 secondary multipath peak on the leading edge calculation. In this experiment, the LOS metal reflector model was used to generate the characteristics of the signal containing the secondary multipath peak. In simulating the worst case, a multipath peak with relative powers ranging from -40 dB to 1.9 dB with respect to the LOS peak was shifted by different time offsets. The algorithm has a very consistent error profile (< 2.0 mm) when the multipath peak offset exceeds 220 ps (66 mm spatially) as shown in Figure 66. Considering the final column in Table 20, the maximum error for a worst-case signal such as +1.9 dB relative to the LOS peak is only 7.65 mm. Compared to 133 mm, 58.5 mm, and 37.1 mm for RSS, First Peak, and Iterative Peak Subtraction, respectively, this result is a considerable improvement.

Once the final pulse position has been calculated, optionally, a fixed-point interpolant can be used to compute a sub-sample pulse position. Using this method and varying the pulse between




Figure 66: Max-Ratio Leading Edge Simulation: Error in positioning due to multipath peaks with various power levels (1.9 dB to -40 dB) as detected by max-ratio leading edge algorithm. Note the maximum values are recorded in Table 20




Figure 67: Max-Ratio Leading Edge Simulation: Error in positioning due to sub-sample shifts in peak location, noting that very little error results in this case

successive pulse indices has exceedingly small errors ($\ll 1$ mm), as seen in Figure 67. This can be taken to mean that should the correct leading edge location occur between two samples, the error due to the sample spacing is bounded.

In order to characterize this pulse detection method in the presence of Additive White Gaussian Noise (AWGN) the following noise model was adopted as an extension to (8)

$$h_o(t) = \sum_n a_n(t) s(t - \tau_n(t)) + n(t) \quad (36)$$

where $n(t)$ is drawn from a zero mean homoskedastic i.i.d. process. In order to report the results with the highest degree of specificity, the SNR of the received signal $h_o(t)$ was measured using the technique,

$$h_o(t) = r(t) + n(t) \quad (37)$$

$$V_{RMS}^r = \sqrt{\frac{\sum_{i \in T_p} r(i)^2}{N}} \quad (38)$$

$$V_{RMS}^n = \sqrt{\frac{\sum_{i \in T_p} n(i)^2}{N}} \quad (39)$$

$$SNR_{meas} = 20 \log_{10} \left(\frac{V_{RMS}^r}{V_{RMS}^n} \right) \quad (40)$$

where the received signal is composed as signal $r(t)$ and noise $n(t)$ components. The RMS power for each component can then be found using (38) and (39), where T_p represents the segment of time containing only the LOS pulse and N is the number of samples in region T_p . Finally, the SNR is computed the usual way. By using this procedure, effects such as multipath interference and pulse ringing do not artificially inflate the SNR values, and thus, offer a more precise picture of noise degradation effects on this algorithm. A simulation was enacted, by creating 10 instantiations of random noise using SNR values that ranged from roughly -2 dB to 95 dB. As demonstrated in Figure 68, error as a result of noise degradation has a fairly linear decrease from 2 dB to 15 dB, with the exception of one outlier. For this algorithm, error was bounded by 1 cm for SNR values greater than 14 dB, and error was bounded by 1 mm for SNR values greater than 30 dB.




Figure 68: Max-Ratio Leading Edge Simulation: Error resulting from degradation in SNR due to AWGN, 10 random trials. Note that with the exception of one outlier, error decrease is roughly linear for SNR values between 2.5 to 15

In Figure 69, an instantiated signal is shown which represents the case of the algorithm operating in a high SNR environment (~ 30 dB) and a correspondingly low error (0.56 mm) when compared to a noise-free reference signal. Shown in Figure 70 is a medium SNR (~ 15 dB) environment, with increased error (4.3 mm). Also, Figure 71 demonstrates an instantiated signal in a low SNR environment (3.2 dB), with correspondingly high error of (28.7 mm). Also displayed in each of these examples is a noise threshold which is determined by sampling the noise prior to the onset of the pulse. For these simulations, the maximum of a 100 sample window is used as the noise threshold. Considering this threshold, any peaks that are found while the signal level is strictly below the threshold are discarded. Note that this threshold computation is realistic in the finalized system (and has been implemented as such). This is due to the fact that the pulse repetition period (PRP) after sequential sampling is 1 ms, which allows for sufficient noise sampling between successive pulses. Additionally, in scenarios where round-robin or TDMA protocols are used to compute many tag locations, the noise level may be dynamically set on an individual tag per pulse basis, thus increasing the robustness of the system.

5.4.2 Digital Design of Max-Ratio Leading Edge Detection

The Max-Ratio Leading Edge Detection algorithm has been implemented on an FPGA using Verilog HDL and been verified using a digital signal software simulator (ModelSim, Mentor Graphics [127]). The design methodology for the FPGA implementation follows a pipeline approach facilitating real-time peak detection and minimizing overhead typically associated with memory storage or convolution steps. In this pipeline design methodology, data is input to the device on the rising edge of a global clock signal, and after a fixed time delay, the corresponding determination of the leading edge point is available as a binary signal taken off the device. The




Figure 69: Max-Ratio Leading Edge Simulation: Instantiated signal (blue) with high SNR (30.0) and associated low error (0.56 mm), showing calculated noise threshold (black)




Figure 70: Max-Ratio Leading Edge Simulation: Instantiated signal (blue) with medium SNR (14.6) and associated medium error (4.30 mm), showing calculated noise threshold (black)




Figure 71: Max-Ratio Leading Edge Simulation: Instantiated signal (blue) with low SNR (3.2) and associated higher error (28.7 mm), showing calculated noise threshold (black)

time delay of the pipeline is solely dependent on the nominal values of parameters of the Max-Ratio Algorithm. Another core philosophy driving the design of this device is reconfigurability, allowing the device to be used in different configurations or for perhaps different applications, such as IR communication schemes. Based on the capacity limits of the FPGA itself, two Max-Ratio algorithm modules can be placed on one chip. As shown in Figure 72, the schematic shows the general layout of the input and output pins of the device, while all of the processing is encapsulated in the Base Station Module. At the basic level, the Base Station Module has two input buses—one for each data stream inputted into the Max-Ratio algorithm implemented on the device. Also, the Base Station Module has several other debugging and tuning inputs, such as a threshold adjustment input, push button reconfiguration of the PRP, the system clock, and a synchronous reset. At the exit, the device displays the current PRP on the 4-segment LCD display, the noise threshold is displayed dynamically on an 8 LED array, and most importantly the module emits a binary valued signal indicating a leading edge has been detected. This




Figure 72: Xilinx schematic describing pin inputs and outputs when Base Station Module is configured in one sample channel (I) per base station configuration as depicted in Figure 73

detected signal has the same period as the PRP so that the resultant signal can be used as a trigger for other functionality. Essentially, the leading edge of this signal is at a slight delay with respect to the detected leading edge, but with a fixed offset as previously discussed. For the primary configuration, the Base Station Module handles the signal processing and leading edge detection for two base stations as can be visualized in Figure 73, where the DSP boxes in this case represent instantiations of the Max-Ratio Algorithm. In this setup, each base station has its corresponding in-phase down converted signal (I) sampled by each channel of a two channel ADC. Thus, the inputs are each digitized signal and the outputs are the leading edge of each pulse, respective to each base station. The downside to this approach is that each of the two base stations feeding the FPGA must be located in close enough proximity to be wired to the same device.

A second configuration as illustrated in Figure 74, shows the device connected such that the in-phase (I) and quadrature (Q) down sampled signals originating from the same base station are sampled using a two-channel ADC. The signals are independently passed to the Max-Ratio leading edge detection modules on the FPGA, and independent values of detected leading edges are passed out of the device. In subsequent processing, the time-locations of these leading edges can be averaged or processed in some other filtered manner to produce a final TOA result. This




Figure 73: One sample channel (I) per base station configuration

configuration offers an attractive option to processing I/Q data, given that the I and Q channels have been found to be non-overlapping by a fixed offset (see Figure 12). The downside to this configuration is that it requires one ADC, one FPGA, and two sampling mixers per base station. In the testing and discussion that follows, the previous configuration (Figure 73) is the configuration that was chosen to perform the final system testing and validation, due to a lack of availability of two sampling mixers per base station.

Supposing that the I/Q channel offset can be calibrated, a third configuration is possible. In this configuration, the signal envelope can be directly calculated on the FPGA as $I^2 + Q^2$. In this situation, two base stations can be served by one FPGA in conjunction with two ADCs. It is further noted that by using this configuration, the discrete time low-pass filter (LPF) module and the rectification module in the description that follows could be safely eliminated.

As implemented, the digital signal processing (DSP) steps will now be described in more detail. In the description that follows, it is assumed that the device is operating on non-enveloped single channel data (i.e. I or Q only). Shown in Figure 75 are the main sub-modules of rectification (denoted by I^2), low pass filtering (LPF), peak detection, and signal hold, as well as the auxiliary modules of function select and noise detection. As the signal is received by this device, it represents the raw digitized data transmitted by the ADC. The data width is nominally 10-bits




Figure 74: Two sample channels (I and Q) per base station configuration




Figure 75: Overall diagram of FPGA implementation of DSP module

and while the architecture is parameterized to this width, it can be readily changed in the design software. The 10-bit midpoint of the raw data is subtracted and the result is squared, resulting in strictly positive data.

Subsequently, the data is passed to the LPF module as depicted in Figure 76, which implements a 7-pole equiripple FIR filter (shown is only a 4 pole filter). In this step, the data is converted to 16 bits for full resolution of the filter, and subsequently returned to 10 bit values at the output. Due to the limited number of full bit multipliers available on the FPGA, the 7-pole filter is the optimal design considering two channels per FPGA.

Depending which function is selected by the Function Select Module, either the Peak Detect Module (Figure 77) is enabled when the system is in Peak Search Mode, or the Noise Detect Module is enabled when the DSP Module is in Noise Sampling Mode. Supposing that the Peak Detect Module is enabled the LPF data is viewed by the Peak Detect Module as potentially containing a leading edge. Also, a positive valued noise estimate is also passed to the module, representing the maximum valued sample of LPF data when the DSP Module was last in Noise Sampling Mode. The data is passed to each of two MaxWindow Modules, one with a window




Figure 76: Module-level diagram showing internal structure of FIR LPF




Figure 77: Module-level diagram showing internal structure of Peak Detection functionality

size of 272 and one with a window size of 16. Because MaxWindow Modules differ in window size (by design) the data exiting the modules is misaligned by the difference of the window sizes. In this case a simple first-in-first-out (FIFO) buffer of nominal length 256 is used to realign the data. In order to compare the data according Max-Ratio Algorithm fixed threshold of 4 (Algorithm 2, step 5), the MaxWindow signal corresponding to the filter size of 272 is bit-shifted right and compared to the delayed signal corresponding to the MaxWindow filter of size 16. This intermediate result is combined using a logical AND operation with another intermediate comparison, which ensures that both signals are above the Noise Estimate input. The final binary valued signal is passed to the Hold Module where a positive value is retained for half of the PRP. After the Hold Module has retained the signal for half of the PRP, the Hold Module triggers the Function Select Module to switch states placing the DSP Module in Noise Sampling Mode. This state change enables the Noise Detect Module shown in Figure 78. This module also has a parameter value representing the PRP as a total number of samples (not shown), which can be set by user depression of a push-button on the evaluation board. Since this module is triggered near the midpoint of the pulse repetition cycle (i.e. between two successive pulses), the noise sampling can use any fraction up to one half of the PRP to perform noise sampling, though in




Figure 78: Module-level diagram showing internal structure of Noise Detect module

practice PRP/4 or PRP/5 is sufficient for noise characterization. At the onset of the noise sampling, the MaxWindow module is reset to eliminate any leftover data. The MaxWindow module window size is set to the appropriate fraction of the PRP and sampling takes place for exactly the length of one window. Upon completion the maximum valued sample is stored in a register which provides the Noise Estimate value when the DSP Module exits the Noise Sampling Mode. Should an extreme noise sample value artificially inflate the Noise Estimate value, the Peak Detection Module will likely not detect the next pulse. If no pulse is detected, the control signals of the Hold Module will keep the DSP Module in Peak Search Mode, consequently, the Noise Detect Module will not be enacted to adjust its erroneous Noise Estimate value. In this case, a Counter Module signals a single right bit-shift operation of the Noise Estimate, thus reducing the threshold by a factor of two, at the end of each PRP. In the typical case when a pulse is detected each cycle this functionality is never utilized. In cases when the SNR degrades completely, the Noise Estimate value is reduced *ad infinitum* until a detected pulse again places this module in Noise Sampling Mode. Optionally, switches located on the FPGA board have been used to implement a user-settable lower limit on the Noise Estimate value. When used in combination with tag switching protocols, this functionality would need to be carefully interleaved to ensure that the Noise Estimate is not corrupted by pulses of the previous tag transmission.

5.4.3 Digital Design of Max-Ratio Leading Edge Detection

5.4.3.1 Device Simulation Testing

Upon design completion of the Max-Ratio algorithm in Verilog HDL, verification was needed to complete the design and ensure that the design matched the results of the simulation in Section

5.4.1. Using ModelSim, test signals were constructed from raw data collected from the ADC under a previous FPGA architecture. Shown in Figure 79 are the simulation results showing input signals and resulting leading edge output signal.

Across multiple signals, the simulated result differed from the ModelSim result by a constant difference of -1, as tabulated in Table 21. The constant offset originated in a fundamental difference in how each program handled indexing. Essentially, any fixed difference between two base stations would be eliminated using TDOA or TOA techniques which are discussed in more detail in Chapter 6.

5.4.3.2 Device Experimental Testing

The Max-Ratio algorithm has led to vast improvements over the RSS-type algorithm that was previously implemented for the system. The main flaw with the RSS algorithm, as with many signal strength dependent algorithms is that the RSS algorithm was extremely sensitive to a




Figure 79: ModelSim signal simulation showing Max-Ratio Algorithm as implemented in Verilog HDL targeted for FPGA

Table 21: Comparison of Max-Ratio algorithm simulation result and ModelSim result, demonstrating a consistent leading edge result

Signal Name	Simulation Result	ModelSim Result	Difference
I1	243	242	-1
Q1	251	250	-1
I2	250	249	-1
I3	262	261	-1
I4	258	257	-1

manually defined threshold set on the FPGA board by the prospective user. This threshold had nefarious effects on the 1D positioning values due to the changing received signal level as the tag was translated. Essentially, there was a large system non-linearity as the tag was translated, due to the previous algorithm's dependence on the frequency offset between the tag and system LOs. Since the Max-Ratio algorithm relies on dynamic thresholding, sensitivities to the current noise level are minimized by dynamically setting the threshold on every pulse cycle (Figure 80).

5.5. Genetic Programming Edge Detection

Genetic programming (GP) is an automatic method for developing expressions or programs, first proposed by John Koza [128]. The literature is rich with genetic programming used in signal processing tasks. Such uses include discovering signal peptides in peptide chains [129], customizing control programs for prostheses, solving DSP problems of channel equalization, noise cancellation, interference removal [131], and vector-based GP for symbol rate detection[132].

Recently, computational intelligence techniques have been used in UWB applications. Hatscher and Diskus [133] used a genetic optimization technique in a see-through-wall [134] [133] application to fit model parameters to multipath echoes, ultimately for accurate signal




Figure 80: Average 1D error in mm versus sample index for prior fixed-threshold algorithm (red) and Max-Ratio algorithm (blue), noting that both the maximum error and the average error show significant improvement

reconstruction. These imaging applications essentially have the same problems as positioning systems, which is finding acceptable signal characters to map to the spatial domain. Positioning accuracy in 1D and similarly axial resolution in imaging applications can be related to time domain resolution. In imaging applications, cross-range accuracy can be improved via use of synthetic aperture techniques, whereas in positioning applications 3D accuracy is improved by reducing GDOP by the addition of well-positioned base stations. Khanbary and Vidyarthi used genetic optimization techniques to provide channel allocation optimization in an array of base stations [136]. Also, Colman and Willink used genetic algorithms to optimize base station signal processing on overloaded networks. In this technique, the GA was initialized using soft biased method to ensure adequate population diversity upon algorithm commencement [137]. A motivation for using GP techniques in this body of research is to allow for the automated testing, developing, and optimizing of novel algorithms to perform the peak/leading edge detection.

5.5.1 Genetic Programming Peak Detection

To use genetic programming for the purpose of peak detection, it is important to consider the representation of the genetic program in order that the signal processing steps are capable of implementation on an FPGA. In Figure 81, a tree structure is shown which indicates the connectivity between several nodes each of which is a function or module which can be instantiated on FPGA hardware. In this example, the terminal nodes, or inputs, to this signal processing algorithm are the IQ channel data streams as well as several constants (i.e. 1, 2, 4, etc.). Since the data from the ADC continuously flows into the FPGA device, simulating this algorithm requires vectors of data that are representative of the IQ data stream. At each time instant, the current values for the IQ data are set at each terminal node, and the tree is evaluated using a post-order traversal method.

Figure 82 highlights some of the high-level GP algorithm choices that have been selected for this study. Given that the purpose of any peak detection algorithm used for the system is aimed at minimizing the potential for error, the objective function was chosen as the Mean Negative Variance (MNV) of the computed peak location over several IQ channel realizations. This metric was chosen as opposed to a metric which designates a certain position on the signal. This choice




Figure 81: Example of a signal processing algorithm in tree representation for use in genetic programming

Objective Function	Initialization	Crossover	Mutation	Selection
• Mean Negative Variance (MNV) of Peak Locations • MNV + Minimum Description Length	• Ramped $\frac{1}{2}$ & $\frac{1}{2}$	• Subtree switching (Maximum Depth Protection)	• Random Sub-tree Deletion with replacement	• Roulette (Fitness Proportionate) • Ranked

Figure 82: Genetic Programming Algorithm Choices

was to allow the GP result to find the most useful portion of the signal. The algorithm progresses using crossover and mutation between various trees and selection of the best performing trees to increase peak detection performance.

5.5.2 Functional and Terminal Node Selection

For any algorithm to be able to be implemented on an FPGA, only functions which can be implemented as Verilog modules are included for use. Table 22 gives a partial list of functional nodes which have been tested as part of this study. The number of parameters of a function dictates the number of child nodes a particular functional node contains.

