

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Toward Self-Contained Pedestrian Inertial Navigation: Fused Silica Dual-Shell Gyroscopes and Navigation Algorithms

Permalink

<https://escholarship.org/uc/item/0rs5c44m>

ISBN

9798190948592

Author

Parrish, Austin

Publication Date

2026-09-17

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Toward Self-Contained Pedestrian Inertial Navigation: Fused Silica Dual-Shell Gyroscopes
and Navigation Algorithms

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Mechanical and Aerospace Engineering

by

Austin R. Parrish

Dissertation Committee:
Professor Andrei M. Shkel, Chair
Professor Camilo Velez Cuervo
Professor Alexandra Voloshina

DEDICATION

To my family: Angelica McDaniel, Alyssa Zabel-Parrish, Brett Parrish, Lana Parrish,
Craig Parrish, and Poppy Parrish.

TABLE OF CONTENTS

	Page
LIST OF ABBREVIATIONS	viii
LIST OF FIGURES	xii
LIST OF TABLES	xx
ACKNOWLEDGMENTS	xxi
VITA	xxiii
ABSTRACT OF THE DISSERTATION	xxviii
1 Introduction	1
1.1 Motivation	1
1.2 Navigation Methods	3
1.3 Pedestrian Inertial Navigation Literature Review	4
1.3.1 Meso-Scale Inertial Platforms	7
1.3.2 Stance-Phase Detection and ZUPT	8
1.3.3 Non-Standard Gait and Multi-IMU Aiding	9
1.4 Hemispherical Resonator Gyroscope Literature Review	11
1.4.1 Structural Asymmetry Compensation	13
1.4.2 Reduction of Energy Dissipation	14
1.4.3 Femtosecond Laser Induced Chemical Etching	16
1.4.4 Transduction of Hemispherical Resonator Gyroscopes	17
1.5 Navigation Accuracy and Gyroscope Grade for ZUPT-aided INS	20
1.5.1 Summary of Commercially Available IMUs	20
1.5.2 Navigation Error in ZUPT-Aided INS	22
1.5.3 Idealizations of the Position-Uncertainty Model	25
1.6 Problem Statement	28
1.6.1 Transduction of the Dual-Shell Hemispherical Resonator Gyroscope	28
1.6.2 Realization of a Compact, Self-Contained Navigation Platform	31
1.6.3 Navigation Across Firefighter Walking and Crawling Gaits	32
1.7 Dissertation Overview	34

2	Background	38
2.1	Pedestrian Inertial Navigation	38
2.1.1	Strapdown Inertial Navigation	38
2.1.2	Representation of Navigation Information	39
2.1.3	Coordinate Transformations	41
2.1.4	Inertial Measurements and Gravity	42
2.1.5	Velocity Update	42
2.1.6	Attitude Update	44
2.1.7	Error Model	44
2.1.8	Zero-Velocity Update (ZUPT)	45
2.1.9	Extended Kalman Filter Implementation	47
2.2	Coriolis Vibratory Gyroscopes	48
2.2.1	Allan Deviation	49
2.2.2	Lumped-Parameter Model of a Vibratory Gyroscope	50
2.2.3	Mechanization of the Coriolis Vibratory Gyroscope	53
2.2.4	Energy Loss Mechanisms	54
2.2.5	Fabrication of the Dual-Shell Hemispherical Resonator Gyroscope . .	56
I	Fused-Silica Dual-Shell Gyroscope Development	59
3	Laser-Defined Electrodes for Fused Silica Resonators	60
3.1	Fabrication Process	63
3.1.1	FLICE Process	64
3.1.2	In-Plane Electrodes Patterning	65
3.1.3	Metallization	65
3.2	Optimization of FLICE Capacitive Gap	66
3.2.1	Pulse Energy Analysis	67
3.2.2	Etching Duration Analysis	68
3.3	Capacitance Analysis	69
3.4	Device Characterization	71
3.4.1	Cross-Section of IPE-DSG	72
3.4.2	Metallization	73
3.4.3	Capacitance	74
3.5	Gyroscope Operation	75
3.6	Discussion	78
3.7	High-Capacitance Electrode Design	79
3.8	Conclusion	82
4	Non-Metallized DSG	84
4.1	Non-Metallized Driving	85
4.2	Fabrication Process	87
4.3	Non-Metallized Tuning	88
4.3.1	Mode-Shape Characterization	90
4.3.2	Direct-Rim Trimming	93

4.3.3	Teeth-Like-Tines Trimming	94
4.4	Non-Metallized Sensing	98
4.5	Gyroscope Operation	99
4.6	Conclusion	101
5	Analysis of DSG Geometry for Indirect Piezoelectric Actuation	104
5.1	Transduction Efficiency and the γ Ratio	105
5.2	Modeling Approach	106
5.2.1	DSG Designs Considered	107
5.2.2	Electrode Shaping	107
5.2.3	Rapid Analytical Finite Element Analysis Technique	110
5.2.4	Piezoelectric Analysis	112
5.3	Theoretical Analysis	113
5.4	Fabrication	114
5.5	Cleaning of PZT-Coated Dual-Shell Gyroscopes	117
5.6	Experimental Evaluation	119
5.6.1	Quality Factor Evaluation	119
5.6.2	IPA Displacement Evaluation	120
5.7	Discussion	122
5.8	Conclusion	124
II	Self-Contained Pedestrian Navigation	125
6	Meso-Scale Inertial Navigation Platform	127
6.1	System Architecture	128
6.2	Navigation Algorithm	129
6.3	Hardware Implementation and System Constraints	130
6.3.1	Computational Constraints	132
6.3.2	Power and Thermal	132
6.3.3	Packaging and Mechanical Loading	133
6.4	Performance Evaluation	134
6.5	Conclusion	134
7	Stance-Phase Detection and IMU Placement for Firefighter Crawling	137
7.1	Stance-Phase Detection of Crawling	138
7.1.1	Data Collection	138
7.1.2	Data Processing	139
7.1.3	Stance-Phase Detectors	140
7.1.4	Window Size	141
7.1.5	Parameter Search Results	141
7.1.6	Detector Statistic Comparison	142
7.2	Evaluation of Optimal Mounting Position for Firefighter Crawling	146
7.2.1	Crawling Methods	147
7.2.2	Data Collection	148

7.2.3	Data Processing	151
7.2.4	Results	152
7.3	Conclusion	153
8	Dual-IMU ZUPT-Aided Inertial Navigation for Seamless Firefighter Walking and Crawling	155
8.1	Dual-IMU Architecture	156
8.2	EKF Formulation and Kinematic Constraint	156
8.2.1	Filter Initialization	158
8.2.2	ZUPT Measurement Model	162
8.2.3	Zero-Position-Difference Constraint	162
8.3	Multi-Gait Dataset	164
8.4	Gait Classification	165
8.5	Navigation Results and Discussion	168
8.6	Conclusion	173
9	Conclusions	175
9.1	Contributions	176
9.1.1	Fused Silica Dual-Shell Gyroscope Contributions	176
9.1.2	Self-Contained Pedestrian Navigation Contributions	178
9.2	Answers to the Research Questions	179
9.3	Projected Gyroscope Performance	181
9.4	Bridging Sensor Development and Pedestrian Navigation	184
9.5	Limitations	185
9.6	Future Research Directions	186
9.6.1	Fused Silica Dual-Shell Gyroscope	186
9.6.2	Self-Contained Pedestrian Navigation	190
	Bibliography	191
	Appendix A Microfabricated Optically Pumped Gradiometer	215
A.1	Background	216
A.2	Energy Level Splitting of Nuclei	218
A.3	Optical Pumping	219
A.4	Spin Exchange	220
A.5	Electron Paramagnetic Resonator Magnetometer Dynamics	220
A.6	Microfabricated Optically Pumped Gradiometer	225
	Appendix B Lab Procedures	234
B.1	1700°C Uniform-Temperature Glassblowing Furnace	234
B.2	Lapping/Release Traveler	241
B.3	Annealing Furnace	250
B.4	Automated PAE Characterization of Dual-Shell Resonators	255
	Appendix C MATLAB Code for Dual-IMU Constraint	259

Glossary

AGC Automatic Gain Control , xv, 53, 76, 100

ALD Atomic Layer Deposition , 16, 18, 20, 83, 88, 188

AMV Acceleration Moving Variance , xviii, 8, 141–144, 151, 190

ARW Angle Random Walk , xii, xix, 23–25, 30, 49, 53, 54, 181–185

CEP Circular Error Probable , xix, 36, 37, 134, 135, 139, 141, 142, 145, 152–154, 168, 170, 171, 173, 178–180

CTE Coefficient of Thermal Expansion , 12, 14, 55, 188, 189, 238

CVG Coriolis Vibratory Gyroscope , 11, 12, 19, 34, 38, 48, 86, 188

DSG Dual-Shell Hemispherical Resonator Gyroscope , xii, xiv–xvii, xx, xxi, 12–16, 19, 20, 23, 29, 30, 33–35, 38, 49, 56–58, 60, 61, 64–67, 69, 71–75, 78, 79, 82–85, 87, 89–92, 94, 97–102, 104–124, 175–177, 179, 181, 182, 184–188, 190, 237, 238, 240, 241, 255, 260

EKF Extended Kalman Filter , xx, 8, 36, 46–48, 128, 129, 132, 134, 139, 151, 153, 156, 158, 159, 163, 173, 174, 178, 179, 181, 190, 259