One of several problems encountered during the selection of functional nodes during this study, is the digital verifiability of the functional node itself. While the digital verification of nodes such as addition, subtraction, and multiplication remains trivial, functions such as delay are intractable. This is primarily due to the abstraction of memory provided by the HDL design software. For example, in the case of a delay node the 2nd parameter governs the delay length, which may be variable. Since the on-board memory is limited, variable delay length in the Verilog implementation has a fixed cut-off, which in practice causes divergence between the

Table 22: Partial List of functional nodes used with genetic programming		
Function	# of Parameters	Description
delay	2	Delay 1 st by 2 nd
maxwindow	2	Max of 1 st in window size 2 nd
dot	2	Cumulative dot product
+	2	Element-wise sum
*	2	Element-wise product
sq	1	Element-wise square
ema	2	Exponential Moving Average (FIR filter)

software simulation of the GP performance and on-chip performance. Secondly, it is important for every operation performed on the data stream to be closed. In the software simulation, this requirement is fulfilled by each node accepting vector types as input and producing one vector type as output. The Dot Product Node, for example, is implemented as a cumulative dot product, thus, the intermediate results of a full dot product operation are available as the output vector. Again, practical implementation constraints inherent to FPGA design force the use of fixed width integers to serve as the underlying data type. Thus, while the final dot product value may conform to the closure property, intermediate results may experience overflow or underflow. Even if the operation were to be implemented internally with higher bit width integers or fixed point numbers, the intermediate results could still suffer from overflow/underflow problems. If the synthesized algorithm is dependent on properly-valued intermediate results to function correctly, then the software simulation and the FPGA implementation would yield different results.

Terminal nodes serve as the input to the signal processing program. In one set of simulations, the terminal nodes consist of the IQ digital channels as well as the integer constants: 0, ± 1 , ± 2 , ± 4 . This IQ simulation represents the idealized case, as properly aligned and sampled I and Q channels should enable unambiguous envelope calculation. To facilitate this simulation, 30

instances of synthetic I and Q channel data were created and corrupted with AWGN to an SNR of 30 dB. A second simulation used 60 instances of actual I-channel data, which were acquired from the system in a static LOS experiment. This data was used to provide a realistic scenario in which only one channel was available.

5.5.3 Genetic Programming Parameters

While several runs of the GP algorithm were performed to test parameter settings with which the GP run produces adequate results, the ultimate goal of these simulations was to find a solution which adequately provided a comparable program to the previously discussed Max-Ratio algorithm. Table 23 gives the primary parameters controlling a single GP run. These nominal values achieved results that were consistent with algorithm convergence, although occasionally the resultant program was too complex to be implemented on the FPGA. Often this came in the form of too many gates being required for the FPGA implementation, but sometimes it was a result of the GP-generated program consuming too much of one of the finite resources on the chip. Since any delay functionality is essentially implemented as a FIFO buffer, on-board memory was occasionally consumed by an excess of delay nodes. Also, the particular chip used in this study only supports ten 18x18-bit multipliers, so excessive amounts of multiply or dot product nodes could exceed this maximum.

Table 23: Parameters studied controlling one GP Run

Parameter	Settings Tested
Generations	50
Population Size	100
Percent Mutation	30%
Percent Crossover	25%,55%,75%
Max Tree Height	6,8,10,12
Repetitions	10 times
Training, Testing Percentages	33%, 33%
Probability of Terminal in “grow” mode	50%

The number of new members per generation is given by:

$$N_{indiv} = P_{size} (Mutation_{per} + Xover_{per}) \quad (41)$$

5.5.4 Genetic Programming Challenges

Using Negative Mean Variance for the objective function, sometimes the algorithm converges to consistent values, but just uses deterministic combinations of the constants (i.e. does not fully utilize IQ signals). A solution to this problem is to use randomized small shifts (Figure 83) to the peak locations (± 10 sample positions) to prevent this happening. By subtracting off the small peak shifts from final peak location, deterministic solutions will be eliminated due to adverse score.

This solution was essentially required for proper solution convergence of a GP run. Without this functionality, any deterministic program would satisfy a MNV value of 0, which is the highest possible score. Deterministic programs found in this manner would be excessively propagated to future generations corrupting the results significantly.

Another problem plaguing the GP paradigm in general is that of program bloat. Over the course of the GP run the resulting programs have a tendency to grow excessively large. The solution to




Figure 83: Diagrammatic view of shifts applied to noisy signals, to avoid deterministic genetic programs

this problem is two-fold. A hard cap is used to prevent any trees from exceeding a certain height (“Max Tree Height” in Table 23) under mutation and crossover operations. A soft cap is also used to put trees that are increasingly complex at a scoring disadvantage relative to programs which are simpler. In order to implement the soft cap the principle of MDL is utilized as a score adjusting mechanism. As given by

$$MDL = \log_{10}(1 + N_{nodes}) \quad (42)$$

$$S = MNV - MDL \quad (43)$$

where the MDL term consists of the $\log_{10}$ of one plus the total number of nodes in the tree. To compute the adjusted score, the objective function subtracts the MDL from the MNV. Thus, if the number of nodes is increased by a factor of ~ 10 , then the MNV would have to decrease by one unit for the solution to be viable.

5.5.5 Genetic Programming with Boosting

In order to achieve a robust algorithm that can deal with various situations, the boosting meta-algorithm is proposed as a solution. The boosting algorithm is useful for combining previous solutions of GP output into a weighted combination. As implemented for the UWB positioning system, this consists of weighting the final peak position output (as opposed to weighting the intermediate signals) (Figure 84).




Figure 84: Example of how to apply the boosting meta-algorithm to improve genetic programming results

While boosting was examined briefly, the trees produced by combining either 4 or 8 sub-trees were typically too large to implement on the FPGA, while not offering a significant improvement in error. Also, as will be discussed in Chapter 6, non-linear pulse position filtering was found to have excellent statistical characteristics versus the linear combinatorial form illustrated in Figure 84.

5.5.6 Tree Simplification Engine

An interesting effect of implementing programs created using GP on an FPGA device is that intermediate signals, while solved for recursively in software, may be reused in hardware. Thus, simplification routines were developed to reduce the nodes from the final solution of a GP run, to a representation which is useable as a connected set of Verilog modules. This engine takes into account commutativity rules associated with each functional node. The routines that provide this functionality are highlighted in Appendix C.

5.5.7 Genetic Programming Simulation Results

Using the Max-Ratio Algorithm for reference and IQ synthetic data as input, Table 24 shows the GP resulting program which has a higher score of -1.478 vs. -5.15 (indicating a higher accuracy over a battery of noisy input signals). Note that the GP result has almost half the number of nodes required by the Max-Ratio Algorithm, and the tree height is also smaller. The tree height

Table 24: Synthetic Data Genetic Programming Results

Case	Tree Height	Number of Nodes	Score	Signal Processing Program
Max-Ratio Algorithm	9	57	-5.15	$(((((sq(I) + sq(Q)) \text{ ema } 4) \text{ delay } (sq(sq(4)) - sq(4))) * (sq(sq(sq(4))) * (4 * 2))) > 1) \text{ AND } ((((((sq(I) + sq(Q)) \text{ ema } 4) \text{ maxwindow } sq(4)) * 4) \text{ delay } (sq(sq(4)) - sq(4))) > ((sq(I) + sq(Q)) \text{ ema } 4) \text{ maxwindow } sq(sq(4))))))$
GP Result	8	29	-1.478	$(((((sq(I) + sq(Q)) \text{ ema } sq(sq(4)) \text{ delay } first(first(cumsum(I))) * (sq(sq(sq(4))) * (4 * 2))) > 1) \text{ AND } (3.1)))$

can be thought of as the “critical path” when the program is implemented, as this is the minimum number of sequential steps from the input data to the final output.

Poor algorithm convergence clearly results if the maximum tree depth is set to six as shown in Figure 85. Across ten runs using the synthetic dataset, the best population member remains nearly constant, while the population mean makes progress in only several of the runs. The average best member is computed using only the MNV portion of the score. In this case, an average best member score of -4.70 corresponds to a standard deviation of 1.30 mm, assuming sample spacing of 0.6 mm per sample. Despite the algorithm’s poor convergence, this is still an acceptable solution to the peak detection problem.

When the maximum tree depth is increased to eight, immediately noticeable are progression of the best member and population mean (Figure 86). Unusually, past the 20-30 generation mark,




Figure 85: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 6, 30% mutation rate, 55% crossover, MDL scoring, 100 individuals.




Figure 86: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 8, 30% mutation rate, 55% crossover, MDL scoring, 100 individuals

the population mean typically decreases, while the best member consistently increases. This may be an artifact of program bloat despite using MDL as a scoring metric.

When the maximum tree depth is increased to ten (Figure 87), the resultant average best member shows marked improvement over the case when the tree depth is limited to six. This indicates that a tree depth of six is insufficient to fully represent a solution to the signal processing problem. Also, the best member convergence is more rapid and consistent across the ten runs.

Comparing crossover rates between tree depths of eight and ten, lowering the crossover rate slightly improves the average best member of the depth eight run (Figure 88), while slightly increasing the average best member of the depth 10 run (Figure 89).

5.5.8 Genetic Programming Experimental Data Results

For this section of the genetic programming simulation, the programs operated on 60 instances of captured I-channel data in addition to the constant terminal nodes of 0, ± 1 , ± 2 , and ± 4 .




Figure 87: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 10, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals.




Figure 88: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 8, 30% mutation rate, 25% crossover, MDL scoring, 100 individuals




Figure 89: Genetic programming run on synthetic IQ data (30 signals), (blue best member, green average member, red worst member), for maximum tree depth 10, 30% mutation rate, 25% crossover, MDL scoring, 100 individuals.

Consequently, the signal envelope cannot be directly calculated, thus, ambiguities in the pulse position pose a much more difficult problem for the signal processing program to solve. In comparing these results to the previous section, it is clear that the positioning programs do not perform as well. Examining Figure 90, a tree depth of eight, with 30% mutation rate, 75% crossover rate, and MDL scoring, results in an average best member of -326.1 or correspondingly a standard deviation of 10.8 mm. An example Verilog implementation generated by GP is shown in Appendix B, which was generated from the best member (score: -134.4) of the Figure 90 run. When the maximum tree depth is increased to ten (Figure 91), a slight performance increase is noted in the average best member. Also, the convergence of the population mean is rapid and more consistent than the previous case. However, increasing the tree depth to size twelve, shown in Figure 92, fails to demonstrate any significant advantage over the depth ten case.




Figure 90: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 8, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals.




Figure 91: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 10, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals.




Figure 92: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 12, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals. This run utilized the full set of functional nodes

The best member (score: -142.18) from the run depicted in Figure 92 was selected for FPGA implementation. This member was 31 nodes after simplification, and included the fullest possible set of functional nodes: add2, sub2, mult2, neg, square, dot, delay, cumsum, ema, maxwindow, first, AND2, OR2, and greater. The resulting peak detection algorithm was synthesized and integrated onto the FPGA in the single-channel per base station mode. The resultant peak detection algorithm failed to produce a single detected peak when attached to the RF microwave system. This could occur for any number of reasons previously mentioned, such as overflow/underflow under the operations cumulative summing or cumulative dot product. Also, the genetic program may have been tuned to one specific amplitude where the peak typically occurred.

On a second attempt at producing a valid peak detection signal processing routine, the functional node set was reduced. In this secondary experiment, the functional nodes were: add2, sub2, mult2, neg, square, delay, cumsum, maxwindow, AND2, OR2, and greater. The parameters of

this run, shown in Figure 93, were identical to the run shown in Figure 92, with the exception of the functional node set. The result of this implementation was tested for 1D static ranging error (Figure 94) compared to the Max-Ratio algorithm. One of the striking differences is that the Max-Ratio algorithm has lower maximum error as well as consistently lower average error. One of the potential reasons is the possible reliance of the GP algorithm on a fixed-height pulse, or loss of precision issues previously discussed. In light of these challenges, further work would need to be added to the GP portion of this study, to ensure that the simulation and scoring portion of the GP software exactly matches the constraints and data representations of the FPGA. Also, the experimental data provided to the GP software simulator should have a larger swing of SNR values, pulse positions, and types of multipath corruption to ensure a robust solution.




Figure 93: Genetic programming run on experimental I-only data (60 signals), (blue best member, green average member, red worst member), for maximum tree depth 12, 30% mutation rate, 75% crossover, MDL scoring, 100 individuals. This run utilized a sub-set of the possible functional nodes




Figure 94: Average 1D error in mm versus sample index for GP algorithm (red) and Max-Ratio algorithm (blue), noting that both the maximum error and the average error are worse for the GP algorithm than the Max-Ratio algorithm or the previous fixed threshold method (Figure 80)

6. ROBUST 3D POSITIONING IN MULTIPATH ENVIRONMENTS

6.1. Background

In the field of 3D positioning, a significant body of research exists that demonstrates various techniques in transforming 1D range values from a number of base stations into a 3D location.

The most basic of these are RSS measurements at each base station. In this scheme, the actual amplitude of the received signal is used to infer the distance between the tag and the base station.

The RSS positioning algorithm must assume a path-loss function to perform this inference, and thus, blockages and varying transmitter power levels will disrupt this type of positioning. The methods of multidimensional scaling (MDS) [89] and maximum-likelihood estimation (MLE) are used by Li [138] to study collaborative localization in received signal strength (RSS) sensor networks. In this study the purpose was to evaluate positioning of multiple tags through the collaborative measurement of each tag. The TDOA algorithm is well known [116] and used in a variety of wireless and wired positioning applications. This algorithm is well-suited to UWB positioning because it does not require that the transmitter and the receiver be synchronized. Chan and Ho examined the problem of time scale and TDOA positioning, in the case that the mobile tag movement causes a Doppler shift in received signals [88] similar to [139]. TOA algorithms [62] rely only on time-of-flight information and are heavily used in the field of GPS. Similar to the TDOA algorithm, TOA algorithms do not require explicit synchronization between transmitter and receiver. Both TDOA and TOA algorithms degrade when sufficient short-term jitter or long-term wander between clocks alters the perceived nominal time scale. If

receiving antenna arrays are used, AOA [140] methods may be used as a standalone algorithm or in conjunction with TOA methods [141].

6.1.1 Robust Positioning Methods

According to differing system conditions, various positioning methods have advantages and disadvantages. A frequently used algorithm for UWB positioning is the TDOA method, which has been previously discussed. This algorithm works well when all base stations are synchronized with the same PRF clock and, in general, offers higher performance when the PRF clock is also coherent with the tag transmission PRF clock. In GPS, the TOA algorithm is widely used, because the satellites are highly synchronized via atomic clocks, but the user receiver remains unsynchronized with the system. Thus, the TOA algorithm provides the time bias of the user with respect to the system as part of its positioning solution. In the Microwave Positioning Subsystem the transmitter and receiver use high-grade 10.0 MHz crystals with ± 0.5 ppm stability. Despite the clock's excellent stability, when these two clocks are used in conjunction with a Direct Digital Synthesizer (DDS) to provide for the sequential sampling, an unintentional time scaling of $\pm 11\%$ is theoretically possible when a scale factor of $\alpha = 100,000$ is used [42]. To deal with these wide fluctuations in scale, a new algorithm is proposed, termed scaled-time-of-arrival (STOA). The TDOA algorithm, by comparison, ignores any time offset between transmitter and receiver and assumes no frequency bias between clocks [62]. While the TOA directly solves for the time bias, the STOA algorithm interprets any systematic lengthening or shortening of 1D range values as a change in scale. Table 25 compares the input vector and the output solution for all three algorithms, where Δt is the time bias and Δs is scale figure based on

Table 25: Comparison of various positioning algorithms, the input vector represents the measured values passed to the algorithm, while the solution vector represents the available solution upon algorithm completion

	TDOA	TOA	STOA
Input Vector	$\Delta p = \begin{bmatrix} r_1 - r_2 \\ r_1 - r_3 \\ r_1 - r_4 \end{bmatrix}$	$\Delta p = \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \end{bmatrix}$	$\Delta p = \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \end{bmatrix}$
Solution Vector	$\Delta x = \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \end{bmatrix}$	$\Delta x = \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \\ \Delta t \end{bmatrix}$	$\Delta x = \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \\ \Delta s \end{bmatrix}$

current frequency offset between tag and receiver. It remains difficult to solve for both Δt and Δs simultaneously due to poor matrix conditioning.

6.1.2 TDOA Algorithm

Algorithm 3: TDOA Algorithm

There are four receivers at known positions R_{x1} or (x_1, y_1, z_1) , R_{x2} or (x_2, y_2, z_2) , R_{x3} or (x_3, y_3, z_3) , and R_{x4} or (x_4, y_4, z_4) , and a tag at unknown position (x_u, y_u, z_u) (Figure 95). The measured distance between four known receivers to the unknown tag can be represented as ρ_1, ρ_2, ρ_3 , and ρ_4 , which is given by

$$\begin{aligned} \rho_i &= \sqrt{(x_i - x_u)^2 + (y_i - y_u)^2 + (z_i - z_u)^2} + ct_u \\ &= f(x_u, y_u, z_u, t_u) \end{aligned} \quad (44)$$

where $i = 1, 2, 3$, and 4 , c is speed of light, and t_u is the unknown time delay in hardware.

The differential distances between four receivers and the tag can be written as

$$\begin{aligned}
\Delta\rho_{1k} &= \rho_1 - \rho_k \\
&= \sqrt{(x_1 - x_u)^2 + (y_1 - y_u)^2 + (z_1 - z_u)^2} \\
&\quad - \sqrt{(x_k - x_u)^2 + (y_k - y_u)^2 + (z_k - z_u)^2}
\end{aligned} \tag{45}$$

where $k = 2, 3$, and 4 , and the time delay t_u in hardware has been cancelled.