FEA Finite Element Analysis , 13, 55, 110, 122

FLICE Femtosecond Laser Induced Chemical Etching , 16, 17, 34, 35, 61, 63–66, 68, 71, 72, 74, 78, 79, 82, 83, 96, 175–177, 180, 187, 189

FRB Force-to-Rebalance , 14, 53, 181, 182, 184, 189

HRG Hemispherical Resonator Gyroscope , xx, 11–23, 25, 29, 49, 55, 56, 60, 83, 84, 90, 101, 102, 123, 177, 184, 238

IMU Inertial Measurement Unit , xii, xviii–xx, 6–10, 20–24, 27, 28, 31–34, 36, 46, 126–135, 137–142, 145, 146, 148, 150–168, 171–174, 178–181, 184, 185, 190, 215, 259

INS Inertial Navigation System , 3, 4, 6, 7, 9–11, 25, 31, 34, 36, 45, 83, 102, 126, 128, 129, 134, 135, 139, 141, 151, 157, 173, 178, 179, 184–186

IPA Indirect Piezoelectric Actuation , xiv–xvii, xix, 20, 29, 35, 36, 85, 87, 89, 100–102, 104–107, 110, 111, 113, 114, 117, 119–124, 175–177, 180, 182, 183

IPE In-Plane Electrodes , xix, 34, 61, 64–67, 69, 71–75, 78, 82, 83, 175, 176, 182, 183, 188

LDV Laser Doppler Vibrometer , xiv, xv, xvii, xix, 35, 85, 90–92, 94, 98–103, 118–121, 123, 177, 180, 182, 183, 255–258

MEMS Microelectromechanical Systems , xx, 4, 8, 13, 17, 19–23, 25, 48, 79, 150, 190

MZI Mach-Zehnder Interferometer , 19, 123, 177

PAE Principal Axis of Elasticity , 90, 92, 93, 116, 255

PCB Printed Circuit Board , xvii, 64, 99, 121, 128, 130, 131

PE-ALD Plasma-Enhanced Atomic Layer Deposition , 66, 73, 75, 83, 188

PHT $\text{PbHf}_x\text{Ti}_{1-x}\text{O}_3$, xv, 20, 35, 88, 100, 101, 113

PIC Photonic Integrated Circuit , xix, 35, 124, 177, 183

PLL Phase Locked Loop , xv, 76, 90, 92, 93, 100, 255, 257, 258

PPE Personal Protective Equipment , 9, 148, 164, 173

PZT $\text{PbZr}_x\text{Ti}_{1-x}\text{O}_3$, xvi, xvii, xx, 20, 35, 106–109, 112, 115–122, 124, 177

Q-factor Quality Factor , xvii, xx, 12, 14–18, 20, 29, 30, 34, 35, 48, 49, 53, 54, 61, 71, 72, 83, 85, 86, 92–94, 96, 97, 102, 105, 110, 111, 117–120, 122, 124, 176, 177, 180–183, 185–187, 189, 255, 258

RAFT Rapid Analytical Finite Element Analysis Technique , xvii, 35, 105, 106, 110, 120, 122, 124, 176, 177

RMSE Root-Mean-Squared Error , 9, 36, 134, 141, 142, 152–154, 168, 171, 173, 179

RRW Rate Random Walk , 24, 25, 50, 185

SHOE Stance Hypothesis Optimal dEtection , 8, 46, 129, 140–142, 151, 153, 190

SNR Signal-to-Noise Ratio , xv, 14, 35, 73, 101, 118, 123, 188

SWaP+C Size, Weight, Power, and Cost , xii, 4, 7, 10, 20, 22–25, 28, 184

TED ThermoElastic Damping , 14, 55, 84, 189

WA Whole-Angle , 12, 14, 19, 48, 52, 54, 84, 86, 189

ZPD Zero-Position-Difference , xviii, xx, 10, 36, 156, 157, 159–164, 168, 171–174, 179, 181

ZUPT Zero-Velocity Update , xii, 4, 6, 7, 9–11, 25, 26, 28, 31–34, 36, 45, 46, 83, 102, 126, 128, 129, 134, 137–139, 141, 142, 145, 151, 153–155, 160, 161, 163, 165, 171–175, 178, 179, 181, 184, 185

LIST OF FIGURES

	Page
1.1 Summary of common navigation and aiding techniques and their ability to operate without prior infrastructure, indoors, and in fire conditions (smoke, high temperature, noise), together with their suitability for use in firefighting scenarios, which requires all three.	5
1.2 Comparison of commercially available Inertial Measurement Units (IMUs) by performance grade, showing Angle Random Walk (ARW) against Size, Weight, Power, and Cost (SWaP+C). Numeric labels correspond to the entries in Table 1.1. Inspired by Lutwak [1].	24
1.3 Perpendicular position uncertainty $\sigma_{\perp}$ over time for Zero-Velocity Update (ZUPT)-aided pedestrian INS, evaluated from Eq. (1.2) for three gyroscope grades (consumer, tactical, and near-navigation), together with the parallel uncertainty $\sigma_{\parallel}$ from Eq. (1.1), which is common to all grades. Simulation parameters are listed in Table 1.2.	26
2.1 Illustration of i-, e-, and n- coordinate frames (adapted from [2]), where P denotes the position of the navigation system, O is the center of mass of the Earth, and ω_{ie} denotes the rotation of the Earth.	40
2.2 Schematic of log-log plot of Allan Deviation.	50
2.3 Illustration of the lumped 2-DoF mass resonator used to represent gyroscopes.	50
2.4 Fabrication process of the Dual-Shell Hemispherical Resonator Gyroscope (DSG): (a,b) cavity pre-etching, (c,d) plasma-assisted bonding of the triple wafer stack, (e) high-temperature glassblowing, (f) release of the inner shell by lapping. At right, the released resonator assembled to a planar electrode substrate for electrostatic transduction.	58
3.1 Comparison of DSG assembly methods: (a) eutectic bonding to separate electrode substrate (prior art) [3], and (b) FLICE-patterned in-plane electrodes in the bulk of the fused silica (this work).	61

3.2	Fabrication process of IPE-DSG [4]: (a/b) cavity pre-etching of cap and substrate wafers, (c/d) triple-stack plasma-assisted wafer bonding and annealing, (e) glassblowing in a high-temperature furnace at 1700 °C, (f) substrate removal by lapping until 300–400 μm of material remains before the second cavity is released, (g) FLICE patterning of KOH release trenches and electrical isolation with KOH etching duration of >10 hours, (h) FLICE patterning of capacitive gap followed by KOH etching for approximately 6.5 hours, (i) metallization with Pt via PE-ALD, (j) electrical isolation using laser ablation of PE-ALD, and (k) wire bonding to actuation PCB. Note: steps (g-j) are a zoomed-in view of the red box in step (f).	62
3.3	Cross-section of partially released dual-shell resonator illustrating the region where the electrodes are formed with the FLICE process.	63
3.4	Fully Released IPE-DSG with features labeled. A close-up view of the capacitive gap is provided on the right illustrating the inner shell, capacitive gap, and electrical isolation. The perpendicular lines for electrical isolation are patterned to ensure alignment between the two separate FLICE patterning steps (steps (g) and (h) of Fig. 3.2)	64
3.5	Illustration of FLICE patterning a 2D channel from bottom to top. The laser follows a serpentine pattern stepping in the z-direction in 6 μm increments. The cross-sections discussed in Section 3.2 were prepared by lapping the material, as indicated in green.	67
3.6	SEM cross-sectional images of 2D channels patterned in 300 μm -thick fused silica samples using varying laser pulse energies. All samples were etched in KOH for 6.5 hours. Cross sections were prepared as described in Fig. 3.5.	68
3.7	Measured gap widths for 2D channels formed in three 300 μm -thick fused silica samples as a function of laser pulse energy. Each sample was measured at three locations, resulting in a total of nine data points. The plot displays the mean and standard deviation for each pulse energy. Only one sample fully released for the lowest pulse energy and the data reported for the center corresponds to this sample.	69
3.8	(Top) Cross section of six 300 μm -thick fused silica samples prepared with 0.54 μJ pulse energy with KOH etching time ranging from four to fourteen hours in two hour increments. (Bottom) Summary of gaps observed for three laser pulse energies: 0.540 μJ , 0.612 μJ , and 0.684 μJ . Each etching time is a separate sample. Note: no channels patterned with 0.684 μJ pulse energy and 4 hours etch time fully released.	70
3.9	Analytical capacitance of the IPE-DSG as a function of capacitive gap and electrode thickness.	71
3.10	Experimental Q-factor measurement of sample A to determine the Q-factor of the n=2 wineglass mode as-fabricated.	73
3.11	Cross-section of a partially released dual-shell resonator illustrating the in-plane electrode design. The inner shell is shown on the left.	74
3.12	SEM images of the left and right side cross-sections. The three images on each side illustrate the top, middle, and bottom of the capacitive gaps. . . .	75