Differentiating this equation will give

$$\begin{aligned}
d\Delta\rho_{1k} &= \frac{(x_1 - x_u)dx_u + (y_1 - y_u)dy_u + (z_1 - z_u)dz_u}{\sqrt{(x_1 - x_u)^2 + (y_1 - y_u)^2 + (z_1 - z_u)^2}} \\
&\quad + \frac{(x_k - x_u)dx_u + (y_k - y_u)dy_u + (z_k - z_u)dz_u}{\sqrt{(x_k - x_u)^2 + (y_k - y_u)^2 + (z_k - z_u)^2}} \\
&= \left(\frac{x_1 - x_u}{\rho_1 - c\tau_u} + \frac{x_k - x_u}{\rho_k - c\tau_u} \right) dx_u \\
&\quad + \left(\frac{y_1 - y_u}{\rho_1 - c\tau_u} + \frac{y_k - y_u}{\rho_k - c\tau_u} \right) dy_u \\
&\quad + \left(\frac{z_1 - z_u}{\rho_1 - c\tau_u} + \frac{z_k - z_u}{\rho_k - c\tau_u} \right) dz_u .
\end{aligned} \tag{46}$$

In equation (46), x_u , y_u , and z_u are treated as known values by assuming some initial values




Figure 95: Calculation of tag position based on TDOA approach

for the tag position. dx_u , dy_u , and dz_u are considered as the only unknowns. From the initial tag position the first set of dx_u , dy_u , and dz_u can be calculated.

These values are used to modify the tag position x_u , y_u , and z_u . The updated tag position x_u , y_u , and z_u can be considered again as known quantities. The iterative process continues until the absolute values of dx_u , dy_u , and dz_u are below a certain predetermined threshold given by

$$\varepsilon = \sqrt{dx_u^2 + dy_u^2 + dz_u^2} \quad (47)$$

The final values of x_u , y_u , and z_u are the desired tag position.

The matrix form expression of equation (3) is

$$\begin{bmatrix} d\Delta\rho_{12} \\ d\Delta\rho_{13} \\ d\Delta\rho_{14} \end{bmatrix} = \begin{bmatrix} \alpha_{11} & \alpha_{12} & \alpha_{13} \\ \alpha_{21} & \alpha_{22} & \alpha_{23} \\ \alpha_{31} & \alpha_{32} & \alpha_{33} \end{bmatrix} \begin{bmatrix} dx_u \\ dy_u \\ dz_u \end{bmatrix} \quad (48)$$

where

$$\begin{aligned} \alpha_{k-1,1} &= \frac{x_1 - x_u}{\rho_1 - c\tau_u} + \frac{x_k - x_u}{\rho_k - c\tau_u} \\ \alpha_{k-1,2} &= \frac{y_1 - y_u}{\rho_1 - c\tau_u} + \frac{y_k - y_u}{\rho_k - c\tau_u} \\ \alpha_{k-1,3} &= \frac{z_1 - z_u}{\rho_1 - c\tau_u} + \frac{z_k - z_u}{\rho_k - c\tau_u} \end{aligned} \quad (49)$$

The solution of equation (49) is given by

$$\begin{bmatrix} dx_u \\ dy_u \\ dz_u \end{bmatrix} = \begin{bmatrix} \alpha_{11} & \alpha_{12} & \alpha_{13} \\ \alpha_{21} & \alpha_{22} & \alpha_{23} \\ \alpha_{31} & \alpha_{32} & \alpha_{33} \end{bmatrix}^{-1} \begin{bmatrix} d\Delta\rho_{12} \\ d\Delta\rho_{13} \\ d\Delta\rho_{14} \end{bmatrix} \quad (50)$$

where $[J]^{-1}$ represents the inverse of the α matrix. If there are more than four receivers, the least-squares approach can be applied to find the tag position.

6.1.3 TOA Algorithm

Algorithm 4: TOA Algorithm

1. Again, there are four receivers at known positions R_{x1} , R_{x2} , R_{x3} , and R_{x4} , and a tag at unknown position (x_u, y_u, z_u, t_u) , with an unknown time bias between two clocks t_u . The measured distance between four known receivers to the unknown tag can be represented as ρ_1 , ρ_2 , ρ_3 , and ρ_4 , which is given in the same form as (44).

2. By using an approximate position $(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{t}_u)$, an approximation to the range value can be found

$$\hat{\rho}_i = \sqrt{(x_i - \hat{x}_u)^2 + (y_i - \hat{y}_u)^2 + (z_i - \hat{z}_u)^2} + c\hat{t}_u \quad (51)$$

3. Assuming the position consists of an approximate location and time bias, plus an offset

$$\begin{aligned} x_u &= \hat{x}_u + \Delta x \\ y_u &= \hat{y}_u + \Delta y \\ z_u &= \hat{z}_u + \Delta z \\ t_u &= \hat{t}_u + \Delta t \end{aligned} \quad (52)$$

4. The first order partial derivatives yield

$$\begin{aligned}
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{t}_u)}{\partial \hat{x}_u} &= -\frac{x_i - \hat{x}_u}{\hat{r}_i} \\
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{t}_u)}{\partial \hat{y}_u} &= -\frac{y_i - \hat{y}_u}{\hat{r}_i} \\
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{t}_u)}{\partial \hat{z}_u} &= -\frac{z_i - \hat{z}_u}{\hat{r}_i} \\
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{t}_u)}{\partial \hat{t}_u} &= c
\end{aligned} \tag{53}$$

where

$$\hat{r}_i = \sqrt{(x_i - \hat{x}_u)^2 + (y_i - \hat{y}_u)^2 + (z_i - \hat{z}_u)^2} \tag{54}$$

5. To find the step which minimizes the error with respect to the current location $\hat{\rho}_i$, we define

$$\begin{aligned}
\Delta \rho_i &= \hat{\rho}_i - \rho_i \\
a_{i1} &= \frac{x_i - \hat{x}_u}{\hat{r}_i} \\
a_{i2} &= \frac{y_i - \hat{y}_u}{\hat{r}_i} \\
a_{i3} &= \frac{z_i - \hat{z}_u}{\hat{r}_i}
\end{aligned} \tag{55}$$

6. Thus,

$$\begin{aligned}
\Delta \boldsymbol{\rho} = \begin{bmatrix} \Delta \rho_1 \\ \Delta \rho_2 \\ \Delta \rho_3 \\ \Delta \rho_4 \end{bmatrix} \mathbf{H} &= \begin{bmatrix} a_{11} & a_{12} & a_{13} & 1 \\ a_{21} & a_{22} & a_{23} & 1 \\ a_{31} & a_{32} & a_{33} & 1 \\ a_{41} & a_{42} & a_{34} & 1 \end{bmatrix} \\
\Delta \mathbf{x} &= \begin{bmatrix} \Delta x_u \\ \Delta y_u \\ \Delta z_u \\ \Delta t_u \end{bmatrix}
\end{aligned} \tag{56}$$

describes the linear system

$$\Delta \rho = H \Delta x \quad (57)$$

which can be solved as in (50) for the update step that minimizes the range errors. If there are more than four receivers, the least-squares approach can be applied to find the tag position.

6.1.4 STOA Algorithm

In this novel algorithm, compared with the previous TOA algorithm, the clocks are not assumed to be biased, but rather operate with an unknown frequency bias. As previously discussed, this frequency bias creates the effect of unintentional time scaling. In this derivation, an unknown parameter s will be used to represent the scale, where a scale value of unity will represent the case when the frequency offset is zero.

Algorithm 5: STOA Algorithm

1. Again, there are four receivers at known positions R_{x1} , R_{x2} , R_{x3} , and R_{x4} , and a tag at unknown position (x_u, y_u, z_u, s_u) , with an unknown frequency bias between two clocks represented as a scale factor s_u . The measured distance between four known receivers to the unknown tag can be represented as ρ_1 , ρ_2 , ρ_3 , and ρ_4 , which is given in a similar form as (44) and (58).

2. By using an approximate position $(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{s}_u)$, an approximation to the range value can be found

$$\hat{\rho}_i = \hat{s}_u \sqrt{(x_i - \hat{x}_u)^2 + (y_i - \hat{y}_u)^2 + (z_i - \hat{z}_u)^2} \quad (58)$$

3. Assuming, the position consists of an approximate location and scale factor, plus an offset

$$\begin{aligned}
x_u &= \hat{x}_u + \Delta x \\
y_u &= \hat{y}_u + \Delta y \\
z_u &= \hat{z}_u + \Delta z \\
s_u &= \hat{s}_u + \Delta s
\end{aligned} \tag{59}$$

4. The first order partial derivatives yield

$$\begin{aligned}
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{s}_u)}{\partial \hat{x}_u} &= -\frac{x_i - \hat{x}_u}{\hat{r}_i} \\
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{s}_u)}{\partial \hat{y}_u} &= -\frac{y_i - \hat{y}_u}{\hat{r}_i} \\
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{s}_u)}{\partial \hat{z}_u} &= -\frac{z_i - \hat{z}_u}{\hat{r}_i} \\
\frac{\partial f(\hat{x}_u, \hat{y}_u, \hat{z}_u, \hat{s}_u)}{\partial \hat{s}_u} &= \hat{r}_i
\end{aligned} \tag{60}$$

where

$$\hat{r}_i = \sqrt{(x_i - \hat{x}_u)^2 + (y_i - \hat{y}_u)^2 + (z_i - \hat{z}_u)^2} \tag{61}$$

5. To find the step which minimizes the error with respect to the current location $\hat{\rho}_i$, we define

$$\begin{aligned}
\Delta \rho_i &= \hat{\rho}_i - \rho_i \\
a_{i1} &= \frac{x_i - \hat{x}_u}{\hat{r}_i} \\
a_{i2} &= \frac{y_i - \hat{y}_u}{\hat{r}_i} \\
a_{i3} &= \frac{z_i - \hat{z}_u}{\hat{r}_i} \\
a_{i4} &= -\hat{r}_i
\end{aligned} \tag{62}$$

6. Thus,

$$\Delta \boldsymbol{\rho} = \begin{bmatrix} \Delta \rho_1 \\ \Delta \rho_2 \\ \Delta \rho_3 \\ \Delta \rho_4 \end{bmatrix} \mathbf{H} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{34} & a_{44} \end{bmatrix} \quad (63)$$

$$\Delta \mathbf{x} = \begin{bmatrix} \Delta x_u \\ \Delta y_u \\ \Delta z_u \\ \Delta s_u \end{bmatrix}$$

describes the linear system







$$\Delta \boldsymbol{\rho} = \mathbf{H} \Delta \mathbf{x} \quad (64)$$

which can be solved as in (50) for the update step that minimizes the range errors. If there are more than four receivers, the least-squares approach can be applied to find the tag position. It must be noted that if the ranging values provided to this algorithm contain both a scale factor and a significant time bias, then any bias will be approximated by the scale factor that minimizes the resulting error. As a consequence to this property, this algorithm must be used in coherent systems where any hardware delays are known and fixed. As such this algorithm was not tested in the final non-coherent system testing.

6.1.5 Positioning Algorithms Simulation Results

Using several realistic configurations of base stations, each algorithm was tested with $\pm 11\%$ scaling and 1D jitter values with a standard deviation 10.0 mm; this was repeated 25 times for each of 84 tag positions. The system was assumed to be previously calibrated for cable length offset. Since the STOA algorithm was able to accommodate a changing scale figure in the provided solution, it provided strictly lower error values for each configuration (Table 26).

Table 26: Comparison of simulation error results across differing positioning algorithms when 11.0% scaling error is present using various base station configurations. *indicates the algorithm failed to converge

	Configuration (# of Base Stations)	TDOA (mm)	TOA (mm)	STOA (mm)
	Flat 9	136.1	132.1	0.833
	Curved 9	4.23	4.12	0.569
	Flat 7	124.3	123.7	0.920
	Curved 7	4.47	4.37	0.765
	Flat 5	143.9*	144.7*	3.83
	Curved 5	5.64	5.53	2.11

Repeating the experiment above using the same base station configurations except a 1.1% scaling error was applied to each raw 1D range (Table 27). In this case, the STOA algorithm did not perform as well. The TOA algorithm performed generally better than the STOA algorithm with lower number of base stations in the curved configuration. In the flat configuration, the STOA algorithm provided superior results.

The difference in PDOP values between these configurations is apparently a large source of the error in these experiments, as can be seen in Figure 96. This can be directly seen in Table 27, as the large factor of error in the TDOA column between the Flat 9 configuration and the Curved 9 configuration.

6.2. Kalman Filtering

Kalman filtering is widely used in GPS [62] and positioning systems [45] since it offers an attractive dynamical model which can change filter weights based on system characteristics. In the Microwave Positioning Subsystem, there are two distinct ways with which to incorporate a Kalman filter. The first option is to use the Kalman filter to filter the individual range values, while the second is to Kalman filter the 3D position as returned by the positioning algorithm. Preliminary investigation has found roughly equal accuracy performance from these two options,

Table 27: Comparison of simulation error results across differing positioning algorithms when 1.1% scaling error is present using various base station configurations.

Configuration (# of Base Stations)	TDOA (mm)	TOA (mm)	STOA (mm)
Flat 9	4.54	3.82	0.789
Curved 9	0.649	0.540	0.600
Flat 7	4.23	3.88	0.890
Curved 7	0.705	0.634	0.703
Flat 5	11.84	11.52	3.47
Curved 5	1.92	1.85	2.12




Figure 96: For the TDOA algorithm, the degradation in PDOP due to geometric differences between the Flat 9 configuration and the Curved 9 configuration

with the later representing less computation time. Since the Kalman filter requires two matrix inversion steps, the Kalman filter operating on a 3D position represents a more compact representation than in the case of the Kalman filter operating on the raw range values.

In preliminary testing the Kalman filter was applied to raw 3D position values returned by the TDOA algorithm when the tag was maintained in a static position. The Kalman filter assumed a propagation model of constant velocity and filter error estimates were set using the 3D variances of the raw data. Figure 97 shows the raw system error, with extreme ranging errors of greater than 10^4 mm. After the Kalman filter was applied to the same data it is clear in Figure 98 that the Kalman filter has excellent performance in filtering extrema. Upon filter stabilization, around sample 700, the error remains on the order of 10-12 mm.

In this example the Kalman weights were set with an overly pessimistic view of the data, with the measurement noise covariance matrix set with such extreme values, effectively only static points will be able to be acquired due to lengthened acquisition time (e.g. 700 samples in the




Figure 97: Raw static positioning error using TDOA algorithm and a cable length calibration




Figure 98: A Kalman filtered result of the exact data contained in Figure 97, note that after the filter stabilizes around sample 700 error approaches 10-12 mm

previous example). With two techniques presented subsequently, the Kalman filter measurement noise covariance may be set lower with the added benefits of protecting the Kalman filter from entering erroneous points.

6.3. Protected Kalman Filtering

Kalman filtering can handle error-prone points, but extreme outliers can cause an over-estimate of the process error covariance matrix, especially in high dynamic situations. A new approach is to discard (or pre-filter) outliers using the Mahalanobis distance, or other distance metric, before the point is fed to Kalman filter

$$D_M = \sqrt{(y_t - x_{t-1})^T R^{-1} (y_t - x_{t-1})} \quad (65)$$

$$\text{discard}(y_t) \quad \text{on } [6 \leq D_m < \infty)$$

where y_t is an observed 3D point calculated by TDOA, TOA, etc. and x_{t-1} is the previous state of a Kalman filter, R can be the measurement error covariance or can be reweighted based on geometry, and D_M is the Mahalanobis distance. With D_M set at 6, Figure 99 shows elimination of the largest spikes as would be expected using this metric. If D_M were set even lower (i.e. 3), several of the remaining large spikes would be removed, but dynamic motion would be restricted.

6.3.4 Position Dilution of Precision Combined with Kalman Filtering

The topic of Position Dilution of Precision (PDOP) or more generally known as Geometric Dilution of Precision (GDOP) is widely analyzed in the field of GPS [62]. A discussion of PDOP and its relation to this Microwave Positioning Subsystem is presented as follows. The use of PDOP values in conjunction with Kalman filtering is a topic that has not been completely




Figure 99: Dynamic simulation comparing Kalman Filtering 3D position data (red) with Kalman Filtering with Mahalanobis Protection (blue). Note that the largest spikes have been eliminated.

investigated. A previous method weighted Kalman filter output from subsets of GPS satellite configurations by the inverse of the configuration's PDOP value to yield a combined solution [142]. A novel approach is to use the PDOP as a geometry-specific method to pre-filter points entering into the Kalman filter. This method utilizes PDOP values to measure the impact of base station constellation geometry on final system error given 1-D ranging errors. There are two main ways this can be accomplished. The first is to use element-wise reweighting of the base measurement error covariance by the current scalar PDOP value as

$$R_w = R_0 \cdot PDOP \quad (66)$$

where R_0 represents the initial or base error measurement covariance matrix and R_w is the re-weighted version which would be used to filter this particular 3D point. Secondly, the measurement error covariance is set using the appropriate components of the inverse of the TDOA or TOA error covariance matrix D as

$$R_w = D\sigma_{1D}^2 = (H^T H)^{-1} \sigma_{1D}^2 \quad (67)$$

where σ_{1D}^2 is an initial estimate or a system maintained estimate of the 1D ranging variance. Since this representation carries a relatively large amount of information regarding the predicted error covariance for a particular geometry, this representation is used for the primary testing of this technique. Essentially, (66) is a scaled approximation to (67) using the square root of the trace of the D matrix (as in (23)). As shown in Figure 100, a dynamic simulation showed that D matrix protection in the form of (67) significantly reduced error versus an unprotected Kalman filter. This is intuitive when considering that in areas where the PDOP values of one coordinate axis degrade severely, the Kalman filter can selectively filter these values with a lower degree of sensitivity.

6.4. Non-Linear Filtering

Upon further inspection of the raw time indices determined by the Max-Ratio algorithm in




Figure 100: Dynamic simulation comparing Kalman Filtering 3D position data (red) with Kalman Filtering with D-Matrix Protection (blue), note that the RMS error has been significantly reduced

experimental testing, it was found that the noise characteristics of the data were distinctly non-Gaussian. As can be visualized in Figure 101, the ripples in the I-only channel data are caused when non-coherent LOs are used for up-conversion and down-conversion of the transmitted pulse. Despite the fact that the Max-Ratio algorithm dynamically calculates a threshold based on 25% of the maximum height of the pulse, this threshold is still sensitive to the non-linearities of the leading edge of the pulse. As a consequence, the computed leading edge location can oscillate between two rather stable positions on the pulse.

As displayed in Figure 102, the raw data (blue) originating from the Max-Ratio algorithm maintains fairly consistent bounds. However, the data oscillates primarily between two values (-225 and -295), corresponding to an overall swing of roughly 35 mm, which is unacceptable given the target system accuracy. When a windowed box filter was applied to the data (green), the filtered data exhibited less variance but exhibited short term wander which is undesirable. A windowed median filter, which is a simple rank order filter, was also applied to the data,




Figure 101: Synthesized I^2 channel data corrupted by transmitter-receiver 8 GHz LO frequency offset

shown in dark red. Finally, by plotting the cumulative distribution of the same raw data (Figure 103), it is clear that the stable regions of the distribution are approximately located at the first and third quartiles. As such, a windowed Mid Inner Quartile Range (MIQR) filter was devised to remove the effects of this consistent oscillation. The results of the MIQR filter are shown in Figure 102 in yellow and exhibit the highest level of stability and correspondingly the lowest amount of variance.

From the previous work in this chapter, it is noted that the TDOA algorithm solves sets of non-linear equations. Given this, another hypothetical filtering technique was tested termed Multi-Base filtering. In this technique, one or more range values originating from the same base station were collected. Using these extra ranging values the TDOA algorithm solution (50) was solved using the Moore-Penrose pseudo-inverse, where redundant base station positions were considered with the duplicate ranging data. Table 28 and Table 29 show the 3D RMS errors and 3D 95% confidence intervals, respectively, for these types of non-linear and linear filtering




Figure 102: Effect of several types of windowed TOA index filters on the raw data (blue), also represented in Figure 103, each of the windows are size 16, the MIQR filter (yellow) has the lowest variance of the three types

techniques when used with the TDOA algorithm. In this case, the Multi-Base filtering technique, the window size refers to the number of ranging value assigned to each unique base station position. It should be highlighted that the two filtering techniques with the lowest RMS error and smallest 95% confidence interval are the averaging filter and the MIQR filter.




Figure 103: Cumulative distribution of pulse TOA index values showing non-Gaussian structure

Table 28: 3D Positioning RMS errors in mm for various index filtering techniques used in conjunction with the Max-Ratio Algorithm

Window Size	Raw	Median	MIQR	Ave	Multi-Base
1	16.018	16.018	16.018	16.018	16.018
5	16.018	7.781	6.973	7.508	12.036
9	16.018	7.506	4.162	5.434	11.667
13	16.018	7.113	2.888	4.298	11.05
17	16.018	6.65	2.572	3.713	10.517
21	16.018	6.769	2.434	3.382	10.509
25	16.018	6.479	2.233	3.108	10.66
29	16.018	6.015	2.047	2.86	10.563

Table 29: 3D Positioning 95% C.I. in mm for various index filtering techniques used in conjunction with the Max-Ratio

Window Size	Raw	Median	MIQR	Ave	Multi- Base
1	15.295	15.295	15.295	15.295	15.295
5	15.295	11.562	8.62	8.31	19.009
9	15.295	10.696	6.914	6.029	19.291
13	15.295	10.338	4.919	4.601	18.691
17	15.295	9.777	4.334	4.254	17.943
21	15.295	9.543	4.229	4.143	18.474
25	15.295	9.416	3.797	4.855	18.49
29	15.295	8.757	3.307	4.928	18.219

7. SYSTEM CALIBRATION AND VERIFICATION

This chapter focuses on the challenges of calibrating systematic error in order to produce results which not only offer the highest level of precision, but also ensure that the total system error relative to a reference optical tracking system is minimized. In a rudimentary fashion, a primary contributor to the final system error is the unknown lengths of the cables that connect each base station to the control station. These cables are of fixed length and any length differences between the electrical lengths of base stations directly corrupt TDOA values. A second contributor of systematic error is that of nominal PRF clock frequency bias and nominal ADC clock values differing from system calculated values. In this case, a calibration algorithm is developed which simultaneously solves for small scale factor differences as well as cable length offsets. This algorithm also relies on a known positioning reference. Supposing that a rudimentary calibration based on one of these techniques has been performed during system fabrication, a final algorithm is simulated to demonstrate the feasibility of a final installed system calibration which does not require a positioning reference. This final algorithm can solve for small bias values as well as base station orientation, which contributes to non-linear errors with respect to tag trajectories.

7.1. System Cable Length Calibration

The system requires a cable length calibration step due to electrical length differences that cause constant bias in measured time difference (Figure 104). To perform positioning once the calibration step has been completed, the biases are subtracted from measurements using offset values stored in the system software. These bias values can be inferred from those offsets which minimize error when a UWB tag is in a known and static location.




Figure 104: Diagram showing cable length offsets and consequent need for system calibration

7.1.4 System Cable Length Calibration Algorithm

A step-wise error minimizing calibration algorithm is shown in Algorithm 6. This algorithm works by minimizing system error when the tag is in a known location. Since the cable length problem affects each positioning algorithm, each positioning algorithm requires an offset calibration step.

Algorithm 6: Remove Cable Length Bias from System Measurements

1. Initialize Offset:

$$O := [d_{01} , d_{02} , d_{03}] \quad (68)$$

2. Measure System Error against true value:

$$E_{SYS} := P_{uwb} - P_{true} \quad (69)$$

3. Sum the squared system error vector over 20 iterations, and take the mean value:

$$MSE = \sum (E_{SYS} .^2) / 20 \quad (70)$$

4. Adjust the offsets in each direction given step size:

$$O_{new} := [d_{01} \pm d_s , d_{02} \pm d_s , d_{03} \pm d_s] \quad (71)$$

5. Record O_{new} with lowest E_{SYS} then alter $D_s = 0.9 * D_s$.

7.1.5 System Cable Length Calibration Results

The results of the calibrated system are shown in Figure 105. In the best case, the calibration algorithm reduced the system error for the TDOA algorithm from roughly 1000 mm to 54.95 mm in 200 iterations and an initial step size of 150 mm. When a 10-point averaging filter was applied to the points, the error was further reduced to 11.92 mm.

The calibrated system was used in combination with a 1000-point averaging filter. The results shown in Figure 106 highlight the system drift that is present due to clock frequency bias (between asynchronous 10 MHz clocks), which causes time-scaling. It should be noted that for this experiment α was set to 100,000, thus a $\pm 11\%$ scale factor error is the likely contributor of this systematic error.

The potential downside of this algorithm is that while it effectively can calibrate offsets if the tag is fixed, it cannot be used if the tag is in dynamic motion or if the location of the tag is unknown. Additionally, if the nominal scale factor is significantly different than the system expectation, then any scale factor differences get masked by the calibrated offsets. As the clocks drift relative to each other, the nominal scale factor will change and thus be evidenced.




Figure 105: Experimental results showing the error trend of sampled data positioned using the cable length calibration algorithm combined with the TDOA algorithm with no filtering




Figure 106: Experimental results showing the error trend of sampled data positioned using the TDOA algorithm with a 1000-point averaging filter, and alpha =100,000 scale factor, noting the possibility for $\pm 11\%$ scale factor error

7.2. System Time-Scale and Offset Calibration

As noted above in Figure 106, the system may have some long-term drift associated with a frequency bias between the 10 MHz PRF clock located on the tag and the 10 MHz clock PRF source that feeds each of the base stations. Additionally, a 50 MHz clock drives each of the FPGAs and associated ADCs. Consequently, due to the unintentional time scaling phenomenon described in Chapter 3, the period of the sample spacing as interpreted by the positioning software may be 2-3% different than the calculated nominal value when $\alpha = 10,000$. In this form of calibration, not only is any nominal offset between the base stations removed, but also a scale factor is calculated for each of the received time differences. In contrast to the cable length calibration algorithm, this algorithm cannot be enacted at a solitary point, but rather requires a diversity of points be taken throughout the working volume for the most accurate results. Similarly, however, a positioning reference is needed to record the points taken with a high degree of accuracy. To compute this transform using a least squares approach, first

$$\Delta r_{1k}^t = r_1^t - r_{k+1}^t \quad (72)$$

$$\Delta \rho_{1k}^t = \rho_1^t - \rho_{k+1}^t \quad (73)$$

where Δr_{1k}^t is the range differences computed by the reference positioning system and $\Delta \rho_{1k}^t$ are the range differences computed by the UWB system as in (45) and t represents the time index of the acquired ranges. Next the transformation H is solved for in the linear system,

$$\Delta P = H \Delta R \quad (74)$$

$$\Delta P = \begin{bmatrix} \Delta \rho_{11}^1 & \cdots & \Delta \rho_{1k}^1 & 1 \\ \vdots & \ddots & \vdots & 1 \\ \Delta \rho_{11}^N & \cdots & \Delta \rho_{1k}^N & 1 \end{bmatrix} \quad (75)$$

$$\Delta R = \begin{bmatrix} \Delta r_{11}^1 & \cdots & \Delta r_{1k}^1 \\ \vdots & \ddots & \vdots \\ \Delta r_{11}^N & \cdots & \Delta r_{1k}^N \end{bmatrix} \quad (76)$$

$$H = \begin{bmatrix} s_{11} & s_{12} & s_{13} \\ s_{21} & s_{22} & s_{23} \\ s_{31} & s_{32} & s_{33} \\ d_1 & d_2 & d_3 \end{bmatrix} \quad (77)$$

In general this technique is realistic if a reference system can be used to perform the calibration. Since a wide diversity of points must be used to ensure that ΔP and ΔR are full rank, this calibration technique is termed dynamic calibration, as it requires movement of the tag and reference positioning system.

7.3. Base Station Calibration Algorithm using Arbitrary Path Motion

A critical step that must be undertaken for final system error to yield mm or sub-mm range accuracy is the accurate calibration of the base station orientation. The orientation is important due to the Vivaldi base station antenna sensitivity to phase-center (and thus positioning) variation with respect to the angle-of-arrival (AOA). As Figure 107 shows, as the AOA varies from 0° to 45° the phase center error variation is roughly 10 mm. A base station calibration

algorithm is developed in Section 7.3.5.3 which uses an arbitrary path motion to calibrate base station orientation if the base station position is known.

7.3.4 Introduction

A method has been proposed to correct three-dimensional (3D) positioning error due to base station antenna orientation that takes into account non-boresight electrical length differences due to antenna phase center errors. An automated algorithm is used to calibrate an ultra-wideband (UWB) system using only a starting estimate of the base station position and acquired positions central to the base stations. The true positions of the acquired 3D calibration points are unknown to the calibration algorithm. Upon completion of the algorithm, the base station orientation is estimated, along with estimates of electrical length offsets due to potential cable length differences. This method is designed to minimize small errors due to base station position and orientation uncertainty. The algorithm is shown to be robust given that the availability of accurate 1D ranging can be provided by the system.




Figure 107: Measured Vivaldi phase center error versus angle for E-cut, H-cut, average of E and H cuts, and centered quadratic fit

In designing and constructing a high accuracy ($\sim 1\text{mm}$) UWB positioning system [42] [76], several types of errors and biases that are safely ignored on the decimeter scale become significant to overall system performance. One such effect is that of phase center errors originating from transmitter-receiver mal-alignment. These errors can be successfully removed using the TDOA algorithm with knowledge of the receiving antenna orientations. However, ignoring this type of error proves impossible for the TDOA algorithm to remove. Even in a trial case involving 100 base stations, phase center errors are correlated and create considerable system error. Thus, the accuracy of the TDOA algorithm is highly dependent on the accuracy of the differences in the TOA of the UWB pulses.

Using an iterative approach, an algorithm has been developed and tested under simulated line-of-sight (LOS) conditions where ranging error was modeled as a zero mean Gaussian random variable with a variance equal to the amount of ranging error present in the UWB system.

Currently, commercial UWB positioning systems of Sapphire DART (Multispectral Solutions, Inc.) and Ubisense have indoor positioning accuracy of 10 cm and 15 cm, respectively [70]. Interesting results presented by Zetik et al. and Meier et al. indicate that ranging has the potential to achieve mm or even sub-mm accuracy levels [72] [45]. Recent work in UWB positioning algorithms has focused on increasing the accuracy of TOA measurements in Non-Line-of-Sight (NLOS) conditions [143], increasing TOA resolution [144], and increasing TOA accuracy in the presence of interference [108] .

7.3.5 Materials and Methods

The purpose of the experiments described previously was to expose and eliminate the error due to antenna phase center and simultaneously positioning error due to the antenna bore sight

direction uncertainty. First, a short discussion of the phase center error is included for relevance to this calibration procedure.

7.3.5.1 Antenna Phase Center Error

The single element Vivaldi antenna is used on the receiver side of the system. The design of the antenna is shown in Figure 108. The z-axis points in the direction of the “end fire” radiation pattern typical of antipodal antennas. This particular design of antenna is particularly sensitive to phase center variation in both the E-plane (parallel with the plane of the antenna, xz-plane in Figure 108) and H-plane of the antenna radiation pattern as shown in Figure 107 [42]. The average error between the E-cut and H-cut are used to approximate the error as a function of any angle of incidence relative to the boresight direction. A quadratic approximation was chosen and is of the form

$$\Phi = E(\theta) = 0.00000364(\theta)^2 \quad (78)$$

where θ is calculated as the full angle in degrees between the angle of incidence and the boresight direction and Φ is positioning error in meters.

7.3.5.2 Orientation Error Simulation

An experiment was undertaken to discover the effects of the phase center error on the TDOA algorithm when combined with a large amount of base stations. In this case, it was hypothesized that the redundancy in the large number of base stations would tend to cancel out the phase center errors.

In this simulation a virtual scene is created with $N=100$ base stations randomly distributed on a sphere of radius $R=3$ m. In this experiment base stations are assumed to all be oriented directly




Figure 108: Designed Vivaldi antenna, with the +z axis being the boresight direction

toward the center of the sphere, which is the origin of the coordinate system (0,0,0). The tag travels in a helical path between x-axis values of $\pm R/2$ according to

$$P_i^T = (x_i, r_i \sin(w_H x_i), r_i \cos(w_H x_i)) \quad (79)$$

$$r_i = \min\left(\frac{x_i}{2}, r_{MAX}\right) \quad (80)$$

where x_i are $M=81$ uniformly sampled points between $x=\pm R/2$. A particular instance of this recreated scene is shown in Figure 109. Given the tag and base stations positions, the true 1D range is calculated for each tag position relative to each base station. Also, the true angle of incidence is known and calculated for each tag position-base station pair. Synthetic range values are thus calculated as

$$d_{ij}^{TBS} = \|P_i^T - P_j^{BS}\| + \Phi_{ij} + b_j + n_{ij} \quad (81)$$

where $1 \leq i \leq M$, $1 \leq j \leq N$ and P_j^{BS} represents the position of the j^{th} base station, Φ_{ij} is the phase center error between the i^{th} tag position and j^{th} base station. A Gauss-distributed random bias term b_j ($\mu=0$, $\sigma=2.5$ mm) is added to the range to account for potential electrical length differences, and the slight uncertainty in measuring the base station position phase center along the boresight direction. The nominal value for the standard deviation was chosen to be an order




Figure 109: Instance of the 100 base station experiment

of magnitude higher than the uncertainty of 3D localization. Finally, a random noise figure of n_{ij} ($\mu=0$, $\sigma=1.24$ mm) is added to each range to account for uncertainty in 1D range measurements. This nominal value was chosen based on the actual 1D measurement RMS error of the system [76]. In this simulation, the standard TDOA algorithm is applied to the synthetic ranges generated in (81), and the error of the calculated position vs. the true position as a function of the x_i position of P_i^T is shown in Figure 110.

According to Figure 110 and Figure 111, redundancy in the number of base stations is not sufficient to remove phase center error. This can also be understood by considering that the geometry at both extremes of the tag path represent cases where many of the base stations will simultaneously have significant phase center error that creates a positive length bias that the TDOA algorithm cannot remove.

7.3.5.3 Calibration Algorithm

Considering the same geometry as described in the previous simulation, it is desirable to calibrate the system to accurately determine the orientation of the base stations only through




Figure 110: 3D Positioning Error vs. Tag x Position for 100 base station experiment illustrating the difference between Φ included (dashed) and not included (solid) in equation (4)




Figure 111: Base station specific results of errors shown in Figure 110, note that base stations oriented along the tag track experience the least amount of error due to phase center variation

measured data. Assuming that the base stations are again oriented roughly toward the center of the work area and that the base stations positions are known, the proposed algorithm will estimate both the orientation and boresight bias (b_j) according to the error model given in equation (90). Also, to perform this calibration there must be $N > 4$ base stations, as one base station must be left out. This leaves a minimum of 4 base stations to provide unambiguous hyperbolic 3D positioning.

Algorithm 7: Phase Center TDOA Calibration

1. Using the raw 1D ranging values, use the TDOA algorithm on $N-1$ base stations (leaving one base station out of the calculations). Iterate so that each base station is left out in turn. This is performed at all tag positions. The result is MN position estimates $\hat{P}_{ij}^T$, for the i^{th} position corresponding to the j^{th} base station that was removed. All of the base stations are given an initial orientation estimate $\hat{O}_j^{BS}$, which is a random unit vector, representing the boresight direction.
2. Each position estimate $\hat{P}_{ij}^T$ is used to calculate the angle θ_{ij} between the tag position-base station and the current orientation estimate $\hat{O}_j^{BS}$.
- 3a. The true 1D ranging values $\hat{R}_{ij}$ are estimated based on the calculated θ_{ij}

$$\hat{R}_{ij} = R_{ij} - E(\theta_{ij}) - \hat{b}_j \quad (82)$$

where R_{ij} is the measured range and $\hat{b}_j$ is the current estimate of the boresight bias (initially set to zero).

- 3b. Alternatively, the error due to the phase center $\hat{\Phi}_{ij}^{TBS}$ is estimated without consideration of the current orientation

$$\hat{\Phi}_{ij}^{TBS} = R_{ij} - \|\hat{P}_i^T - P_j^{BS}\| - \hat{b}_j \quad (83)$$

and subsequently,

$$\hat{\theta}_{ij} = \begin{cases} E^{-1}(\hat{\Phi}_{ij}^{TBS}) & \hat{\Phi}_{ij}^{TBS} > 0 \\ 0 & \hat{\Phi}_{ij}^{TBS} \leq 0 \end{cases} \quad (84)$$

noting that this step can be performed independently of step 3a (diagrammatically shown in Figure 112).

4. The estimated ranging values $\hat{R}_{ij}$ are then used in the standard TDOA algorithm to re-estimate position values $\tilde{P}_i^T$ using all N base stations in the calculation.