3.13	Measured capacitance of each of the 24 electrodes of sample A. The estimated coplanar capacitance has been removed from these measurements. Red illustrates capacitance with Cr-Au sputtering and Blue with PE-ALD of Pt. The dashed lines represent the average values of 102fF (Cr-Au) and 187fF (PE-ALD of Pt).	76
3.14	Schematic of control loops implemented on the Zurich HF2LI Lock-in Amplifier. AGC and PLL were established along the drive axis (indicated as X-axis) for gyroscope operation.	77
3.15	Characterization of the scale factor of sample A by applying a sinusoidal excitation at a frequency of 0.05Hz and an amplitude of $3.1[^\circ/s]$. The resulting scale factor is $0.07 [mV/(^\circ/s)]$	77
3.16	Projection of scale factor improvement based on optimized FLICE parameters from Section 3.2. The star indicates the measured scale factor of sample A using dimensions from Fig. 3.12. All values are normalized to sample A's measured scale factor.	78
3.17	Comparison of electrode configurations for the DSG: out-of-plane electrodes, single-cut in-plane electrodes, and pseudo-parallel plate in-plane electrodes. The top row illustrates the design, the middle row the corresponding capacitance model, and the bottom row the calculated capacitance as a function of capacitive gap, shown for electrode thicknesses of 100 μm , 200 μm , 300 μm and 400 μm in the two in-plane cases.	81
4.1	Schematic of the Indirect Piezoelectric Actuation (IPA)-driven DSG. The outer shell is coated with a piezoelectric film and selectively excited using patterned electrodes. Stress applied to the outer shell is transferred to the vibrating element, the inner shell, through the junction, introducing a directional standing wave in the $n = 2$ wineglass mode. The film and electrode materials used in this chapter are given in Section 4.2; those used in Chapter 5 differ.	87
4.2	Fabrication process for the IPA-driven DSG: (a,b) cavity pre-etching of cap and substrate wafers, (c,d) triple-stack plasma-assisted wafer bonding and annealing, (e) glassblowing in a high-temperature furnace at 1700 $^\circ\text{C}$ to form the dual-shell structure, (f) backside flattening via lapping, (g) deposition of the Ti/Pt bottom electrode and the piezoelectric film, (h) continued lapping to fully release the inner shell, and (i) selective sputtering of the outer shell electrodes through a shadow mask. The film and electrode materials differ between this chapter and Chapter 5; those used here are given in the text. Laser ablation trimming, Section 4.3, is performed between steps (h) and (i).	89
4.3	Shadow masking method for deposition of the outer electrodes, step (i) of Fig. 4.2. (a) Characterization of the $n = 2$ wineglass mode principal axis orientation using an Laser Doppler Vibrometer (LDV) with shaker excitation. (b) Application of the laser-ablated shadow mask, fabricated from cleanroom tape on a 3D-printed holder, aligned and applied conformally to the DSG outer shell. (c) Completed DSG with outer electrodes selectively deposited through the shadow mask.	90

4.4	Conceptual illustration of the mode-shape characterization setup. The DSG is mounted flat by its substrate with the rim facing upward, driven into resonance by a bulk piezoelectric stack, and probed point by point along the rim by an LDV on a three-axis positioning stage, all inside a vacuum chamber.	91
4.5	Summary of direct-rim laser-ablation trimming iterations performed on the two piezoelectrically coated samples, sample 1 (top) and sample 2 (bottom).	91
4.6	Experimental setup for semi-automated mode-shape characterization inside the vacuum chamber, realizing the arrangement of Fig. 4.4, with the DSG mounted on the three-axis stage.	92
4.7	Illustration of the direct-rim laser ablation trimming approach. The mode shape of the lower-frequency mode is obtained and laser ablation is performed at the antinodes.	95
4.8	Frequency response characterization of the two piezoelectrically coated DSG samples, trimmed directly on the rim, using shaker piezoelectric stack excitation and LDV for detection. The top/bottom figures show the frequency response of samples 1/2, respectively.	95
4.9	(Left) Cross-sectional illustration of the teeth-like-tines geometry. The nominal dimensions used in this dissertation are as follows: the thickness of the tines (T) was approximately 400 μm , the angular width of the tine (α) was approximately 5.25°, and the radial width of the tine (W) was approximately 220 μm . (Right) Illustration of tines undergoing ablation-based material removal; the trimming locations were aligned to the antinodes as detailed in Fig. 4.7.	96
4.10	(Left) Summary of trimming iterations across 17 teeth-like-tines DSGs, reducing the frequency split from an average of 36.5 Hz to 1.0 Hz. (Right) Frequency response of the DSG with the lowest achieved frequency split of 0.10 Hz. . .	97
4.11	Mean quality factor before and after laser-ablation trimming, for shells trimmed directly on the rim and for shells trimmed with teeth-like-tines. Error bars denote one standard deviation. All shells were cleaned and annealed before the post-trimming measurement.	98
4.12	Experimental setup for optical probing of the DSG sample. Both the vertical and horizontal LDVs were precisely positioned and corrected for tilt to maximize the reflected signal from the sample.	99
4.13	Schematic of the IPA-driven DSG in open-loop gyroscope configuration. One LDV is aligned to the drive axis, which is used to establish the Automatic Gain Control (AGC) and Phase Locked Loop (PLL) control loops. The other LDV is aligned to the sense axis and demodulated for the gyroscope response.	100
4.14	Frequency response probed by the drive-axis LDV for the IPA-driven DSG. The intended drive axis is observed to have ≈ 25.6 dB Signal-to-Noise Ratio (SNR) and is 9.9 dB larger than the observed sense-axis displacement. . . .	101
4.15	Open-loop gyroscope response of the $\text{PbHf}_x\text{Ti}_{1-x}\text{O}_3$ (PHT)-coated DSG to sinusoidal input rotation with a period of 4 s and an amplitude of 15.7° s ⁻¹ . .	101

5.1	Overview of DSG geometry (cross-sectional view of right half of the device as shown in Fig. 4.1). The dimensionless γ ratio used to characterize DSG designs in this study is defined as $\gamma = D/L$, where D is the distance from the base of the central stem (or the equivalent location for stemless designs) to the junction between the inner and outer shells and L is the length of the inner shell.	106
5.2	Ten DSG designs considered for IPA transduction analysis, organized into two groups of five by increasing γ ratio: stemmed designs (1s–5s) and stemless designs (1n–5n). Designs 2s/2n and 4s/4n were fabricated with $\text{PbZr}_x\text{Ti}_{1-x}\text{O}_3$ (PZT) deposited on the outer shell and experimentally characterized.	107
5.3	Two adjacent differential volume elements taken from (right) DSG outer surface and a piezoelectric layer: (top left) the excited piezoelectric material, (bottom left) DSG outer shell experiencing deformation dictated by a given mode shape. The traction forces illustrated correspond to a state of biaxial stress in the piezoelectric element and plane stress in the modal element. The coordinate frames from global X-Y-Z frame to local piezoelectric poling frame (1-2-3) are illustrated. The local 3-direction is normal to the outer surface at each point. 1- and 2-directions are defined to align with the invariant principal strains at each point. Adapted from [5].	108
5.4	(a) Distribution of β values on the outer shell of DSG 4n for the $n = 2$ wineglass mode. Red regions ($\beta = +1$) and blue regions ($\beta = -1$) indicate locations where electrodes are driven with opposite polarities. Green regions ($\beta = 0$) indicate regions where some in-plane stress cancellation occurs. (b) Illustration of $\sigma_1 + \sigma_2 - \sigma_3$. This metric identifies regions where, despite partial cancellation indicated by $\beta = 0$, significant modal force contribution remains. (c) Final electrode placement for IPA of the $n = 2$ wineglass mode. Red and blue regions denote electrodes requiring opposite voltage polarity to generate compatible stress states for the $n = 2$ wineglass mode.	110
5.5	Illustration of the contributions to the modal force from the e_{31} , e_{32} , and e_{33} components for DSG design 4n. The legend applies to all three component plots.	113
5.6	(Left) Electrode group E_1 is located closest to the axis of symmetry; E_2 is located further in the radial direction. (Right) The modal force of E_1 and E_2 for the ten DSGs considered. Assumes an applied electric field of $E = 200 \text{ kV cm}^{-1}$	115
5.7	(Left) Cross-section view of the first principal invariant of stress for a low γ ratio design (top, DSG 1n) and a high γ ratio design (bottom, DSG 5n). (Right) The relationship between the γ ratio and r_{E1} , defined as the outer radius of electrode group E_1 , is shown for all DSG designs considered.	115
5.8	Modal force (blue) and modal stiffness (orange) versus design number for stemmed (1s–5s) and stemless (1n–5n) DSG geometries. Modal force increases by a factor of $2,224\times$ comparing stemmed designs 1s to 5s, while modal stiffness increases by a factor of 14.3. Modal force values assume an applied electric field of $E = 200 \text{ kV cm}^{-1}$	116