5. The collection of position estimates $\tilde{P}_i^T$ is then used to update the boresight bias estimate

$$\hat{b}_j = \hat{b}_j^{old} + \frac{1}{M} \sum_{i=1}^M (\|\tilde{P}_i^T - P_j^{BS}\| - \hat{R}_{ij}) \quad (85)$$

which can be considered as an update to the existing bias estimate using the mean ranging error across all measured positions.

6. Next the orientation estimate $\hat{O}_j^{BS}$ is updated according to:

$$K_j = \min \arg_i \left(\sum_{k=1}^N \hat{\theta}_{ik} \right) k \neq j \quad (86)$$

$$\hat{O}_j^{BS} = \frac{(P_{K_j}^T - P_j^{BS})}{\|P_{K_j}^T - P_j^{BS}\|} \quad (87)$$

thus, the final step in this iterative algorithm is to assign the base station orientations according to the points along the point track where the $N-1$ base station positions yield the lowest amount

of angular error. Under the assumption that the base stations are oriented roughly towards the center of the work area, this condition yields algorithm convergence.

7. The algorithm converges when the relative RMS error between position estimates falls below a threshold.

7.3.5.4 Calibration Simulations

Several simulations under realistic system configurations were used to verify the accuracy and convergence characteristics of Algorithm 7, for calibrating the base station orientations and electrical length differences. These trials were enacted under a varied number of base stations randomly distributed on a sphere. Several fixed parameters set in this study were: the minimum resolvable distance (0.1 mm), the noise in 1D measurement (1.24 mm), and the number of trials




Figure 112: Visual depiction of Step 3b of Algorithm 7

for each configuration (20 trials). The 1D ranging noise was taken from observed 1D ranging noise in [76].

7.3.6 Results

The algorithm was monitored and averaged across the 20 trials to give a relative indicator of success. Two such trials are shown in Figure 113. In the first case which represented only 6 base stations, there was not a sufficient level of redundancy for the algorithm to converge to sub-mm range calibrated error. This is due in part to occasional dilutions of geometric precision which causes instability in the algorithm to ascertain correct ranging values. While on average the $N=6$ case did not converge to sub-mm levels of accuracy, most of the runs did converge with the mean being negatively influenced by a few cases of non-convergence.

From Figure 113 and Figure 114 it is clear that the $N=10$ base station case represented enough system redundancy to drastically reduce noisy 1D ranging errors (1.24 mm) with the limit being the (0.1 mm RMS or 0.7 mm max error) minimum resolvable distance being preserved.




Figure 113: Comparison of $N=6$ base stations and $N=10$ base stations on the calibration algorithm convergence of RMS error. Note algorithm instability for $N=6$ case. Sub-mm calibrated accuracy is possible with $N=10$ base stations




Figure 114: Comparison of N=6 base stations and N=10 base stations on the calibration algorithm convergence of maximum error. Note algorithm instability for N=6 case. Sub-mm calibrated accuracy is possible with N=10 base stations

7.3.7 Discussion

The system calibration procedure described in this paper is necessary to push the limits of UWB positioning beyond the current levels of commercial accuracy. The algorithm was tested in simulated LOS conditions where measurement error is well-known to be minimal. Further work would involve testing the algorithm in situations of non-line-of-sight (NLOS) conditions, as well as situations where the base station orientations cannot be assumed to be oriented toward the center of the work area. While much work in UWB positioning has been performed in driving raw 1D TOA accuracy, with many associated algorithms being developed in the literature, the author's wish is to motivate a careful look at antenna phase center effects on overall system accuracy. The algorithm presented here describes a calibration procedure that can be translated into TDOA based positioning systems where AOA effects result in system degradation of accuracy. In such systems, the orientation of the base station relative to the mobile tag can be inferred as well as biases resulting in fixed electrical length differences. Such differences could

arise if different lengths of cable are used to connect the base stations, or if the base stations retransmit pulse arrival times. The prime feature of the proposed algorithm is that arbitrarily captured points can be used to calibrate the system, which indicates that this algorithm may be able to be implemented as a continuous online process or to be purely performed offline. The drawback to this technique would be the relatively large degree of redundancy that must be present in the system.

8. SYSTEM CHARACTERISTICS AND ERROR ANALYSIS

When considering the vast number of parameters, settings, and techniques inherent to the UWB system, experimental testing can assist in determining optimal system configuration under certain operating conditions. The goal of the experimental testing is to highlight the tradeoffs that exist in building a real-time positioning system and offer guidance on various configurations where acceptable performance might be found supposing different system requirements. For example, one large system tradeoff is the accuracy of positioned point versus the frequency at which points may be acquired. If the accuracy requirements of the end-user are minimal, then point acquisition frequency can be increased. Another tradeoff is accuracy versus whether the system must be tethered to the control station (i.e. wired or wireless). Another tradeoff is that of system accuracy versus the number of base stations that share a coherent LO. In all of these tests, the primary motivating goal is RMS error when compared to a reference optical system (Optotrak) which has an error of roughly 0.1 mm RMS in-plane and 0.3 mm RMS out-of-plane errors. While the Optotrak errors can be considered to be uncorrelated with the UWB errors, the additive effect of the Optotrak system has been deemed negligible, and therefore is not subtracted from the calculated system accuracy of the UWB system.

All of the tests shown in this section have roughly similar PDOP values for the coordinate axes namely: 2.3, 0.9, and 0.6, for the coordinate specific $PDOP_x$, $PDOP_y$, and $PDOP_z$, respectively. This represents an overall PDOP value of 1.95. Since the base station configuration remains the same, the error values shown below can be directly compared, although, if more base stations were acquired, these PDOP values would be able to be decreased in practice. Since the overall

PDOP value describes the error variance the following equation can be used to adjust the error values presented in the next section

$$E_{RMS}^{adj} = E_{RMS}^{orig} \frac{PDOP_{adj}}{PDOP_{orig}} \quad (88)$$

where E_{RMS}^{adj} is the resultant 3D RMS error after adjustment, E_{RMS}^{orig} is the original 3D RMS error, $PDOP_{adj}$ is the resultant scalar PDOP value resulting from a different base station configuration, and $PDOP_{orig}$ is the original configuration with a nominal value of 1.95. In an eight base station configuration, for example, the PDOP value was found to be below 0.95 for all positions within a 1 m cube. In this example, the error values would be reduced by more than a factor of two.

8.1. System testing

One interesting test of the system was that of testing the scale factor α , which corresponds inversely to the analog sample spacing. When the typical value of $\alpha=10,000$ is used, the corresponding sample spacing is 10 ps, and 1 ms is required to acquire one set of pulses from the system. A second test was performed when $\alpha=40,000$, corresponding to an analog sample spacing of 2.5 ps. At this scale factor, the system requires 4 ms to acquire a set of pulses to produce a single 3D position measurement. As shown below in Figure 115, the RMS error can be decreased by almost 1.5 mm when using a scale factor of 40,000.

Another experimental set up examined the positioning error difference between a wired and a wireless setup. These error values are depicted in Figure 116, indicating the robustness and high-precision that is achievable in a wired system. The value of 0.73 mm in the wired case is comparable to the system presented by Meier et al. with a stated static accuracy of 0.1 mm [45],




Figure 115: Comparison of different scale factors: 10000 (red) and 40000 (blue) on 3D RMS error, in a non-coherent, MIQR=17, 5x point averaging, using TDOA algorithm




Figure 116: Comparison of wireless and wired PRF (10MHz) clocks: wireless (red) and wired (blue) on 3D RMS error, static, MIQR=41, 5x point averaging, using TDOA algorithm

although their result degrades to 2.0 mm in dynamic tracking. In the wireless case, 5.59 mm of RMS error appears to be the lowest error found in the published literature.

One of the key factors in decreasing the error to this level was the connection of the 4 base stations to the same 8GHz LO on the control station. Note that this is a realistic case as the base stations are already wired to transmit the pulse positions to the control station. Prior to this experiment, the clock stability of independent LOs, one per base station, was not thought to affect system performance. As Figure 117 clearly shows, drastic error improvement is found with both 2 coherent base stations and also markedly with 4 coherent base stations.

When the raw 3D point data similar to the data from Figure 117, in the non-coherence case, is passed through varied types of Kalman filters, the RMS error values decrease as visualized in Figure 118. While a conventional Kalman filter reduces the error from 65.75 mm to 54.79 mm, the developed techniques of D matrix protection and Mahalanobis protection further reduce the




Figure 117: Raw data captured using wired PRF clocks and varying levels of LO (8GHz) coherence on each of the 4 base stations




Figure 118: 3D RMS error values resulting from application of various 3D Kalman filtering techniques with non-coherent base stations LOs, compare with Figure 117

error to 44.03 mm and 19.15 mm, respectively. This last technique, Mahalanobis represents more than a three-fold reduction in error against the raw data.

Since coherent base station LOs offer a convincingly lower degree of error, the next experiment (Figure 119) was enacted to test median and MIQR index filtering for a fixed window size of 17. This experiment exhibits the same trend as Table 28 and Table 29 also with the less distinct advantage of MIQR filtering over Median filtering. Considering that the MIQR filter and the Median filter are algorithmically $O(1)$, constant-time algorithms with respect to the window size, larger window sizes would be able to be used without incurring additional computational overhead. In general, the time constant of the Median filter is lower than that of the MIQR filter. Considering metal interference, an experiment was developed to measure the system performance in the presence of non-occluding metal interferers (two metal bars) placed near the transmitter. As is shown in Figure 120, metal interferers placed close (1 cm) to the transmitter




Figure 119: 3D RMS error values resulting from index filtered data captured using wireless PRF clocks, coherent base stations, and the TDOA algorithm




Figure 120: 3D RMS error values resulting data captured dynamically in the presence of metal interferers using wireless PRF clocks, coherent base stations, MIQR=17, ave=5, and the TDOA algorithm

cause large spikes in error values primarily due to multipath interaction. When the metal interferer was removed to a distance of approximately 8 cm away from the transmitter, the error bounds in terms of RMS and maximum error decreased.

For many applications, the goal with any positioning system is to achieve a low rate of error while the tracked device is in free-form motion (i.e. not a prescribed path). One criticism of the work of Meier et al. stems from their reporting of error values originating from the fact that the testing was performed on a calibrated track, which can provide constant velocity motion. Using a first-order Kalman filter in this case can obscure true system performance under varied dynamic conditions. In another experiment, shown in Figure 121, the tag was held by an observer and waved in several circular patterns with different orientations. The final error for this experiment was 6.95 mm RMS, which demonstrates system tolerance to dynamic environments.




Figure 121: Free-form tag motion, comparing the 3D tracked points of the Optotrak (red) and UWB systems representing 6.95 mm RMS error, IQR=17, no averaging, wireless, non-coherent

Following this result, a more complete set of experiments were undertaken to study system hysteresis, linearity, sensitivity, and repeatability errors [146]. One of these experiments seen in Figure 122 in three orthogonal views is the optical and UWB tags traversing down a track with a second DOF (the tag can translate laterally slightly down the track).

From the side, the hysteresis error is visualized the most succinctly as in Figure 123 in two trials. Here the hysteresis error between 3.3-4.7% is determined by the maximum 3D error between positive travel and negative travel normalized by the track length.

Considering linearity, Figure 124 shows a robotic traversal of three linear tracks, with a 3 second pause at 20 distinct points. These three linear tracks were adjusted and a least-squares fit line was passed through the Optotrak points and the UWB points, and the maximum error between these two lines was recorded as less than 6 mm over the track length of 650 mm. Thus, the linearity error is calculated as less than one percent.

Using the same data previously acquired, only the static points (velocity less than 15 mm/s) were considered (Figure 125). The point positions were measured twice in this configuration, and it was determined that the static repeatability error was 7.00 mm. Again, the linearity is quite low when considering only the static points as the robot moves in slightly curved paths. Thus, the linearity measurement is slightly higher than if only static points are included.

Finally, the highest performance configurations for both static and dynamic accuracy are displayed to indicate the settings for the best found system performance. The best configuration found for dynamic tracking was: MIQR=41, ave=5, and dynamic calibration. The best configuration found for static accuracy was MIQR=17, ave=5, $\alpha=40,000$, and dynamic calibration.




Figure 122: Three orthographic views a) Top b) Front c) Side view of dynamic track experiment. All dimensions are in mm, where UWB data (blue) is plotted alongside Optotrak data (red). Optotrak data in the top view exhibits translation in cross-track direction due to a loose sliding bracket




Figure 123: Hysteresis testing showing positive (blue) and negative (red) motion down a track on two occasions. Note local nonlinearities which appear to be correlated in each direction. Hysteresis error is 4.7% and 3.3% of full-scale, respectively




Figure 124: 3D large range linearity testing with dynamic robot motion. Optotrak data (red) is notably smoother, however, the UWB (blue) system exhibits only a slight linearity error (~1%)




Figure 125: Robot repeatability experiment, where each of 20 points was visited twice. UWB data (blue) and Optotrak points are plotted for direct comparison; RMS 3D error is 7.00 mm




Figure 126: Best results found: 3.26 mm for static positioning and 5.14 mm for dynamic positioning

8.2. Error Budget

8.2.4 Accuracy

To fully understand system accuracy, one must understand all of the factors that can degrade the final positioning accuracy. Table 30 gives an outline of system error contributions. The errors are listed as observed error if the error was observed through experiment or directly measured. The theoretical worst case calculation results from

8.2.4.1 Sampling Error

The errors originating from sampling effects are two-fold. First, analog sampling performed by the sampling mixer compresses the high-bandwidth down-converted signal into a low bandwidth that can be sampled adequately by the ADC. In a basic sense, the sampling mixer exerts a larger effect on the sampling resolution. However, since inter-sample analog signals can approximate

Table 30: Breakdown of system errors with observed and theoretical magnitudes

Error Type	Observed 3D RMS Error	Theoretical Worst Case
Sampling Error		
- ADC (50 MS/s)	-	0.3 mm
- Sampling Mixer (100 GS/s)	1.49 mm vs. 400 Gs/s	1.5 mm
- Sampling Mixer (400 GS/s)	-	0.375 mm
Pulse Width (300 ps)	-	0.001 mm
Base station		
- Position	1-2 mm	5 mm (manual measurement)
- Orientation	3-5 mm	10 mm @ $\pm 45^\circ$ boresite angle
PRF non-coherence	4.7 mm vs. wired	1-2 mm (jitter measurement)
- 10k scale factor	-	1.1 cm / m
- 40k scale factor	-	4.0 cm / m
- 100k scale factor	-	11 cm / m
LO B.S. non-coherence	180 mm	-
AWGN		
- 30 SNR	-	0.550 mm
- 15 SNR	-	4.30 mm
- 3 SNR	-	28.7 mm
Multipath	5-10 mm	7.65 mm
Dynamic Motion	0.08 / (mm/s) mm	0.0004 / (mm/s) mm

the original signal, the over-sampling on the part of the ADC (nominally at 5 times oversampling), can partially retrieve this information.

8.2.4.2 Pulse Width

The width of the Gaussian pulse can affect the positioning accuracy, primarily due to errors in accurately performing the leading edge detection step. The slope of the rising edge, known as the slew rate, dictates the probability that the next sample will be greater than the previous sample under known noise conditions. The calculated positioning point is determined by the accuracy of assigning the correct sample location as the pulse time-of-arrival. For example, when the noise conditions yield a SNR of 30dB combined with the pulse, which has a leading edge slew rate of 1.3×10^{10} V/ps, the error probability of the peak detection algorithm being off by one sample is

0.32%, while the two sample probability is negligible. This phenomenon results in an expected error of significantly less than one sample.

8.2.4.3 Error due to Base Station Positioning and Base Station Orientation

The accuracy of positioning is highly dependent on the accuracy of the known locations of the base stations. Additionally, if the base station orientations are unknown significant errors can arise, which has been shown to be as high as 10 mm, if the broadside Transmitter-Receiver angle exceeds 45° . If the base stations positions are known to the millimeter range, a calibration algorithm has been devised to calibrate base station positions and orientations. This algorithm has been treated in depth and has been found to eliminate these errors to levels less than the absolute magnitude of the sampling error.

8.2.4.4 Error originating from pulse repetition rate non-coherence

For this type of error the theoretical value differed from the observed value by a factor of two. Since this type of error is related to a scale difference, the actual error value must be related to the physical location of the tag relative to the base stations. The incongruity between the theoretical and observed results must be examined in terms of percentage scale range errors rather than 3D RMS errors to be consistent.

8.2.4.5 Error originating from local oscillator non-coherence

Again, the theoretical examination of this topic in Chapter 3 suggested that LO non-coherence should only cause a DC bias to the incoming down-sampled signals. In the case when only a signal channel is used for each base station, this error is increased with the effect of the carrier offset. Any LO shift (nominal or dynamic) toward a sub-harmonic value of the carrier LO will

create larger frequency leakages in the down-converted, and thus the sub-sampled signals. This carrier leakage can cause non-linear effects in 1D position values in static positioning and more troublesome, non-linearities in dynamic tracking scenarios.

8.2.4.6 Overall Error

8.2.5 Dynamic Error

Maximum dynamic error can be defined as the maximum difference in position between a tag/probe and the values displayed on the screen. Under the design consideration of 24 Hz update rate per tag (with roughly 85 tags, thus 2.04 kHz system wide), the upper bound on dynamic error is dictated by the maximum speed the tags are typically moved. If accuracy is desired while moving the tags at 0.5 m/s, for example, the dynamic error will be 0.21 mm. If higher accuracy is desired, slower speeds must be used when acquiring points. If fewer than 100 tags are active, a higher rate of accuracy is also possible, due to the higher sampling frequency of individual tags.

8.2.6 Bandwidth of system

Considering the difference between narrowband and UWB positioning systems, bandwidth is a critical parameter. Narrowband systems are traditionally limited by 1/10 of wavelength as the physical limit of 1D positioning accuracy. When greater amounts of the spectrum are utilized for positioning, position accuracy can be increased linearly with bandwidth. Given a certain carrier frequency (or central frequency) of transmission, the bandwidth cannot be increased without limitation, however, because the bandwidth can be a maximum of twice the carrier frequency (in this case the used spectrum would go from DC- $2F_c$). The system has a carrier frequency of 8 GHz, and a -10dB bandwidth of 6 GHz. Also, the bandwidth is directly (Table 31) related to the

maximum slew rate of the system and thus forms the underlying physical phenomenon by which algorithms such as the Max-Ratio algorithm are highly dependent.