5.9	Predicted displacement amplitude at the rim antinode for the $n = 2$ wineglass mode versus design number for stemmed (blue, 1s–5s) and stemless (orange, 1n–5n) DSG geometries. Displacement is maximized for intermediate γ ratios where the F_m/k_m ratio is optimized. Assumes $Q_m = 100\,000$ and an applied electric field of $E = 200\text{ kV cm}^{-1}$	117
5.10	Quality Factor (Q-factor) versus γ ratio for 39 DSG samples grouped into low, medium, and high γ categories, with maximum variation of $\gamma \pm 0.07$ within each group. The dataset includes 28 DSGs from [6] and 11 DSGs fabricated in this study, five of which were coated with PZT.	120
5.11	(a) Experimental setup for IPA DSG displacement characterization. The DSG was mounted on a Printed Circuit Board (PCB) in vacuum and the outer electrodes were driven via wire bonds. An LDV aligned with an antinode of the $n = 2$ wineglass mode measured displacement amplitude. (b) Time-domain displacement measurement of sample 4n under IPA with incrementally decreasing applied electric field.	121
5.12	Displacement amplitude at the $n = 2$ wineglass mode antinode versus applied electric field, 0 kV cm^{-1} to 200 kV cm^{-1} , for five PZT-coated DSG samples actuated via IPA.	121
5.13	Comparison of predicted displacement from the Rapid Analytical Finite Element Analysis Technique (RAFT) method at $Q_m = 1$ with measured displacement divided by Q_m for all PZT-coated DSGs. The experimental measurement and the modeling were both evaluated at an applied electric field of 200 kV cm^{-1}	122
6.1	“Sugar-Cube” conceptual architecture including a microcontroller, separate sensor PCBs, a Bluetooth module, battery regulation, and a 3D-printed case for stress relief and environmental protection.	128
6.2	Illustration of “Sugar-Cube” architecture platform integrated in the sole of a shoe.	129
6.3	Flowchart of the runtime framework. The framework runs in a real-time loop as new sensor data is collected.	130
6.4	The thermal profile of the “Sugar-Cube” after one hour of operation is shown with the corresponding color map.	133
6.5	Navigation Solution and Circular Error Probable for the “Sugar-Cube” board over a 125m trajectory spanning flat planes, an elevator, and stairs. The experiment was performed 10 times.	136
7.1	Diagram of the mounting positions used in the crawling experiments and crawling method. The IMU is attached to the side of the knee using the fabric of the knee-pad. The IMU is rigidly attached to the heel using a custom 3D-printed mount. The IMU used is the VectorNav VN-200 [7].	139
7.2	Parameter search for heel-mounted (a) and knee-mounted (b) IMU crawling. The window size was swept along with detector threshold values (not shown) for all 4 detectors considered. The threshold that rendered the minimum CEP over 10 crawling trials for a given window size was selected and displayed. . .	143

7.3	SHOE Detector statistic comparison with a window size of 0.05 s (40 samples). Three typical gait cycles are shown for (a) walking with heel-mounted IMU, (b) crawling with heel-mounted IMU, and (c) crawling with knee-mounted IMU. The estimated stance-phase for all figures is highlighted in red; this was derived visually for (a). In order to improve accuracy the stance-phase was derived for (b) using Acceleration Moving Variance (AMV) detector with $e^{4.51}$ and $W = 0.9$ s (720 samples) and (c) using AMV detector with $e^{-2.51}$ and $W = 0.05$ s (40 samples). Both thresholds are derived from the parameter sweep discussed in Section 7.1.5. Note the stance-phase for (b) and (c) has been smoothed to appear continuous during these periods.	143
7.4	Knee-mounted IMU crawling detector statistic comparison. Three typical gait cycles are shown for (a)/(b) ARE detector with a window size of $W = 0.07$ s (56 samples) and $W = 0.3$ s (240 samples) respectively and (c)/(d) AMV detector with a window size of $W = 0.05$ s (40 samples) and $W = 0.3$ s (240 samples) respectively. The thresholds selected are all derived from the parameter sweep. Finally (e) shows one gait cycle and highlights the detected stance-phase periods for all four detectors.	144
7.5	Final position errors from the heel and knee-mounted IMUs are shown with the optimal parameters from the parameter sweep. The heel-mounted IMU trajectory uses the ARE detector with a window size of 0.07 s (56 samples) and renders a CEP of 0.95 m. The knee-mounted IMU trajectory uses AMV detector with a window size of 0.9 s (720 samples) and renders a CEP of 0.22 m.	146
7.6	Illustration of hand/knee crawling, duck crawling, and four mounting positions for the IMU and electromagnetic position sensors used in this experiment. A close up of the IMU and electromagnetic position sensor and the corresponding magnetic source is provided on the right. The IMU is the VectorNav VN-200 [7], and the electromagnetic tracker is the Polhemus Patriot [8].	147
7.7	Final position error for all mounting positions: outside knee, center knee, inside knee, and heel. The final position is shown in red for hand/knee crawling and blue for duck crawling. Each series of experiments was performed 10 times and the resulting CEP and vertical RMSE are detailed in Tables 7.1 and 7.2.	149
7.8	Histogram of residual stance-phase velocity for the heel-mounted IMU performing hand/knee crawling. The standard deviation of the fitted normal distribution (in red) is defined as the velocity uncertainty.	150
8.1	IMU mounting positions on the foot and knee for the dual-IMU setup. Lever arms from each IMU to the ankle were measured per subject with a ruler and are expressed in the respective IMU body frames shown. The average values were $\vec{L}^s = [0 \ 0 \ 7.6]^\top$ cm and $\vec{L}^k = [-38.9 \ 0 \ 2]^\top$ cm, where the 2 cm Z^k -component accounts for the knee-pad standoff. Residual lever-arm error is absorbed by the Zero-Position-Difference (ZPD) measurement noise.	157

8.2	Temporal alignment of shoe (blue) and knee (orange) IMU data, shown as gyroscope-derived change in pitch angle $\Delta\theta$ for walking data. (Top) $\Delta\theta$ over 11 gait cycles; the alignment landmark, circled, is where the knee-pitch maximum coincides with the shoe-pitch minimum. (Bottom) The 11 cycles superimposed, each centered on the knee-pitch peak, $t = 0$	160
8.3	Sensitivity of the shoe trajectory to error in the knee mounting-yaw offset $\Delta\psi$. Each curve is a complete estimator run, identical except for $\Delta\psi$, displaced from its calibrated value by 0, ± 3 , ± 6 , ± 12 , and $\pm 24^\circ$. Trajectories are rotated to a common heading at 10 m, which removes any rigid heading offset; circles mark maximum displacement.	161
8.4	Confusion matrix for the random forest gait classifier, trained on all three subjects and applied to Subjects I, II, and III, showing accuracy above 86.4% except for transitions between walking/crawling.	167
8.5	Stance-phase detection statistic (log-likelihood) and the resulting gait-classifier output for the shoe IMU, (a), and the knee IMU, (b). Each IMU's statistic is most stable in its anchor gait, the shoe during walking and the knee during crawling. The colored bar at right encodes the classified state, and the green labels above (a) mark the true gait, with the small gaps corresponding to the walk-to-crawl transitions.	168
8.6	Multi-gait North-East trajectories over the nine trials under the three filter configurations described in the text; black stars mark the true trajectory endpoint.	169
8.7	Multi-gait vertical trajectories versus time over the nine trials for each of the three filter configurations.	169
8.8	Final horizontal position error over the nine multi-gait trials for each of the three filter configurations; red markers are the per-trial trajectory endpoints, the black star is the true endpoint, and the red circle indicates the Circular Error Probable (CEP), summarized in Table 8.4.	170
8.9	Easting error accumulation across the multi-gait protocol for all three navigation configurations. Each trajectory is heading-aligned so the corridor lies along the north axis; the plotted quantity is the median absolute cross-track error over the nine trials at five checkpoints: start, the three labeled gait transitions (walk to crawl, crawl to walk, walk to duck), and the endpoint. The shaded band shows the interquartile range of the corresponding value.	170
9.1	Projected ARW of the two transduction methods versus frequency split. (a) IPA: detection-limited ARW for chip-scale Photonic Integrated Circuit (PIC) and benchtop LDV optical floors, with the mechanical-thermal (Brownian) floor in black; the drive amplitude is scaled with Q from the measured Q-normalized displacement. (b) In-Plane Electrodes (IPE) on Sample A: electronic-noise-limited ARW estimated from the noise model of [9], evaluated at the Q-factor values and capacitances shown, with the Brownian floor in black; the drive amplitude is set by the gap and the out-of-plane to in-plane displacement ratio.	183

LIST OF TABLES

	Page
1.1 Comparison of some commercial gyroscopes by performance grade. A comparison of the gyroscope performance of reported Microelectromechanical Systems (MEMS) fabricated Hemispherical Resonator Gyroscopes (HRGs) is provided below as well.	23
1.2 Parameters used in the position-uncertainty model of Eqs. (1.1)–(1.2). Values marked “measured (this work)” were characterized by the author.	27
1.3 Summary of the research questions and the chapters in which they are addressed.	33
3.1 Experimentally measured parameters of IPE-DSG samples	72
4.1 Characteristics of the two piezoelectrically coated DSGs, trimmed directly on the rim.	94
5.1 Characterized Q-factor of PZT-coated DSG samples before and after cleaning, for the $n = 2$ and $n = 3$ wineglass modes.	118
6.1 Specifications of sensors and other components used on the “Sugar-Cube” Navigation Board.	131
7.1 Summary of velocity uncertainty and stance-phase duration across 50 gait cycles, as well as navigation accuracy for the 10 trials of length 42.6 m for the hand/knee crawling method.	148
7.2 Summary of velocity uncertainty and stance-phase duration across 50 gait cycles, as well as navigation accuracy for the 10 trials of length 42.6 m for the duck crawling method.	148
8.1 Process-noise parameters for the Extended Kalman Filter (EKF), based on an 800 Hz IMU update rate.	158
8.2 Per-gait ZPD position uncertainty (one standard deviation). The ZPD constraint is symmetric, so a single isotropic covariance $\mathbf{R} = \sigma^2 \mathbf{I}_{3 \times 3}$ is applied, with σ the per-gait value listed.	164
8.3 Random forest gait-classification accuracy (%) on the held-out multi-gait dataset by training-subject composition and sensor configuration.	167
8.4 Final-position error over the nine multi-gait trials for each filter configuration.	172

ACKNOWLEDGMENTS

I begin with a quote from my late Aunt Megan who captured the nature of undertaking a PhD: “One should never begin a project the magnitude of a dissertation without projecting the cost involved: emotional, spiritual, physical, monetary. But then one rarely appreciates how much of all of those attributes will be required to finish a dissertation.” This observation proved accurate in the course of my PhD, which would not have been completed without the vast support system I have been fortunate to have.

First and foremost, I would like to thank Professor Andrei Shkel for his unwavering support over my PhD. His encouragement, patience, guidance, and expert knowledge were essential in the course of this work.

Collaborating with Dr. Nicholas Strnad, my mentor from DEVCOM US Army Research Lab, provided one of the most rewarding opportunities of my PhD. The piezoelectric materials referenced in this dissertation were deposited by him and his team at ARL SEMASC Cleanroom.

I am grateful for all the conversations and long hours I shared with my labmates in the lab. Thank you to all lab mates: Dr. Sina Askari, Yu-Wei Lin, Dr. Radwan Noor, Dr. Mohammad Asadian, Dr. Yusheng Wang, Dr. Doreen Hii, Dr. Daryosh Vatanparvar, Dr. Chi-Shih Jao, Dr. Danmeng Wang, Wei Guan, Dr. Eudald Sengenis-Rafart, Lois Meira-Lopez, Crystal Wai, Esther Wong, Elena Wolgamot, Sina Salmani, Saeid Mozafari, and Junyong Sim. In particular, I am grateful to Dr. Radwan Noor, who met with me weekly for years after his graduation and guided me towards a successful PhD, and Dr. Daryosh Vatanparvar, who provided invaluable advice several years after his graduation. In addition, several students who were undergraduates at the time assisted in releasing HRGs, namely May Maung, Matthew Delmont, and Titus de Jong. I have also been fortunate to collaborate with members of other labs, specifically on photonic transduction of the DSG, and I want to acknowledge Melika Momenzadeh and Professor Maxim Shcherbakov, as well as Ozan Berk Boyraz and Professor Ozdal Boyraz.

Among the most rewarding experiences of my PhD was mentoring a senior design team comprising Zachariah (Zak) J Klutts, Spencer Tyler Herbert, and Matthew Delmont. This team helped create the initial design of the automated mode-shape characterization setup used for correction of DSG asymmetry in this dissertation, which has saved me hundreds of hours over the course of my PhD (~300 mode shapes collected at the time of writing this)! Zak also completed most of the Solidworks modeling for the multi-LDV setup discussed in Chapter 4.

Essential to the fabrication process development was the assistance of Integrated Nanosystems Research Facility (INRF) staff members Jake Hes, Dr. Mo Kebaili, Dr. Chuang-Yuan Lee, and Dr. Richard Chang. The staff members of UC Irvine Materials Research Institute (IMRI) also provided valuable technical support over the years.

This work was supported in part by financial assistance awards 70NANB17H192 and 70NANB21H154 from the National Institute of Standards and Technology (NIST); by the Army Research Laboratory (ARL) and the Army Research Office (ARO) under Collaboration Agreement Numbers W911NF2220069 and 2022137-142232; and briefly by the Defense Advanced Research Projects Agency (DARPA) under Grant No. N66001-16-1-4021 (PRIGM-AIMS). I also gratefully acknowledge support from the Paul and Beverly Holmes Endowed Fellowship and the Rose Hills Foundation Science and Engineering Fellowship.

Reflecting on the path that led me to pursue a PhD, I am especially grateful to a few key mentors: Kenneth Ozuna (Del Norte High School), Professor Jackie Trishman (Cal State San Marcos), Joshua Grant (UC Davis Air Quality Research Center), David Pletcher (NASA Ames Research Center), and Dr. Stephanie Fung (UC Davis).

I would also like to thank all of my friends from UC Irvine, specifically Alexander Nguyen and Donipolo Ghimire, as well as my friends at home in San Diego for their support.

Finally, I want to thank the unwavering support of my partner, Angelica McDaniel; my parents, Craig and Lana Parrish; brother, Brett Parrish; sister, Alyssa Zabel-Parrish; brother-in-law, Matt Zabel; family in Kansas and Iowa; and my partner's parents, Ana and Brian McDaniel. Thank you for listening to me talk about this endlessly for seven years, I'm done now.

VITA

Austin R. Parrish

EDUCATION:

PhD, Mechanical and Aerospace Engineering	Expected 2026
University of California, Irvine	<i>Irvine, CA</i>
M.S., Mechanical and Aerospace Engineering	2020
University of California, Irvine	<i>Irvine, CA</i>
B.S., Mechanical Engineering, Minor: Electrical Engineering	2019
University of California, Davis	<i>Davis, CA</i>

RESEARCH EXPERIENCE:

Graduate Student Researcher	2019-Present
University of California, Irvine	<i>Irvine, CA</i>
Undergraduate Student Researcher	2017-2019
University of California, Davis	<i>Davis, CA</i>

TEACHING EXPERIENCE:

Teaching Assistant - Mechanical Vibrations	2020, 2023
Teaching Assistant - Design of Computer-Controlled Robots	2021
Teaching Assistant - Introduction to Engineering Computation	2020

PROFESSIONAL EXPERIENCE:

Intern - Northrop Grumman Mission Systems	2023
Intern - NASA Ames Research Center	2018

HONORS:

Paul and Beverly Holmes Endowed Fellowship	2021-2023
Rose Hills Foundation Science and Engineering Fellowship	2021-2023

REFEREED JOURNAL PUBLICATIONS:

L. Meira-Lopez, **A. R. Parrish**, and A. M. Shkel, “Understanding and Mitigating Asymmetries in 3D Resonators Fabricated by Micro-Scale Glassblowing,” *IEEE/ASME Journal of Microelectromechanical Systems*, In preparation, 2026.

M. Momenzadeh, L. Meira-Lopez, **A. R. Parrish**, A. M. Shkel, and M. R. Shcherbakov, “Spatially Selective Optothermal Control of Near-Degenerate Modes in Hemispherical Resonators,” *APL Photonics*, In preparation, 2026.

A. R. Parrish, E. Sangenis, S. Mozafari, D. Pachori, and A. M. Shkel, “Dual-IMU ZUPT-Aided Inertial Navigation for Seamless Firefighter Walking and Crawling,” *IEEE Sensors Journal*, Under Review, 2026.

A. R. Parrish, N. A. Strnad, J. S. Pulskamp, and A. M. Shkel, “Analysis of Dual Shell Hemispherical Resonator Gyroscope Geometry for Indirect Piezoelectric Actuation,” *IEEE/ASME Journal of Microelectromechanical Systems*, Accepted, 2026.

A. R. Parrish, E. Wong, L. Meira-Lopez, and A. M. Shkel, “Self-Aligned, High-Aspect Ratio Electrodes for Fused Silica Resonators Digitally Defined by Laser Modification of Material,” *IEEE/ASME Journal of Microelectromechanical Systems*, vol. 34, no. 5, 2026.

O. B. Boyraz, L. Meira-Lopez, **A. R. Parrish**, M. S. Islam, C. Babayigit, A. Atalar, A. M. Shkel, and O. Boyraz, “PIC Based High-Precision Characterization of Fused-Silica Dual-Shell Resonators,” *Optics Express*, vol. 34, no. 3, pp. 4212-4227, 2026.

D. Wang, M. H. Asadian, W. Guan, D. Hii, **A. R. Parrish**, and A. M. Shkel, “Effect of Metallization on Fused Silica Dual-Shell Gyroscopes,” *Journal of Microelectromechanical Systems*, vol. 31, no. 6, pp. 877-887, 2022.

A. R. Parrish, C.-S. Jao, and A. M. Shkel, “Stance Phase Detection for ZUPT-aided INS using Knee-Mounted IMU in Crawling Scenarios,” *IEEE Sensors Letters*, vol. 6, no. 5, 2022.

REFEREED CONFERENCE PUBLICATIONS

M. Momenzadeh, L. Meira-Lopez, O. B. Boyraz, **A. R. Parrish**, M. R. Shcherbakov, O. Boyraz, and A. M. Shkel, “Architecture for a Micro-Scale Fused Silica Hemispherical Gyroscope with Optical Actuation and Interferometric Detection Enabled by Photonic Integrated Circuits,” Miramar Beach, FL, USA. *IEEE Research and Applications of Photonics in Defense (RAPID)*, Aug 2026.

L. Meira-Lopez, **A. R. Parrish**, and A. M. Shkel, “Chemical Etching based Trimming Process for Sub-Hz Permanent Frequency Tuning of Fused Quartz Vibratory Gyroscopes,” Hilton Head, SC, USA. *Solid-State Sensors, Actuators and Microsystems Workshop, Hilton Head*, Jun 2026.

M. Momenzadeh, L. Meira-Lopez, **A. R. Parrish**, A. M. Shkel, and M. R. Shcherbakov, “Spatially Programmable Optical Excitation of Orthogonal Hemispherical Micromechanical Resonator Modes,” Hilton Head, SC, USA. *Solid-State Sensors, Actuators and Microsystems Workshop, Hilton Head*, Jun 2026.

L. Meira-Lopez, M. Momenzadeh, O. B. Boyraz, **A. R. Parrish**, M. R. Shcherbakov, O. Boyraz, and A. M. Shkel, “Early-Stage Feasibility of Optical Excitation and Read-out Enabled by Laser Welding Packaging of a Dual-Shell Quartz HRG,” Destin, FL, USA. *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL) Conference*, May 2026.