8.3. System Time Budget

When examining system performance on a large system level scale, one important parameter is the number of tags that can be supported by the system and the frequency at which points can be taken for each tag. For this analysis, two approaches may be taken. The first is a top-down view of the system. In this analysis, the starting point is the user requirement for a refresh rate of 24Hz to visualize points on a screen. For a system with only one tag, the scaling factor is directly related to the frequency at which 3D points may be acquired. For a scaling factor of $\alpha = 10,000$, 1 ms is required for a single point acquisition, thus this is a physical limit of 41 tags at 24 Hz. If the scaling factor increases the point acquisition time linearly, then the number of tags the system can support decreases inversely. Another limiting factor is the rate at which data can be read off of the FPGA by the serial cable. At this rate, a maximum of 85 tags can be positioned at 24 Hz. Alternatively, a higher speed data link could be employed (i.e. USB) to avoid this constraint. The system software running the TDOA algorithm can also process the incoming time indices and refresh the screen faster than this rate, thus, not posing a constraint at this time.

Table 31: Maximum slew rate compared to system -10 dB band width

Slew rate (V/ps)	Bandwidth (GHz)
6.67x10 ⁸	0.398
1.33 x10 ⁹	0.795
2.22 x10 ⁹	1.33
6.67 x10 ⁹	3.98
1.33 x10 ⁹	7.95
3.33 x10 ⁹	19.9

Using a bottom-up approach, the system must employ a communication scheme to turn tags on and off in quick succession. To position one tag, the following actions must occur: turn off the previous positioned tag, turn on the current tag, and finally allow the microwave system to sample the pulse train. The tag turn off and turn on times of the tag are estimated to be 2.5 μ s and 2.5 μ s, respectively. At this rate the 1 ms sample acquisition time will dominate the time to collect the sample (Table 32). Taking a conservative approach, it is possible to assume that the analog signal will not be fully set up within 1 ms cycle, thus, if two complete cycles are required only about 20 tags may be supported at 24 Hz.

Table 32: System characteristics which determine tag positioning throughput and latency

Effect		Throughput	Latency
Tag on/off	Turn	200 kHz	5.0 μ s
Pulse Repetition Rate		1 – 2 kHz	0.5-1.0 ms
Max-Ratio		~100 kHz	5.44 μ s
Serial Cable		2.04 kHz	490 μ s
TDOA		1.33 kHz	375.5 μ s

9. DISCUSSION AND FUTURE WORK

Demonstrated in this dissertation was a host of surgical tools for applications in pre-planning, intra-operative, and post-operative phases of primarily orthopedic surgeries. The component of this body of research was the development, testing, and theoretical discussion of aspects of a novel UWB microwave positioning technology posited as a subsystem in the overall scope of this research. This positioning technology is, at the time of this writing, the highest performing wireless UWB positioning system that exists in the known literature or as a commercial system. This system exceeds the accuracy of commercially available systems by more than an order of magnitude (~ 5.5 mm accuracy vs. 10-15 cm) (Table 33) and nearly matches the accuracy of the highest performing system when in the wired configuration (0.73 mm this system vs. 0.1 mm of Meier et al.). While this result furthers the state-of-the-art, more work is needed to see use of this positioning technology for its intended application as the backbone of a surgical navigation system. A rich set of solutions exists for extending and improving the accuracy, performance, and robustness of the system.

Table 33: Comparison of demonstrated system with commercially available and research-grade systems in existence

Positioning System	Type of System	Reported Accuracy
Ubisense TM	Wireless	15 cm
Sapphire Dart TM	Wireless	10 cm
Demonstrated System	Wireless	5.5 mm
Zetik et al.	Wired	1.5 cm (2D)
Low et al.	Wired	1.0 cm (1D)
Demonstrated System	Wired	0.73 mm (3D)
Meier et al.	Wired	0.1-2.0 mm

9.1. System Overview

The detailed descriptions and analysis provided in this dissertation serve to highlight some of the challenges of working in the UWB and high frequency RF domains. Before discussing the internals of the UWB positioning system, the early work in Chapter 2 examined and compared two existing technologies used in surgical navigation systems. These systems were tested in a cadaveric hip surgery and in a number of other static and dynamic environments. By framing the discussion in terms of real surgical tasks and goals, discontinuities arose between the performances of the two systems. Next, in Chapter 3 a broad description of our UWB system as it was published was presented. This work introduced the effects of channel modeling, unintentional time scaling, sub-sampling, multipath interference, peak detection, and 3D positioning. Additional work provided a look at the channel modeling work that served as the framework for the peak and leading edge algorithm development in Chapter 4. By studying the effects of multipath interference on a large and static object an in-depth understanding was formed on how this type of interference corrupts the transmitted signal. With this knowledge, several peak and leading edge detectors were studied in Chapter 5, which led to the novel discovery of the Max-Ratio algorithm. Also, the GP technique was demonstrated to simplify and produce new peak detection programs. Although, in the end the GP-produced peak detection results were not superior to the Max-Ratio algorithm, the method provides an efficient way to test new programs should system RF signals change substantially. Next, in Chapter 6, robust positioning methods were examined in the context of differing base station geometries with reducing the final system error as the ultimate goal. Several filtering techniques were studied including Kalman filtering of the 3D positions, as well as non-linear filtering of the indices using

the MIQR filter. To ensure that the system is viable in different environments, three types of system calibration were discussed in Chapter 7. These calibration methods highlight some of the system-level challenges including the unintentional time scaling and the base station phase center error. Chapter 8 points out the system characteristics under a number of parameter and environmental settings. The experimental results in this section demonstrate the final system performance of 3.2 mm static accuracy and 5.1 mm dynamic error.

9.2. Future System Improvements

The future improvements to the system fall into many different categories: performance related, aesthetically related, or commercialization related. Performance related improvements are directed primarily at reducing or mitigating error effects from physical phenomenon, or by improving the robustness of the system with respect to the integration of components of the system itself. Aesthetically driven improvements will yield improvements in system usability, compactness, and ergonomic system appearance. While commercialization related improvements share significant overlap with the previous two categories, some of the motivation in the commercial domain is fabrication cost and quality function deployment. The following ideas are summarized in Table 34.

9.3. Future Work

In tackling the fundamental problems of the system, the tradeoffs that exist between sample acquisition time and accuracy amounts to a Pareto-type multi-objective optimization problem. In multi-objective optimization many new and interesting techniques are currently receiving large amounts of attention in the current computational intelligence literature. The characteristics of

Table 34: Categorized system improvements, with expectations on improvement metrics

Category of System Improvement	Type of Improvement	Description of expected results
Performance	IQ Dual Channel per base station	Doubles fabrication cost per base station, could add 20-30% error improvement
Performance	Coherent FPGA clocks, or mixed-signal ASIC design	Improves noise tolerance, clocking sampling jitter reduction improves performance by 2-5%
Performance	Additional base stations	Improves PDOP values, improves system robustness in dynamic blockage environments
Performance	Probe design: antenna null zone elimination	Optimized antenna orientations will reduce potential for system malfunction in orientation
Performance	System supports multiple tags in TDMA communications	System probe computed error can be decreased with redundant number of tags per probe
Aesthetic	Probe design	Early adopters will enjoy and promote products which suit their needs uniquely
Aesthetic / Commercialization	Tag and base station miniaturization	Significantly increases the positioning potential in constrained spaces, reduces bulky components
Commercialization	Tag-level communications using UWB channel	Less regulatory hurdles as all system functions can be handled at one frequency band
Commercialization	Telemetry: using tag communications	Tag understands and communicates co-located sensor data
Commercialization	Board-level fabrications of base station and control station hardware	Reduces system fabrication costs
Commercialization	Hermetic sealing	Prevent liquids or unsanitary substances from contaminating electronics on the tag

the base station calibration algorithm will continue to be an area where more dedicated research will be focused. Certain aspects of the hardware design improvements, while intriguing, will be out of the scope of this future work, although hardware devices that modify the shape or characteristics of the transmitted pulse must be carefully examined. Additional work related to improving the signal-to-noise characteristics of the digital signal processing is also of interest. Advances in FPGA performance would be a key leveraging point for more computationally intensive algorithms. While this investigation focused on the use of active tags another important vein of research is that of passive tags. In this scenario, tags could be positioned using round-trip time-of-flight information instead of strictly time-of-flight which is used in this system.

REFERENCES

- [1] E. Elfenbein., The Miracle of Warsaw, Indiana. *crossingwallstreet.com*. [Online] 2008.
[Cited: 10 1, 2008.]
http://www.crossingwallstreet.com/archives/2006/03/the_miracle_of.html.
- [2] Hopson, Krista., Big boom in boomer knee replacement surgeries. *UMHS Newsroom*.
[Online] June 2008. [Cited: 10 1, 2008.]
<http://www2.med.umich.edu/prmc/media/newsroom/details.cfm?ID=333>.
- [3] T. M. Wright, C. M. Rimnac, S. D. Stulberg., "Wear of polyethylene in total joint replacement. Observations from retrieved PCA knee implants." *Clin. Orthop.*, 1992, Vol. 276, pp. 126-134.
- [4] M. Mahfouz, R. Booth Jr, J. Argenson, B.C. Merkl, E.E. Abdel Fatah, M. Kuhn., "Analysis of variation fo adult femora using sex specific statistical atlases." *Cote d'Azur, France : Computer Methods in Biomechanics and Biomedical Engineering Conference*, 2006.
- [5] M.R. Mahfouz, B.C. Merkl, E.E. Fatah., "Automatic Methods for Characterizing of Sexual Dimorphism of Adult Femora: Distal Femur." *Computer Methods in Biomechanics and Biomedical Engineering Journal*, December 2007, Issue 6, Vol. 10, pp. 447-456.
- [6] S.C. Faber et al., "Gender difference in knee joint cartilage thickness, volume and articular surface areas: assessment with quantitative three-dimensional MR imaging." *Skeletal Radiol.*, 2001, Vol. 30, pp. 144-150.
- [7] M. Moore, A. Sylvester, B. Merkl, M. Kuhn, M. Mahfouz., "Creating a Statistical Atlas of Femora from Three-Dimensional CT Data." *Anchorage : Annual Meeting of American Association of Physical Anthropologists*, 2006.

- [8] A. M. Badawi, M. A. El-Mahdy, E. E. Abdel Fatah., "4D-Echocardiographic System for Wall Motion and Cardiac Parameters Quantification." s.l. : Proceedings of The 46th IEEE International Midwest Symposium on Circuits and Systems (MWSCAS), 2003.
- [9] M. Kuhn, B. Merkl, M. Mahfouz, R. Komistek., "3D-to-2D Registration of Implant Models to X-ray Images with Special Emphasis on Error Analysis, Segmentation Effects, and Polyethylene Wear." Kyoto, Japan : 18th Annual Symposium of the International Society for Technology in Arthroplasty, 2005.
- [10] M. Kuhn, B. Merkl, M. Mahfouz., "Advanced 3D Analysis of Implanted Joints through 3D-to-2D Registration of Implant Models to Fluoroscopic Images." Antibes, Cote d'Azur, France : 7th International Symposium on Computer Methods in Biomechanics and Biomedical Engineering, 2006.
- [11] M. Mahfouz, W. Hoff, R. Komistek, D. Dennis., "A Robust Method for Registration of Three-Dimensional Knee Implant Models to Two-Dimensional Fluoroscopy Images." IEEE Trans. Med. Imaging, December 2003, Issue 12, Vol. 22, pp. 1561-1572.
- [12] Kuhn M., Mahfouz M., Emam E. Abdel Fatah, Merkl B., "Reconstruction of 3D Patient-Specific Bone Models from Biplanar X-ray Images." Singapore : 12th International Conference on Biomedical Engineering, 2005.
- [13] M.R. Mahfouz, H.E. Dakhakhni, E.E. A. Fatah, R. Tadross, R.D. Komistek., "Three-Dimensional Bone Creation And Landmarking Using Two Still X-Rays." San Francisco : 54th Annual Meeting of Orthopaedic Research Societies, 2008.
- [14] Y. Yoshioka, D. Siu and D. Cooke., "The anatomy and functional axes of the femur." J. Bone Joint Surg. Am., 1987, Issue 6, Vols. 69-A, pp. 873-880.

- [15] M. Mahfouz, E. Fatah, B. Merkl, J. Mitchell., "Automatic and manual methodology for 3-dimensional measurement of distal femoral gender differences and femoral component placement." J. Knee Surgery, 2008. In Revision.
- [16] J.O. Drogset, MD, T. Gøntvedt, MD, PhD, O.R. Robak, MD, A. Mølster, MD, PhD, A. T. Viset, MD and L. Engebretsen, MD, PhD., "A Sixteen-Year Follow-up of Three Operative Techniques for the Treatment of Acute Ruptures of the Anterior Cruciate Ligament." s.l. : Journal of Bone and Joint Surgery, 2006, Vol. 88, pp. 944-952 .
- [17] D.A. Dennis, R.D. Komistek, M.R. Mahfouz, B.D. Haas, J.B.Stiehl., "Multicenter Determination of In Vivo Kinematics after Total Knee Arthroplasty." Clinical Orthops & Related Research., 2003, Vol. 416, pp. 37-57. Coventry Award Paper.
- [18] E. E. Abdel Fatah, M. R. Mahfouz, B. C. Merkl , M. A. Hartzband., "Automatic Method for Calculation Of Clinical Transepicondylar Axis Using Statistical Atlases." San Francisco : 54th Annual Meeting of Orthopaedic Research Societies , 2008.
- [19] G. To, M. Mahfouz, E. Pritchard., "Ligament Balancing during Total Knee Arthroplasty with Wireless Encapsulated Microcantilever Strain Sensors." Singapore : Intl. Conf. on Biomed. Engr. Proceedings,, 2005.
- [20] B.C. Merkl, M.J. Kuhn, M.R. Mahfouz, D. DeBoer., "Surgical Navigation Systems: Evaluating Electromagnetic versus Optical Technology in the OR." San Francisco, CA : Annual Meeting of American Academy of Orthopaedic Surgeons, 2008.
- [21] Hitt, K. and Shurman, J.R. and Greene, K. and McCarthy, J. and Moskal, J. and Hoeman, T. and Mont, M.A., "Anthropometric Measurements of the Human Knee: Correlation to the

- Sizing of Current Knee Arthroplasty Systems." *The Journal of Bone and Joint Surgery*, 2003, Issue 90004, Vol. 85, pp. 115-122.
- [22] Maruyama, M. and Feinberg, JR and Capello, WN and D'Antonio, JA., "The Frank Stinchfield Award: Morphologic features of the acetabulum and femur: anteversion angle and implant positioning." *Clin Orthop Relat Res*, 2001, Vol. 393, pp. 52-65.
- [23] M.R. Mahfouz, E. E. Abdel Fatah, M. A. Hartzband, A. H. Glassman , L. N. Smith., "Three Dimensional Assessment of Proximal Femoral Morphology ? Can It Predict The Bone Quality?" San Francisco : s.n., 2008. 54th Annual Meeting of Orthopaedic Research Societies.
- [24] Medtronic, Inc., Medtronic Ortho Axiem Electromagnetic Navigation. *Medtronic Navigation*. [Online] [Cited: September 3, 2008.]
<http://www.medtronicnavigation.com/procedures/orthopaedic/knee.jsp>.
- [25] Corporation, Stryker., Knee Navigation Surgery eNact : Stryker. *stryker.com*. [Online] [Cited: 7 7, 2008.] <http://www.stryker.com/en-us/products/Orthopaedics/ComputerAssistedSurgery/KneeNavigationSurgery/eNact/index.htm>.
- [26] E. Stindel, J.L. Biard., P. Merloz, S. Plaweski, F. Dubrana, C. Lefevre, J. Troccaz., "Bone Morphing: 3D Morphological Data for Total Knee Arthroplasty." s.l. : Computer Aided Surgery, 2002, Vol. 7, pp. 156-168.
- [27] S. Conley., BRIGIT™ Bone Resection Instrument Guide*. *zimmer.com*. [Online] [Cited: 10 12, 2008.]
<http://www.zimmer.com/z/ctl/op/global/action/1/id/9655/template/MP/navid/7456>.

- [28] Martin Roche, M.D., The MAKOpasty® Procedure . *makosurgicalcorp.com*. [Online]
[Cited: 10 12, 2008.] <http://www.makosurgicalcorp.com/08/makoplasty/index.htm>.
- [29] Intuitive Surgical, Inc., The da Vinci® Surgical System. *intuitivesurgical.com*. [Online]
[Cited: 10 12, 2008.]
http://www.intuitivesurgical.com/products/davinci_surgicalsystm/index.aspx.
- [30] Curexo Technology Corporation., ROBODOC® Surgical Assistant. *robodoc.com*. [Online]
[Cited: 10 12, 2008.] <http://www.robodoc.com/home.html>.
- [31] Greig, K., *Communications about system expectations*. 3 18, 2008.
- [32] E. Berg, M. Mahfouz, C. Debrunner, B. Merkl, W. Hoff., "Fourier Descriptor-Based Deformable Models for Segmentation of the Distal Femur in CT." [book auth.] J. Pejaš ed. K. Saeed. *Information Processing and Security Systems*. New York : Springer, 2005.
- [33] E. Berg, M. Mahfouz, C. Debrunner, W. Hoff., "Automatic segmentation of 3D medical images using Fourier descriptor-based deformable models." Elk, Poland : s.n., 2004. Intl. Multiconference on Adv. Comp. Sys. and Comp. Information Sys. and Industrial Mgt. Applications.
- [34] B.C. Merkl, M. Mahfouz., "Unsupervised Three-Dimensional Segmentation of Medical Images Using an Anatomical Bone Atlas." Singapore : 12th International Conference on Biomedical Engineering, 2005. Outstanding Paper Award.
- [35] E. Berg, M. Mahfouz, C. Debrunner, W. Hoff., "A 2D Fourier approach to deformable model segmentation of 3D medical images." Prague, Czech Republic : s.n., 2004. Workshop on Computer Vision Applications in Medical Image Analysis.