M. Momenzadeh, L. Meira-Lopez, **A. R. Parrish**, A. M. Shkel, and M. R. Shcherbakov, “Programmable Optical Excitation of Micromechanical Resonators,” Charlotte, NC, USA. *Conference on Lasers and Electro-Optics (CLEO)*, May 2026.

O. B. Boyraz, L. Meira-Lopez, **A. R. Parrish**, M. S. Islam, C. Babayigit, A. Atalar, A. M. Shkel, and O. Boyraz, “Dual-channel PIC optical readout for modal characterization of fused-silica hemispherical resonators,” San Francisco, CA, USA. *SPIE MOEMS and Miniaturized Systems*, Jan 2026.

A. R. Parrish, L. Meira-Lopez, M. H. Asadian, E. Wong, and A. M. Shkel, “Laser Patterned Self-Aligned Electrodes for Dual-Shell Hemispherical Resonator Gyroscope,” Orlando, FL, USA. *Solid-State Sensors, Actuators and Microsystems Workshop (TRANSDUCERS)*, Jun 2025.

L. Meira-Lopez, **A. R. Parrish**, D. Wang, E. Wong, and A. M. Shkel, “From Stress to Success: Demonstrating Ultrahigh Quality Factor Dual-Shell Resonator Gyroscopes with Low Stress Topology,” Orlando, FL, USA. *Solid-State Sensors, Actuators and Microsystems Workshop (TRANSDUCERS)*, Jun 2025.

O. B. Boyraz, C. Babayigit, L. Meira-Lopez, **A. R. Parrish**, M. S. Islam, A. M. Shkel, and O. Boyraz, “Fused-Silica Dual-Shell Resonators with Photonic Integrated Readout,” Long Beach, CA, USA. *Conference on Lasers and Electro-Optics (CLEO)*, May 2025.

A. R. Parrish, N. A. Strnad, Z. J. Klutts, L. Meira-Lopez, and A. M. Shkel, “Non-Metallized Dual-Shell Hemispherical Resonator Gyroscope,” Salt Lake City, UT, USA. *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Apr 2025.

L. Meira-Lopez, D. Wang, **A. R. Parrish**, and A. M. Shkel, “A Method to Extract and Model Structural Asymmetries in Dual-Shell Gyroscopes to Study Frequency Asymmetry,” Hilton Head, SC, USA. *Solid-State Sensors, Actuators and Microsystems Workshop, Hilton Head*, Jun 2024.

A. R. Parrish, C.-S. Jao, D. Wang, and A. M. Shkel, “‘Sugar-Cube’ Navigation Board: Pedestrian Hardware Platform that Fits in the Sole of a Shoe,” Kaua’i, HI, USA. *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL) Conference*, Mar 2023.

A. R. Parrish, C.-S. Jao, D. Wang, and A. M. Shkel, “Study of IMU Mounting Position for ZUPT-Aided INS In the Case of Firefighter Crawling,” Kaua’i, HI, USA. *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL) Conference*, Mar 2023.

D. Wang, N. A. Strnad, Y. Wang, **A. R. Parrish**, R. R. Benoit, R. R. Knight, and A. M. Shkel, “Indirect Excitation of micro-HRG Using Segmented Piezoelectric ALD PHT Actuator,” Kaua’i, HI, USA. *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL) Conference*, Mar 2023.

C.-S. Jao, D. Wang, **A. R. Parrish**, and A. M. Shkel, “A Neural Network Approach to Mitigate Thermal-Induced Errors in ZUPT-aided INS,” Avignon, France. *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL) Conference*, May 2022.

C.-S. Jao, **A. R. Parrish**, and A. M. Shkel, “‘Sugar-Cube’: A Real-time Pedestrian Localization Testbed Utilizing Foot-mounted IMU/Barometer/Ultrasonic Sensors,” Virtual Conference. *IEEE Sensors Conference*, Nov 2021.

A. R. Parrish, R. Noor, and A. M. Shkel, “Microfabricated Optically Pumped Gradiometer with Uniform Buffer Gases,” Virtual Conference. *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL) Conference*, Mar 2021.

PATENTS

L. Meira-Lopez, **A. R. Parrish**, and A. M. Shkel, “Laser Assisted Fabrication for Shell Gyroscopes Made from Dielectric Materials,” UC Case No. 2027-0002, *Patent Pending*, 2026.

A. R. Parrish, M. H. Asadian, L. Meira-Lopez, and A. M. Shkel, “Laser Patterned Self-Aligned Electrodes for Hemispherical Resonator Gyroscope,” UC Case No. 2025-819-1, *Patent Pending*, 2025.

A. M. Shkel, D. Wang, and **A. R. Parrish**, “Wafer-Level Fabrication Process of Ferrimagnetic Resonator,” WO2023141492A3, *Patent*, 2025.

ABSTRACT

Toward Self-Contained Pedestrian Inertial Navigation: Fused Silica Dual-Shell Gyroscopes
and Navigation Algorithms

By

Austin R. Parrish

Doctor of Philosophy in Mechanical and Aerospace Engineering

University of California, Irvine, 2026

Professor Andrei M. Shkel, Chair

This Ph.D. dissertation reports on the development of Coriolis vibratory gyroscopes and pedestrian inertial navigation algorithms, toward self-contained localization in environments where satellite signals and pre-installed infrastructure are unavailable. The motivating application is the localization of firefighters during structure fires, where dense smoke, elevated temperature, and no prior knowledge of the building rule out most navigation methods. The rapid drift of a dead-reckoning INS can be curtailed by the Zero-velocity UPdaTe (ZUPT), which estimates the inertial sensor biases during the stance-phase of the gait cycle. The accuracy of ZUPT-aided INS is therefore bounded by two complementary limitations: the noise performance of the yaw-axis gyroscope, which ZUPT cannot correct and sets the ceiling on achievable accuracy, and the reliability of the zero-velocity updates and system integration, which govern how closely that ceiling is met. This dissertation addresses these limitations as parallel advances toward a future integrated system rather than by building that system itself, and is organized in two parts: sensor development and system-level navigation.

The first part advances transduction methods for the Dual-Shell Hemispherical Resonator Gyroscope (DSG). We developed a method to define electrodes for capacitive transduction within the bulk of the fused-silica of the resonator itself using Femtosecond Laser Induced Chemical Etching (FLICE), eliminating the separately fabricated electrode substrate and

the assembly step it requires, and achieving bulk capacitive gaps as small as $5.5\text{ }\mu\text{m}$ on the DSG at a maximum as-fabricated quality factor of 4.57 million. We then demonstrated, to the best of our knowledge for the first time, operation of a microfabricated hemispherical resonator gyroscope as a rate gyroscope with no conductive film on the resonating element, replacing the drive, tune, and sense functions of metallization with indirect piezoelectric actuation, laser ablation trimming, and optical detection, and thereby avoiding the quality-factor degradation that metallization imposes. Finally, we established how the dual-shell geometry governs the efficiency of this indirect actuation, developing a modeling framework and characterizing five samples that together demonstrate a $6.5\times$ improvement in displacement amplitude and a $12.3\times$ improvement in quality-factor-normalized transduction efficiency over non-optimized designs. Each method was operated as a rate gyroscope, though none were incorporated into a navigation platform.

The second part advances system-level navigation. We built a self-contained inertial navigation platform small enough to be embedded in the sole of a shoe, occupying 7.1 cm^3 excluding battery and case, on which barometer-aided ZUPT-aided INS achieved a Circular Error Probable (CEP) of 0.66 m over a 125 m indoor walking trajectory. We then extended ZUPT to the crawling gaits of firefighter smoke-diving operations, for which the stationary interval assumed by conventional foot-mounted implementations is unreliable, by fusing shoe- and knee-mounted solutions in a joint filter coupled through a rigid-leg kinematic constraint and configured online by a random forest gait classifier. Over nine multi-gait trials with three subjects in full firefighter equipment, this reduced the horizontal CEP from 3.86 m to 0.58 m and the vertical root-mean-square error (RMSE) from 3.06 m to 1.26 m.

Although developed toward a common goal, the two parts of this dissertation were not combined into a single experimental system: the DSG was not integrated into either navigation platform, and the navigation studies were conducted with commercial inertial measurement

units (IMUs). The sensor development raises the achievable accuracy ceiling set by the yaw-axis gyroscope; the navigation developments improve how closely it is approached.

Chapter 1

Introduction

1.1 Motivation

Many navigation solutions depend on GNSS or on other RF signals such as LTE, AM/FM radio, digital television, or radar [10]. When available, these signals provide position fixes referenced to an absolute frame, typically accurate to 4.9 m for a consumer cell phone [11], and the resulting error does not accumulate over time as it does in dead reckoning [12]. However, these external signals are less effective underground or indoors, where signal attenuation and multipath degrade accuracy or prevent position updates entirely. In addition, the update rate of these methods may be insufficient to capture the dynamic motion of highly maneuverable platforms or of the human body.