- [36] B.C. Merkl, L.N. Smith, M.R. Mahfouz, A.M. Badawi., "Fuzzy-neural multiple bone sex classification system utilizing genetic algorithms." Cairo, Egypt : 3rd Cairo International Biomedical Engineering Conference, 2006.
- [37] M. Mahfouz, A. Badawi, B. Merkl, E.E. Abdel Fatah, E. Pritchard, K. Kesler, M. Moore, L. Jantz, R. Jantz., "Patella sex determination by 3D statistical shape models and nonlinear classifiers." Forensic Science International, Dec. 2007, Issue 2-3, Vol. 173, pp. 161-170.
- [38] Kuhn, MJ, et al., "Reconstruction of 3D Patient-Specific Bone Models from Biplanar X-ray Images." Singapore : 12th International Conference on Biomedical Engineering, 2005.
- [39] M. Mahfouz, E. Fatah, H. E Dakhakhni, R. Tadross, R.D. Komistek., "Three-Dimensional Bone Creation and Landmarking Using Two Still X-Rays." San Francisco : Annual Meeting of American Academy of Orthopaedic Surgeons, 2008.
- [40] A. Sylvester, B. Merkl, M. Mahfouz., "Assessing AL 288-1 femur length using computer-aided three-dimensional reconstruction." J Human Evolution, 2008, Issue 4, Vol. 55, pp. 665-671.
- [41] B.C. Merkl, A. Sylvester, M. Mahfouz., "A three-dimensional shape comparison of AL129-1a and modern human distal femora." Philadelphia : Annual Meeting of American Association of Physical Anthropologists, 2007.
- [42] M.R. Mahfouz, C. Zhang, B.C. Merkl, M.J. Kuhn, A.E. Fathy., "Investigation of High-Accuracy Indoor 3-D Positioning Using UWB Technology." Microwave Theory and Techniques, IEEE Transactions on, 2008, Issue 6, Vol. 56, pp. 1316-1330.
- [43] Federal Communications Commission., NEW PUBLIC SAFETY APPLICATIONS AND BROADBAND INTERNET ACCESS AMONG USES ENVISIONED BY FCC

AUTHORIZATION OF ULTRA-WIDEBAND TECHNOLOGY. *FCC Web site*. [Online]
[Cited: May 2, 2008.]

http://www.fcc.gov/Bureaus/Engineering_Technology/News_Releases/2002/nret0203.html.

- [44] Z.N. Low, J.H. Cheong, C.L. Law, W.T. Ng, Y.J. Lee., "Pulse detection algorithm for line-of-sight (LOS) UWB ranging applications." *IEEE Ant. and Wireless Prop. Letters*, 2005, Vol. 4, pp. 63–67.
- [45] C. Meier, A. Terzis, S. Lindenmeier., "A robust 3D high precision radio location system." Honolulu, Hawaii : IEEE MTT-S International Microwave Symposium, 2007. pp. 397 – 400.
- [46] N. Nyirongo, W.Q. Malik, D.J. Edwards., "Concatenated RS-Convolutional Codes for Ultrawideband Multiband-OFDM." Waltham, MA : ICUWB, 2006. pp. 137-142.
- [47] Tomlinson, R. Hootor and H., "Delay-Hopped Transmitted-Reference." s.l. : IEEE Conf. on Ultra Wideband Syst. and, 2002. pp. 265-270.
- [48] Qiu, R.C., "A Theory of Time-Reversed Impulse Multiple-Input Multiple-Output (MIMO) for Ultra-Wideband (UWB) Communications." Waltham : ICUWB, 2002. pp. 587-592.
Invited Paper.
- [49] K. Liu, L. Cai, X. Shen., "Multiclass Utility-Based Scheduling for UWB Networks." *IEEE Trans. Vehicular Tech.*, March 2008, Issue 2, Vol. 57, pp. 1176-1187.
- [50] S. Zhao, H. Liu., "Transmitter-Side Multipath Preprocessing for Pulsed UWB Systems Considering Pulse Overlapping and Narrow-band Interference." *IEEE Trans. Vehicular Tech.*, November 2007, Issue 6, Vol. 56, pp. 3502-3510.

- [51] F. Langlotz, L.P. Nolte., "Technical approaches to computer-assisted orthopedic surgery." s.l. : Euro. J. of Trauma, 2004, Issue 1, Vol. 30, pp. 1-11.
- [52] E.A. Hassan, T.R. Jenkyn, C.E. Dunning., "Direct comparison of kinematic data collected using an electromagnetic tracking system versus a digital optical system." s.l. : J. Biomechanics, 2007, Issue 4, Vol. 40, pp. 930-935.
- [53] L.P. Maletsky, J. Sun. N.A. Morton., "Accuracy of an optical active-marker system to track the relative motion of rigid bodies." s.l. : J. Biomechanics, 2007, Issue 3, Vol. 49, pp. 682-685.
- [54] N.B. Schuler, M.J. bey, J.T. Shearn, D.L. Butler., "Evaluation of an electromagnetic position tracking device for measuring in vivo, dynamic joint kinematics." s.l. : J. Biomechanics, 2005, Issue 10, Vol. 38, pp. 2113-2117.
- [55] F. Poulin, L.P. Amiot., "Interference during the use of an electromagnetic tracking system under OR conditions." s.l. : J. Biomechanics, 2002, Issue 6, Vol. 35, pp. 733-737.
- [56] R. Khadem, C.C. Yeh, M. Sadeghi-Tehrani, M.R Bax et al., "Comparative tracking error analysis of five different optical tracking systems." Comp. Aid. Surg., 2000, Issue 2, Vol. 5, pp. 98-107.
- [57] A.M. DiGioia, B. Jaramaz, A.Y. Plakseychuk, J.E. Moody, C. Nikou, R.S. LaBarca, T. J. Levison, F. Picard., "Comparison of a mechanical acetabular alignment guide with computer placement of the socket." J. Arthroplasty, 2002, Issue 3, Vol. 17, pp. 359-364.
- [58] B.M. Jolles, P. Zangger, P.F. Leyvraz., "Factors predisposing to dislocation after primary total hip arthroplasty." J. Arthroplasty, 2002, Issue 3, Vol. 13, pp. 282-288.

- [59] L.D. Dorr, Y. Hishiki, Z. Wan., "Development of imageless computer navigation for acetabular component position in total hip replacement." *Iowa Orthop. J.*, 2005, Vol. 25, pp. 1-9.
- [60] Vossiek, M, et al., "Wireless local positioning." *IEEE Microwave Magazine*. Dec. 2003, Vol. 4, pp. 77-86.
- [61] N. Patwari, J.N. Ash, S. Kyperountas, A.O. Hero III, R.L. Moses, N.S. Correal., "Locating the nodes: cooperative localization in wireless sensor networks." *IEEE Signal Processing Magazine*. July 2005, Vol. 22, pp. 54-69.
- [62] Kaplan, E. D., *Understanding GPS: principles and applications*. Boston, MA : Artech House, 1996.
- [63] A. Stelzer, K. Pourvoyeur, A. Fischer., "Concept and application of LPM-a novel 3-D local position measurement system." *IEEE Trans. Microwave Theory and Tech.*, 2004, Vol. 52, pp. 2664-2669.
- [64] M. Pichler, A. Stelzer, P. Gulden, C. Seisenberger, M. Vossiek., "Phase-error measurement and compensation in PLL frequency synthesizers for FMCW sensors--I: context and application." *IEEE Trans. Cir. and Sys. I*, May 2007, Vol. 54, pp. 1006-1017.
- [65] L. Reindl, C.C.W. Ruppel, S. Berek, U. Knauer, M. Vossiek, P. Heide, L. Oreans., "Design, fabrication, and application of precise SAW delay lines used in an FMCW radar system." *Apr. 2001*, Vol. 49, pp. 787-794.
- [66] Symeo., LPR-A. Symeo. [Online] 2007. http://www.symeo.com/Products_LPR_A.htm.
- [67] L. Wiebking, M. Glanzer, D. Mastela, M. Christmann, M. Vossiek., "Remote local positioning radar." *s.l. : IEEE Radio and Wireless Conference*, 2004. pp. 191-194.

- [68] S. Roehr, M. Vossiek, P. Gulden., "Method for high precision radar distance measurement and synchronization of wireless units." s.l. : IEEE MTT-S International Microwave Symposium, 2007. pp. 1315-1318.
- [69] Fontana, R.J., "Recent system applications of short-pulse ultra-wideband (UWB) technology." IEEE Trans. Microwave Theory and Techn., Sept. 2004, Issue 9, Vol. 52, pp. 2087-2104.
- [70] Multispectral Solutions, Inc., "Sapphire DART (RTLS) Product Data Sheet." *Multispectral Solutions, Inc.* [Online] http://www.multispectral.com/pdf/Sapphire_DART.pdf.
- [71] , "Hardware Datasheet." *Ubisense*. [Online] http://www.ubisense.net/SITE/UPLOAD/Document/TechDocs/Ubisense_hardware_datasheet_May_2006.pdf.
- [72] R. Zetik, J. Sachs, R. Thoma., "UWB localization - active and passive approach." s.l. : Proceedings of the 21st IEEE IMTC, 2004. Vol. 2, pp. 1005-1009.
- [73] C. Zhang, M. Kuhn, B. Merkl, M. Mahfouz., "Accurate UWB indoor localization system utilizing time difference of arrival approach." s.l. : IEEE Radio and Wireless Symposium, 2006. pp. 515-518.
- [74] C. Zhang, A.E. Fathy., "Reconfigurable pico-pulse generator for UWB applications." s.l. : IEEE MTT-S International Microwave Symposium, 2006. pp. 407-410.
- [75] Nguyen, J.S. Lee and C., "Uniplanar picosecond pulse generator using step-recovery diode." s.l. : IEE Electronics Letters, 2001, Vol. 37, pp. 504-506.
- [76] C. Zhang, and A.E. Fathy., "Development of an ultra-wideband exponentially tapered antipodal Vivaldi antennas." s.l. : Proc. IEEE AP-S Int. Symp. , 2005. pp. 1689-1692.

- [77] S.G. Kim, K. Chang., "Ultra-wideband exponentially tapered antipodal Vivaldi antennas." s.l. : Proc. IEEE AP-S Int. Symp, 2004. pp. 2273-2276.
- [78] Y. Yang, C. Zhang, S. Lin, A.E. Fathy., "Development of an ultrawideband Vivaldi antenna array." s.l. : Proc. IEEE AP-S Int. Symp., 2005. pp. 606-609.
- [79] A. Elsherbini, C. Zhang, A.E. Fathy, M. Mahfouz., "notch bandwidth and filter order trade-offs in band-notched monopole antennas for UWB applications." s.l. : IEEE AP-S Int. Symp., 2008. Submitted..
- [80] C. Zhang, A. Fathy, M. Mahfouz., "Performance enhancement of a sub-sampling circuit for ultra-wideband signal processing." IEEE Micro. and Wireless Comp. Lett. , 2007, Issue 12, Vol. 17, pp. 873-875.
- [81] M. Gerding, T. Musch, B. Schiek., "A novel approach for a high-precision multitarget-level measurement system based on time-domain reflectometry." s.l. : IEEE Trans. on Microwave Theory and Tech., 2006, Vol. 54, pp. 2768-2773.
- [82] J. Han, C. Nguyen., "Coupled-slotline-hybrid sampling mixer integrated with step-recovery-diode pulse generator for UWB applications." s.l. : IEEE Trans. on Microwave Theory and Tech., June 2005, Vol. 53, pp. 1875-1882.
- [83] M.M. Zinieris, R. Sloan, L.E. Davis., "A broadband microstrip-to-slot-line transistion." s.l. : Microwave and Optical Technology Letters, 1998, pp. 339-342.
- [84] Datasheet, Tektronix Electrical Sampling Modules., [Online] <http://www.tek.com>.
- [85] Datasheet, Picosecond Pulse Labs 7040 Sampler Module., [Online] <http://www.picosecond.com>.

- [86] M. Shinagawa, Y. Akazawa, T. Wakimoto., "Jitter analysis of high speed sampling system."
Kyoto : IEEE VLSI Circuits Symposium, 1989. pp. 95-96.
- [87] Vectron Intl., *VTC4 series voltage controlled temperature compensated crystal oscillator*.
Hudson, NH : Vectron Intl., 2007. p. 7.
- [88] Y.T. Chan, K.C. Ho., "Joint time-scale and TDOA estimation: analysis and fast
approximation." s.l. : IEEE Trans. on Signal Processing, August 2005, Issue 8, Vol. 53, pp.
2625-2634.
- [89] J.A. Costa, N. Patwari, A.O. Hero., "Distributed weighted multidimensional scaling for
node localization in sensor networks." s.l. : ACM Trans. Sen. Netw., Feb. 2006, Issue 1,
Vol. 2, pp. 39-64.
- [90] Vectron Intl., *Application note VTC4 clipped sine wave TCXO typical phase noise
measurements*. Hudson, NH : Vectron Intl, 2007.
- [91] M. Tuchler, V. Schwarz., "Location accuracy of an UWB localization system in a multi-
path environment." s.l. : IEEE Intl. Conf. on Ultra-Wideband, 2005. pp. 414-419.
- [92] A. Saleh, R. A. Valenzuela., "A statistical model for indoor multipath propagation." IEEE J.
Sel. Areas Commun., Feb. 1987, Vol. 5, pp. 128–137.
- [93] R. J.-M. Cramer, R. A. Scholtz, and M. Z. Win., "Evaluation of the Ultra-Wide-Band
Propagation Channel." IEEE Trans. Antennas Propagat., May 2002, Vol. 50, pp. 561–570.
- [94] Pamp, J., and Kunisch J., "Measurement Results and Modeling Aspects for the UWB Radio
Channels." Baltimore, MD : Proc. IEEE Conf. UWB Sys. and Tech., 2002. pp. 19–23.

- [95] Z. Irahhaute, H. Nikookar, and G. J. M. Janssen., "An overview of ultra wide band indoor channel measurements and modeling." IEEE Microw. Wireless Compon. Lett., Aug. 2004, Vol. 14, pp. 386–388.
- [96] Lee T., Wong T., "Multipath model selection for UWB channels." Waltham, MA : IEEE International Conference on Ultra-Wideband, 2006. pp. 85-89.
- [97] A. Sibille, C. Roblin, S. Bories, and A. C. Lepage., "A channel-based statistical approach to antenna performance in UWB communications." IEEE Trans. Antennas and Propag., Nov. 2006, Vol. 54, pp. 3207-3215.
- [98] A.F. Molisch, D. Cassioli, C. Chong, S. Emami, A. Fort, B.. Kannan, J. Karedal, J. Kunisch, H.G. Schantz, K. Siwiak, M.Z. Win., "A comprehensive standardized model for ultrawideband propagation channels." IEEE Trans. Antennas and Propag., Nov. 2006, Vol. 54, pp. 3151-3166.
- [99] R. J. Vaccaro, C. S. Ramalingam, and D. W. Tufts., "Least-squares time delay estimation for transient signals in a multipath environment." J. Acoust. Soc. Amer., July 1992, Vol. 92, pp. 210–218.
- [100] B. Denis, J. Keignart, and N. Daniele., "Impact of NLOS propagation upon ranging precision in UWB Systems." Oulu, Finland : IEEE Conference on Ultra Wideband Systems and Technologies, 2003. pp. 379-383.
- [101] R. Roy, and T. Kailath., "ESPRIT—estimation of signal parameters via rotational invariance techniques." IEEE Trans. Acoustics, Speech, and Signal Processing, Jul. 1989, Vol. 37, pp. 984-995.

- [102] J. Kusuma, A. Ridolfi, and M. Vetterli., "Sampling of communication systems with bandwidth expansion." New York, NY : IEEE International Conference on Communications, 2002. pp. 1601-1605.
- [103] De Clercq, W., et al., "Removing Artifacts and Background Activity in Multichannel Electroencephalograms by Enhancing Common Activity." Shanghai, China : Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the. pp. 3699-3702.
- [104] R. G. Vaughan, and N. L. Scott., "Super-resolution of pulsed multipath channels for delay spread characterization." IEEE Trans. Commun., Mar. 1999, Vol. 47, pp. 343-347.
- [105] C. Falsi, D. Dardari, L. Mucchi, and M. Z. Win., "Time of arrival estimation for uwb localizers in realistic environments." EURASIP Journal on Applied Signal Processing, 2006. 32082.
- [106] L. Mucchi, C. Falsi, D. Dardari, and M.Z. Win., "Channel parameters estimation for UWB realistic environments." Waltham, MA : IEEE International Conference on Ultra-Wideband, 2006. pp. 155-159.
- [107] Gough, P.T., "A fast spectral estimation algorithm based on the FFT." IEEE Trans. Signal Proc., June 1994, Vol. 42, pp. 1317-1322.
- [108] D. Dardari, A. Giorgetti, M.Z. Win., "Time-of-arrival estimation of the UWB signals in the presence of narrowband and wideband interference." Singapore : IEEE Conference on Ultra Wideband Systems and Technologies, 2007. pp. 71-76.
- [109] H.G. Shantz., "Introduction to ultra-wideband antennas." s.l. : IEEE Conference on Ultra Wideband Systems and Technologies, 2003. pp. 1-9.

- [110] G.M Wubben, M. Schmitz, F. Menge, V. Boder, G. Seeber., "Automated absolute field calibration of GPS antennas in real-time." s.l. : Proc. ION GPS 2000, 2000. pp. 2512-2522.
- [111] A. Elsherbini, C. Zhang, S. Lin, M. Kuhn, A. Kamel, A.E. Fathy, H. Elhennawy., "Enhancing the performance of antipodal Vivaldi antennas using a dielectric rod for high precision localization applications." s.l. : IEEE AP-S Int. Symp, 2007. pp. 1973-1976.
- [112] , "Optotrak 3020." *Northern Digital, Inc.* [Online] <http://www.ndigital.com/optotrak-techspecs.php>.
- [113] T. Zasowski, G. Meyer, F. Althaus, A. Wittneben., "UWB signal propagation at the human head." s.l. : IEEE Trans. Microwave Theory and Tech., June 2006, Vols. 54, part 2, pp. 1836-1845.
- [114] E.J. Zhong, T.Z. Huang., "Geometric dilution of precision in navigation computation." Dalian, China : Int. Conf. on Machine Learning and Cybernetics, Aug. 2006. pp. 4116-4119.
- [115] Nielsen, R., "Relationship between dilution of precision for point positioning and for relative positioning with GPS." IEEE Trans. on Aero. and Elec. Sys., Jan. 1997, Issue 1, Vol. 33, pp. 333-338.
- [116] I. Oppermann, M. Hämäläinen, J. Iinatti, [ed.], *UWB Theory and Applications*. s.l. : Wiley, 2004.
- [117] H. L. Bertoni., *Radio Propagation for Modern Wireless Systems*. Upper Saddle River : Prentice Hall PTR, 2000.
- [118] Pausini, Klaus Witrisal and M., "Statistical Analysis of UWB Channel Correlation Functions." IEEE Trans. Vehicular Tech., May 2008, Issue 3, Vol. 57, pp. 1359-1373.