Methods developed to improve navigation accuracy in these environments are primarily based on beacons such as wireless local area networks, UWB, Bluetooth, and RFID. These solutions have been commercialized by companies such as CloudNav [13] and SmartPoint [14] for applications including asset tracking and customer engagement. While such methods have demonstrated accuracy on the order of 1 m, all reported methods depend on prior infras-

structure, and many are susceptible to jamming and spoofing [15] or to attack [16]. When infrastructure is available it should be used to enhance the navigation solution. However, in mission-critical scenarios, the navigation solution requires components that are always available and therefore self-contained. Examples of mission-critical scenarios in which prior infrastructure may not be available include urban warfare, law enforcement searching facilities, firefighters performing smoke diving in buildings, and emergency response during humanitarian crises [17].

Among these, the localization of a firefighter during a structure fire is the application that motivates this dissertation. In the United States in 2024, 16,275 firefighter injuries were reported at the fireground [18], and 81 % of fireground injuries between 2016 and 2020 occurred at structure fires [19]. Of the fires causing five or more firefighter deaths between 1977 and 2010, 65 % involved structures such as industrial plants, stores, or warehouses [20]. One of the most dangerous activities firefighters perform is smoke diving, where a firefighter enters a building filled with smoke, necessitating the use of a SCBA. In these situations the smoke-diving leader, who monitors the position of the firefighters, must assess the risk level of an operation, and doing so depends on accurate knowledge of each firefighter’s position. This information is crucial in deciding where to send additional units, determining which areas have not yet been searched, and locating injured firefighters in need of rescue [21, 17]. Localization also assists the firefighter in the disorienting, high-stress conditions endured in burning buildings, including darkness, smoke, flames, and noise, as well as the physiological toll of maneuvering through rubble [22]. It can help firefighters locate exits and teammates and avoid searching the same area for survivors twice [21]. Time is of the essence in these scenarios, and localization information makes firefighter response more efficient [23, 24]. Therefore, a localization system designed for firefighters must not depend on prior knowledge of the building or its surroundings, must not require signals from pre-installed infrastructure, and must remain operable in environments filled with airborne particles, smoke, and mist [25, 26, 27].

1.2 Navigation Methods

Many sources of aiding have been used for strapdown Inertial Navigation Systems (INSs), and many reviews of these methods are available in the literature [28, 29, 30, 31, 32]. INSs are frequently used to bridge the outage of other sensors or to aid signal acquisition, as in magnetic or vision-based navigation. One exception to the reliance on pre-installed beacon infrastructure is the recently proposed use of simultaneously localized UWB beacons, in which case inertial navigation is still required both to localize the beacons initially and to navigate between them [33].

Navigation and aiding methods that do not depend on prior infrastructure include vision [34], inter-agent ranging [35, 36], and ultrasonic methods [37, 38, 39]. The primary limitation of ultrasonic ranging sensors is that they require a flat surface for measurement; walls may be angled, or, if the sensor is attached to the shoe, the sensor itself may be angled relative to the wall. Altimeters are effective in estimating the absolute height of a navigation platform from air pressure. Magnetometers have been used to sense magnetic north, but indoors the magnetic field created by the structure of the building skews the measurements. The principal indoor approach is instead magnetic SLAM, in which the ambient magnetic field is mapped and subsequently matched to constrain drift. Magnetic SLAM can aid an inertial solution but cannot provide a navigation solution on its own.

Finally, it is necessary to consider the indoor environment itself in mission-critical navigation cases. For firefighters entering buildings actively on fire, the environment is filled with smoke and steam and the temperature is very high [17, 40, 29]. Several of the self-contained navigation methods described above will have reduced accuracy or will be non-functional under these conditions. A study conducted by Starr et al. [20] demonstrated that LiDAR in dense smoke is less accurate at 4 m smoke visibility and inoperable at 1 m visibility. In general, sensors operating at wavelengths longer than $7.5\text{ }\mu\text{m}$ were found to remain functional in smoke, which eliminates cameras operating in the visible range but permits infrared

cameras. In addition, altimeters provide incorrect height estimates, because rapid heating changes the ambient air pressure [41].

A summary of the navigation and aiding methods is presented in Figure 1.1, indicating whether each is infrastructure free, operable indoors, and operable under fire conditions. Of the methods surveyed, only ZUPT-aided INS, magnetometer aiding, infrared cameras, and ultrasonic ranging remain operable in firefighting conditions. Of these, only ZUPT-aided INS produces a complete navigation solution without prior infrastructure; the remainder aid an existing solution rather than replace it. To date, no system has demonstrated robust, self-contained navigation capable of 1 m accuracy over 15 min of navigation in a mobile, low-power, and relatively inexpensive package. PNSs based on inertial sensing nonetheless remain the most promising route to that capability: they are self-contained, and, owing to MEMS fabrication, they minimally impede firefighter movement [42, 43]. The long-term objective of this line of research is a single microfabricated assembly that integrates multiple navigation methods and aiding sources on the same die, minimizing SWaP+C [44, 45]. This dissertation advances that goal along two complementary fronts rather than by building the integrated assembly itself: the sensor-level drift that bounds long-duration accuracy is addressed through development of a MEMS-fabricated fused-silica gyroscope, and the system-level barriers of packaging and gait-robust estimation are addressed through pedestrian-navigation hardware and algorithms. The two are joined by the noise performance of the yaw-axis gyroscope, which sets the upper bound on the accuracy a ZUPT-aided pedestrian INS can reach and which the sensor work is directed at improving.

1.3 Pedestrian Inertial Navigation Literature Review

An INS comprises three gyroscopes and three accelerometers. Navigation based purely on these sensors is referred to as dead reckoning. For consumer-grade low-cost sensors, the navigation solution will drift on the order of a few meters within seconds without external

Navigation/Aiding Technique	Infrastructure Free	Operable Indoors	Operable in a Fire	Suitable for all Firefighting Scenarios (Infrastructure Free, Indoors, Fire)
GPS/GNSS			☑	
LTE, Radio, Radar, etc.			☑	
WLAN/WiFi • Fingerprinting • Signal Strength		☑	☑	
Beacons (UWB, RFID, Bluetooth, etc.)		☑	☑	
LIDAR	☑	☑		
Altimeter	☑	☑		
Vision (IR, Thermal)	☑	☑	☑	☑
Ultrasonic Ranging	☑	☑	☑	☑
IMU w/ ZUPT	☑	☑	☑	☑
Magnetometer • MARU • Fingerprinting • Odometry	☑	☑	☑	☑

Figure 1.1: Summary of common navigation and aiding techniques and their ability to operate without prior infrastructure, indoors, and in fire conditions (smoke, high temperature, noise), together with their suitability for use in firefighting scenarios, which requires all three.

aiding [46]. A primary reason for the rapid error growth is that attitude error causes a component of gravity to be projected into the horizontal accelerometer channels, where it is then twice integrated into position. For this reason, ZUPT has been an essential aiding element in pedestrian inertial navigation since the seminal paper by Foxlin in 2005 [47, 48]. ZUPT is typically used only in pedestrian navigation scenarios, and specifically in walking scenarios, since it leverages the brief stationary period during the gait cycle in which the foot is on the ground, termed the stance-phase. During the stance-phase the IMU is assumed to have zero translational velocity, and IMU biases are estimated and corrected, greatly reducing navigation error.

The implementation of ZUPT is non-trivial, and to date it has been realized in only a limited number of commercial systems [49, 50], even though patents on the subject have existed since the 1990s [48]. One of the primary reasons is the difficulty of robustly detecting the stance-phase. While this has been studied previously [51], the detection parameters depend strongly on the mode of motion and on how an individual walks. Recent methods incorporate gait speed or frequency into stance-phase detection [52, 53], and the detection of motion type, for example walking or running, using machine learning in order to adjust the stance-phase detection parameters has also become a research focus [54].

Two limitations bound the accuracy that ZUPT aiding can deliver, and both are addressed in this dissertation. The first limitation is fundamental to ZUPT. While the roll and pitch gyroscope biases are observable during the stance-phase, the yaw angle bias is not [48, 55]. Without additional aiding, in-plane drift is therefore governed by the noise performance of the yaw-axis gyroscope, and improving long-duration pedestrian navigation accuracy reduces to improving that sensor. Inertial sensors must additionally provide sufficient dynamic range and bandwidth to capture pedestrian motion while maintaining noise performance [56, 57].

The second is the fragility of the stance-phase assumption under the locomotion methods firefighters actually use. While the navigation performance of the ZUPT-aided INS has been

evaluated in many pedestrian activities, including walking, running, and crawling [48], crawling has received comparatively little attention. During smoke-diving operations, firefighters can be forced toward the floor to escape heat from the smoke and fire, and the floor can also offer improved visibility. When an entire room is filled with smoke, inhibiting visibility, and firefighters must search for survivors, it is also common to crawl and spread out the body so as to cover as much surface area as possible. Since the ZUPT algorithm is predicated on a brief stationary period in every gait cycle, it is imperative to ensure accurate stance-phase detection and to account for the change in dynamics during crawling, as the stationary period may differ in stability and in duration [25]. It has been demonstrated that ZUPT-aided INS accuracy can be improved by optimizing the mounting position of the IMU in the case of walking [58, 59], since mounting position affects stance-phase duration, stance-phase stability, and the shock level experienced by the IMU. There is a corresponding need to explore this question for firefighter crawling.

Realizing such a system also requires that the inertial sensors, aiding sensors, computation, power, and communication be packaged in a form factor that can be worn. Minimizing size favors COTS sensors and a compact microcontroller development board, and exploring the issues associated with packaging a navigation system small enough to fit in the sole of a shoe is the purpose of the meso-scale platform used in this work [38]. Such a platform is a step toward a fully microfabricated navigation assembly. The strapdown inertial navigation formulation and the ZUPT-aided estimation framework on which the later chapters build are reviewed in Chapter 2.