- [119] X. Cheng, H. Shu, Q. Liang, D. Hung-Chang Du., "Silent Positioning in Underwater Acoustic Sensor Networks." *IEEE Trans. Vehicular Tech.*, May 2008, Issue 3, Vol. 57, pp. 1756-1766.
- [120] Nuallain, G. Gaertner and E.O., "Characterizing Wideband Signal Envelope Fading in Urban Microcells Using the Rice and Nakagami Distributions." *IEEE Trans. Vehicular Tech.*, November 2007, Issue 6, Vol. 56, pp. 3621-3630.
- [121] Tan, L. Jiang and S.Y., "Geometrically Based Statistical Channel Models for Outdoor and Indoor Propagation Environments." *IEEE Trans. Vehicular Tech.*, November 2007, pp. 3587-3593.
- [122] Remcom Inc., "Wireless InSite Overview." *remcom.com*. [Online]
<http://www.remcom.com/wirelessinsite/>.
- [123] R. J. Fontana., "Experimental Results from an Ultra-wideband Precision Geolocation System." [book auth.] Edited by P. D. Smith and S. R. Cloude. *Ultra-Wideband, Short-Pulse Electromagnetics*. s.l. : Kluwer Academic/Plenum Publishers, 2002, p. 215.
- [124] Kamil, Y. I., "Master's Thesis: Localization via Ultra-Wideband." *Dept. of Electrical and Electronic Engineering, Imperial College of London*. September 2006.
- [125] R. A. Saeed, S. Khatum, B. Ali, M. Khazani., "Performance of Ultra-Wideband Time-of-Arrival Estimation Enhanced With Synchronization Scheme." *ECTI Trans. on Electrical Eng., Electronics, and Communications* , February 2006, Issue 1, Vol. 4, pp. 78-84.
- [126] K. Pahlavan, X. Li, J-P Makela., "Next-Generation Broadband Wireless Networks and Navigation Services: Indoor Geolocation Science and Technology." *IEEE Communications Magazine*, February 2002, pp. 112-118.

- [127] Mentor Graphics., "ModelSim PE." *model.com*. [Online] [Cited: October 11, 2008.]
http://www.mentor.com/products/fpga_pld/simulation/modelsim_pe/.
- [128] Koza, J.R., *Genetic Programming: A Paradigm for Genetically Breeding Populations of Computer Programs to Solve Problems*. Computer Science Department, Stanford University. 1990. technical report. STAN-CS-90-1314.
- [129] Lennartsson, D. and Nordin, P., "A Genetic Programming Method for the Identification of Signal Peptides and Prediction of Their Cleavage Sites." *EURASIP Journal on Applied Signal Processing*, s.l. : Hindawi Publishing Corporation, 2004, Issue 1, Vol. 2004, pp. 138-145.
- [130] K.A. Farry, J.J. Fernandez, J.S. Graham., *Method of evolving classifier programs for signal processing and control*. 6272479 United States Patent, August 7, 2001. DATA PROCESSING - ARTIFICIAL INTELLIGENCE.
- [131] A.I.E. Alcazar, K.C. Sharman., "Some Applications of Genetic Programming in Digital Signal Processing." [ed.] John R. Koza. Palo Alto, CA : Stanford University, 1996. Late Breaking Papers at the Genetic Programming Conference . pp. 24-31.
- [132] K.L. Holladay, K.A. Robbins., "Evolution of Signal Processing Algorithms using Vector Based Genetic Programming." Cardiff, Wales, UK : Digital Signal Processing, 2007 15th International Conference on, 2007. pp. 503-506.
- [133] S. Hatscher, C. Diskus., "Pulse Based Radar Imaging Using a Genetic Optimization Approach for Echo Separation." s.l. : IEEE Sensors Journal, 2008. In Press..

- [134] Y. Yunqiang, A. Fathy., "Design and Implementation of a Low-Cost Real-Time Ultra-Wide Band See Through-Wall Imaging Radar System." Honolulu : s.n., 2007. IEEE International Microwave Symposium.
- [135] S. Hantscher, B. Praher, A. Reizenzahn, C. Diskus., "Comparison of UWB Target Identification Algorithms for Through-Wall Imaging Applications." Manchester, UK : Proc. of the 3rd European Radar Conference, 2006.
- [136] Vidyarthi, L. Kahnbar and D., "A GA-Based Effective Fault-Tolerant Model for Channel Allocation in Mobile Computing." IEEE Trans. Vehicular Tech., May 2008, Issue 3, Vol. 57, pp. 1823-1833.
- [137] G. Colman, T. Willink,, "Overloaded Array Processing Using Genetic Algorithms with Soft-Biased Initialization." IEEE Trans. Vehicular Tech., July 2008, Issue 4, Vol. 57, pp. 2123-2131.
- [138] Li, X., "Collaborative Localization With Received-Signal Strength in Wireless Sensor Networks." IEEE Trans. Vehicular Tech., November 2007, Issue 6, Vol. 56, pp. 3807-3817.
- [139] K. G. Chhokra, T. Bapty, J. Scott, M. Wilkes., "Accuracy Enhancements for TDOA Estimation on Highly Resource Constrained Mobile Platforms." Unpublished..
- [140] Poisel, R. A., *Electronic Warfare Target Location Methods*. Norwood, MA : Artech House, 2005.
- [141] M. Porretta, P. Nepa, G. Manara, F. Giannetti., "Location, Location, Location." IEEE Vehicular Tech. Magazine, June 2008, pp. 20-29.
- [142] Maki, Stanley., *All DOP GPS optimization. US Patent 5323163* June 21, 1994.

- [143] J. Schroeder, S. Galler, K. Kyamakya, T. Kaiser., "Three-dimensional indoor localization in non line of sight UWB channels." s.l. : Int. Conf. on Ultra-Wideband, 2007.
- [144] Y. Zhou, Y.L. Guan, C.L. Law, C. Xi., "High-Resolution UWB Ranging based on Phase-Only Correlator." Singapore : Int. Conf. on Ultra-Wideband, 2007.
- [145] D. Dardari, A. Giorgetti, M.Z. Win., "Time-of-arrival Estimation of UWWB Signals in the Presence of Narrowband and Wideband Interference." Singapore : Int. Conf. on Ultra-Wideband, 2007.
- [146] R. S. Figliola, D. E. Beasley., *Theory and Design for Mechanical Measurements*.
Singapore : Wiley, 1991.
- [147] Lorenz, C., Krahnstoever, N., "Generation of Point-Based 3D Statistical Shape Models for Anatomical Objects." Computer Vision and Image Understanding, 2000, Vol. 77, pp. 175-181.
- [148] B.C. Merkl, A.E. Fathy, M.R. Mahfouz., "Base Station Orientation Calibration in 3-D Indoor UWB Positioning." Hannover, Germany : ICUWB, 2008.

APPENDICES

APPENDIX A – FISHER DISCRIMINANT PCA ANALYSIS ALGORITHM

To analyze how shape varies between male and female populations of femora, a novel method for automatic feature generation and ranking was developed [5]. This method utilizes the power of PCA both as a means of variable reduction and as a global shape descriptor. This method is designed to find points of high discrimination between male and female femora when normalized against the first Principal Component, which is considered primarily scale. The method relies on weighting the eigenvectors using the Fisher Discriminant Ratio (FDR) defined in (8). The following algorithm shows how the FDR is applied to the PCA coordinates in a summed vector fashion. Algorithm 3 describes the steps necessary to perform this analysis:

Algorithm A-1: Fisher Discriminant Ratio for PCA comparison

1. PCA is performed over all models (both male and female femora). The PCA coordinates are recorded for each class in the matrices $\mathbf{Z}^M$ and $\mathbf{Z}^F$, where rows represent observations and columns are the principal component variables in decreasing order of significance.
2. By computing the means of the columns of $\mathbf{Z}^M$ and $\mathbf{Z}^F$, the group mean vectors are found μ^M and μ^F of length p . Similarly the standard deviations vectors are computed to be σ^M and σ^F of length p .
3. For each of the elements in μ^M , the sign of the mean is computed according to:

$$s_j = \text{sign}(\mu_j^M) \quad 1 \leq j \leq p \quad (89)$$

where s_j can be -1, 0, or +1 and describes the directionality of the j th eigenvector relative to the class means.

4. Next the Fisher Discriminant Ratio is calculated per:

$$d_j = \frac{(\mu_j^M - \mu_j^F)^2}{(\sigma_j^M)^2 + (\sigma_j^F)^2} d_j = \frac{(\mu_j^M - \mu_j^F)^2}{(\sigma_j^M)^2 + (\sigma_j^F)^2} \quad 1 \leq j \leq p \quad (90)$$

where d_j represents the power of the j^{th} principal component to discriminant between the classes of male and female.

5. Subsequently, the weightings are applied to the eigenvectors:

$$F = \sum_{j=1}^p s_j d_j U_j \quad (91)$$

where F is a column vector of length $3N$ and s_j ensures that the summation of the discrimination factor d_j across the p leftmost columns of U consistently point to the same class. Since the eigenvectors contained in U are orthogonal in $\mathbb{R}^{3N}$, but in general do not consistently point towards the same class, it is better to sum the weighted eigenvectors so that consistent differences between the classes are retained.

6. Finally, the column vector F is reinterpreted as a 3D deviation vector for each of the N model points by reorganizing the vector as an $N \times 3$ matrix. The magnitudes of the deviation vectors provide an intuitive mapping of highly discriminating between the classes.

It is worth noting that this procedure highlights areas on models that would be highly discriminating without the use of any other information. That is, if a forensic scientist or anthropologist can accurately locate landmarks identified by this algorithm, the points themselves should provide adequate discrimination without the use of any other landmarks. In the interest of this study, it was desired to examine femoral shape differences independent of the

well-known size differences between male and female femora. Since the first principal component represents most of the variance of femora and is usually considered a scale parameter, an alternative of equations (89)-(91) was formed, where the index j was defined only between 2 and p . This assisted in the location of discriminating areas on the surface of the femur if scale was not considered. Since the top ten PCs contribute to over 99% of the cumulative variance, in these studies p was set to ten.

APPENDIX B – GENETIC PROGRAMMING VERILOG MODULE

GENERATION

```

/////////////////////////////////////////////////////////////////
module BaseStationModule2(
    clk,
    reset,
    pushA,
    S0,S1,S2,S3,S4,
    Thresh,
    PEAKA);
//system-wide characteristics
parameter width=10;
parameter clk_width=32;
parameter hold_length=25000; //exactly half of one period
parameter noise_hold=10000; //period to sample noise
parameter max_hold_length =75000; //period to keep the same threshold
input clk;
input reset;
input pushA;
input [width-1:0] S0;
input [width-1:0] S1;
input [width-1:0] S2;
input [width-1:0] S3;
input [width-1:0] S4;
input [7:0] Thresh;
output PEAKA;           //this is the final output of the peak detection

//wires to describe propagation of intermediate variables up the expression
//tree
wire [width-1:0] S5;
wire [width-1:0] S6;
wire [width-1:0] S7;
wire [width-1:0] S8;
wire [width-1:0] S9;
wire [width-1:0] S10;
wire [width-1:0] S11;
wire [width-1:0] S12;
wire [width-1:0] S13;
wire [width-1:0] S14;
wire [width-1:0] S15;
wire [width-1:0] S16;
wire [width-1:0] S17;
wire [width-1:0] S18;
wire [width-1:0] S19;
wire [width-1:0] S20;
wire [width-1:0] S21;

```

```

wire [width-1:0] S22;
wire clk,reset,clk_slow;
wire reset_outA,reset_outB;
wire valid_outA;
wire PEAKA;
wire RST;
wire [width-1:0] tA;
assign tA = { Thresh,2'b11};
wire readyA;
reg startA;
reg stateA;
parameter up = 1'b1,
          down = 1'b0;

//initial block always gets called first, at FPGA power up.
initial begin
    startA <= 0;
    stateA <= 0;

end

//this block runs continuously
always @(posedge clk) begin

    //synchronous reset
    if(reset) begin
        startA <= 0;
        stateA <= 0;
    end

    //simple two-state structure, to reset GP code after each
    // detected peak. When startA goes high, followed by low, all
    // GP sub-modules get reset.
    case(stateA)
    up: begin
        if(PEAKA == 0) begin
            startA <= 1;
            stateA <= down;
        end
    end
    down: begin
        startA <= 0;
        if(PEAKA == 1) stateA <= up;
    end
    endcase
end

////////// GP MODULES //////////
// general module syntax:
// ModuleType #(parameter) InstanceName(clk,reset,input0,input1,output);
//////////

```

```

first #(width)      Module0(clk,startA,S0,S5);
sq #(width)         Module1(clk,startA,S1,S6);
maxwindow #(width) Module2(clk,startA,S1,S3,S7);
sq #(width)         Module3(clk,startA,S0,S8);
mult2 #(width)      Module4(clk,startA,S3,S3,S9);
sub2 #(width)       Module5(clk,startA,S5,S0,S10);
sq #(width)         Module6(clk,startA,S6,S11);
sq #(width)         Module7(clk,startA,S7,S12);
sub2 #(width)       Module8(clk,startA,S8,S4,S13);
dot2 #(width)       Module9(clk,startA,S10,S11,S14);
sq #(width)         Module10(clk,startA,S12,S15);
intOr2 #(width)     Module11(clk,startA,S13,S9,S16);
ema #(width)        Module12(clk,startA,S14,S2,S17);
sub2 #(width)       Module13(clk,startA,S15,S3,S18);
sq #(width)         Module14(clk,startA,S16,S19);
delay #(width)      Module15(clk,startA,S17,S18,S20);
mult2 #(width)      Module16(clk,startA,S1,S19,S21);
mult2 #(width)      Module17(clk,startA,S20,S21,S22);

assign valid_outA = (S8>tA);
HoldFor #(hold_length) holdA(valid_outA, clk, reset, PEAKA);
////////// End of GP MODULES //////////
endmodule

```

APPENDIX C – C++ CODE FOR GENETIC PROGRAM SIMPLIFICATION

```
vector<Tree*> Tree::compactify(){

    vector<TreeNode*> uniqueList;
    vector<Tree*> subTrees;
    vector<string> symbolNames;

    Tree* newTree = this->clone();

    int height = newTree->getHeight();
    int numNodes = newTree->getTotalSubNodes();

    /*
     * First iterate over all heights, let's start with the smallest nodes
     * of height = 1. After finding each unique node, we'll store them for use
     * in a new set of trees with symbolics for each unique tree we've found
     * thusfar.
     */
    for(int h=1;h<=height;++h){
        for(int kk=0;kk<numNodes;++kk){
            TreeNode* tn=newTree->getTreeNodeIndex(kk);
            if(newTree->getHeight(tn) == h){
                //start with the hypothesis that this node is unique
                bool isUnique=true;
                /* iterate over all of previously found nodes,
                 * see if this one is unique
                 */
                for(int j=0;j<uniqueList.size();++j){
                    if( uniqueList[j]->isEquivalent(tn) ){
                        isUnique = false;
                    }
                }

                /* if this node is unique, then let's give it a name and
                 * push it onto the unique list
                 */
                if(isUnique){

                    Tree* t = new Tree();
                    t->replaceTreeNodeAt(tn->clone(),0);
                    subTrees.push_back(t);

                    //create a symbol name add it to a corresponding list
                    int numEle = uniqueList.size();
                    ostringstream oss;
                    oss << "S" << numEle;
                    symbolNames.push_back(oss.str());

                    //add this node to the unique list
```

```

                                uniqueList.push_back(tn);
                                }
                            }
                        }

for(int k=0; k< subTrees.size(); ++k){
    simplify(subTrees[k]->getHead(), symbolNames, uniqueList);
}

delete newTree;


return subTrees;
}

void Tree::simplify(TreeNode* node, vector<string> symbolNames, vector<TreeNode*>
uniqueList){
    int numChildren = node->getNumChildren();

    //only want to perform this action if this node actually has children
    // the case S1 = Q needs no simplification.
    if(numChildren > 0){
        for(int i=0;i<numChildren;++i){
            TreeNode* tn=node->getChild(i);
            bool isUnique = true;
            for(int j=0;j<uniqueList.size();++j){
                if(uniqueList[j]->isEquivalent(tn)){
                    TerminalFunction* tf = new TerminalFunction (
                        symbolNames[j], NULL);
                    TreeNode* tn2 = new TreeNode(tf);
                    TreeNode* oldChild = node->replaceChild(tn2,i);
                    delete oldChild;
                    isUnique = false;
                }
            }
            if(isUnique){
                simplify(tn,symbolNames,uniqueList);
            }
        }
    }
}

```

VITA



Brandon Cole Merkl was born in Wheat Ridge, CO, on September 10, 1980. Brandon received the B.S. degree in electrical engineering and the B.S. degree in computer science from the Colorado School of Mines, Golden, CO, in 2004. Since 2005, he has worked as a researcher for the Center for Musculoskeletal Research at the University of Tennessee, Knoxville, TN. He is a founding member of Sapientia Technologies, Inc. He has published work in the fields of image processing, forensic anthropology, biomedical engineering, and clinical orthopedics. His research interests consist of such disparate topics as signal processing, high performance computing, machine learning, anatomical modeling/analysis, fuzzy-neural systems, multivariate statistics, computer vision, information theory, and nonlinear optimization. Brandon is currently a member of the IEEE Computer Society, IEEE Computational Intelligence Society, and IEEE Vehicular Technology Society.