1.3.1 Meso-Scale Inertial Platforms

Pedestrian INSs are well suited to firefighter localization because they are self-contained, require no external infrastructure, and impose minimal physical burden on the wearer [42, 43]. Practical deployment is nonetheless governed by SWaP+C [17]. The miniaturization of IMUs

enabled by MEMS has made shoe- and body-mounted platforms feasible [43], and open-source multi-IMU platforms have been developed to support such wearable configurations [60]. Flexible platforms optimized for rapid integration of additional aiding sensors have also been reported [61, 38], though these are generally too large for integration into the sole of a shoe.

Realizing a fully self-contained navigation platform within the sole of a shoe introduces several coupled challenges. The computational load of the EKF must be minimized to sustain a sufficient update rate; prior foot-mounted implementations omitted IMU bias states from the filter specifically to achieve higher update rates [42]. Concentrating this computation into a small package generates heat, which induces bias in the inertial sensors if left uncompensated [62]. This is compounded in firefighting scenarios, where ambient temperature around the responder can change substantially within a short time span [63]. The power draw of the processing unit further constrains achievable battery life, placing it in direct tension with the update rate the filter requires. Packaging is a further limiting concern at this scale: stress transmitted through the package to the MEMS resonators via the sensing-element anchor point manifests as sensor bias, requiring careful load distribution [64]. In a sole-mounted platform this stress is dominated by the subject’s body weight during stance. Finally, tradeoffs between IMU dynamic range and bandwidth constrain the achievable navigation accuracy, as the sensor must capture the shock of foot-ground contact without saturating [58, 57].

1.3.2 Stance-Phase Detection and ZUPT

Alongside the Stance Hypothesis Optimal dEtection (SHOE) detector detailed in Section 2.1.8, stance-phase detectors operating purely on IMU measurements include angular-rate energy detectors, AMV detectors, and acceleration-magnitude detectors [51]. Detector performance depends on the selected window size, detector type, and detection threshold, motivating parameter searches to minimize navigation error [51, 65]. The optimal threshold, however, is

not fixed but varies with the activity being performed [66]. The threshold governs a trade-off: applying too few ZUPTs allows sensor errors to accumulate unchecked, while applying ZUPTs at instances when the true velocity is far from zero introduces modeling errors, so any fixed threshold is calibrated only for a specific set of gait and environmental conditions [48]. The payoff from resolving this trade-off adaptively can be substantial; on a dataset containing both walking and running, alternating between the two activity-optimal thresholds reduced position Root-Mean-Squared Error (RMSE) by more than 50% relative to the best fixed threshold [48]. Early approaches adapted the threshold from signal-derived quantities such as classified velocity [53], and a parallel line of work has instead derived the threshold within a Bayesian formulation, in which it follows from a time-varying hypothesis prior and a loss factor quantifying the cost of incorrect decisions [67]. Machine learning offers a further route to establishing the motion context that determines what the threshold should be [68, 66]. Adaptive and instability-based approaches have also been proposed that bypass explicit stance detection by continuously adapting the measurement covariance [56].

1.3.3 Non-Standard Gait and Multi-IMU Aiding

The navigation performance of ZUPT-aided INS has been evaluated for many pedestrian activities, including walking, running, and crawling, though crawling has received comparatively little attention [17]. As motivated in Chapter 1, crawling is a critical firefighter gait during smoke-diving operations [25, 22]. Common firefighter crawling styles have been characterized in the biomechanics literature, including hand/knee (four-point) crawling [69] and duck crawling [70]. Firefighter Personal Protective Equipment (PPE) is also known to alter gait dynamics, including range of motion and center of gravity, and to increase dynamic instability [71]. When crawling, the foot may not reach a fully stationary state during ground contact, compromising the reliability of ZUPT with a shoe-mounted IMU. Prior work has shown that mounting the IMU on the knee can yield more accurate crawling solutions, since the knee is a primary ground-contact point [25], though the placement on

the knee has not previously been resolved further. For walking, ZUPT-aided INS accuracy has likewise been improved by optimizing the mounting position of the IMU to maximize stance-phase duration and stability [58, 59]. To exploit multiple placements, prior multi-IMU pedestrian systems typically follow one of two approaches. The sensors are either co-located on a single board, providing no spatial diversity across body segments [60], or placed on each foot with cross-limb constraints that require additional ultrasonic [72] or camera-based [39] ranging hardware. Inter-sensor kinematic constraints have been formulated to couple distributed IMUs without added hardware. These include ZPD constraints on foot-shank mountings [73], foot kinematic and shoe-modeling constraints for various motions [74], and tight-coupling schemes exploiting biomechanical motion constraints [75]. Non-zero stance-velocity and other systematic errors have also been modeled and compensated to improve ZUPT-aided accuracy [76].

Three threads from this review are carried into the chapters that follow. The SWaP+C constraints on a wearable platform, and the coupling between computational capability and the states an estimation filter can afford to carry, are addressed in Chapter 6, which develops a meso-scale navigation platform small enough to be embedded in the sole of a shoe. The gait dependence of the stance-phase detector and of its threshold, and the absence of a reliable stance-phase at the foot during crawling, are addressed in Chapter 7, which establishes the detector parameters and the mounting position appropriate to firefighter crawling. The inter-sensor kinematic constraints reviewed above are applied in Chapter 8 to couple shoe- and knee-mounted solutions without additional ranging hardware.

1.4 Hemispherical Resonator Gyroscope Literature Review

The yaw-axis gyroscope is where the sensor and navigation halves of this dissertation converge: its noise sets the upper bound on the accuracy a ZUPT-aided INS can reach. The yaw angle, and the bias of the gyroscope measuring it, are unobservable during ZUPT, so heading error accumulates without bound, causing position error to grow, as discussed in greater detail in Section 1.5. Reducing that drift therefore requires a high-accuracy gyroscope in a package compatible with a foot-mounted platform, which is the sensor developed in Part I; the remainder of this section and the gyroscope background that follows introduce the HRG on which that development builds.

Historically, gyroscopes were formulated using a massive toroidal rotor whose angular momentum preserved its original orientation. Later, optical gyroscopes based on the Sagnac effect [77], such as the ring laser gyroscope (RLG), were proposed and are still widely used today. Fiber optic gyroscopes have high performance, but they are expensive, and their scale factor is quadratically related to the size of the optical cavity (for the RLG), making further miniaturization challenging with high sensitivity [78]. Atomic gyroscopes have also been proposed [79, 80, 81, 82], which sense rotation through the precession of the magnetic moment of nuclei. The scale factor of these gyroscopes is independent of physical dimensions, but they are difficult to fully miniaturize due to the need for magnetic field generation with high uniformity [83]. Coriolis Vibratory Gyroscopes (CVGs) instead exploit the Coriolis force that arises in rotating reference frames, transferring energy between two modes of a vibrating structure proportional to the input rotation rate. While there are many CVG designs, such as the TRG [84] or the DRG [85], one of the most successful is the HRG [86]. Conventionally fabricated HRGs have demonstrated $0.0005^\circ \text{ h}^{-1}$ bias stability [87] and have accumulated over 50 million operating hours on satellites in space without a single mission failure [88]. They consist of a uniform shell and are typically mechanized using the wineglass

modes of operation. Part of the success of the HRG is due to the 3-D geometry, which increases out-of-plane stiffness relative to planar resonators [89] and limits the impact of spurious modes [90]. Fused silica is the preferred structural material, chosen for its low Coefficient of Thermal Expansion (CTE) and correspondingly low intrinsic damping [91], as discussed in Section 2.2.4.

A significant thrust in CVG research is development of a wafer-level fabrication process for microfabricated HRGs meeting the requirements of Whole-Angle (WA), or rate-integrating, operation. In WA operation, two degenerate wineglass modes are used and the precession angle of the effective pendulum is directly related to the angle of rotation. With an ideal resonator, this method of operation has virtually unlimited dynamic range and does not suffer from integration of white noise. However, WA operation places strict requirements on anisoelectricity, anisodamping, frequency split, and the Q-factor of the resonator [92, 93], as developed in Section 2.2.2. Miniaturization compounds this challenge, since it amplifies the effect of fabrication imperfections contributing to frequency split [94].

One of the most successful wafer-level fabrication methods for microfabricated HRGs is micro-glassblowing, which produces uniform 3-D shells [95, 96]. The tolerance of conventional glass microfabrication methods is generally too large to create structures symmetric enough for WA operation [92], motivating the use of micro-glassblowing for creation of the 3-D structure. During glassblowing, the surface energy of the resonator is minimized, leading to minimization of surface roughness and structural asymmetry [97]. The DSG studied in this dissertation is formed by this process; the fabrication sequence is described in Section 2.2.5.

A uniform resonator alone does not constitute a gyroscope, as external driving and sensing are required. Conventional mechanization of the HRG is electrostatic, which imposes two constraints on the device: a separately fabricated electrode array must be assembled to the resonator with micrometer-level capacitive gaps, an approach first recorded in 1969 [98] and still standard practice, and the fused-silica resonating element must be coated with a conduc-

tive film. Both constraints are addressed in this dissertation. The background supporting these developments is presented in Chapter 2. Prior efforts to address these constraints are reviewed in the following sections, motivating the contributions of this dissertation: self-aligned, fabrication-integrated electrodes are developed in Chapter 3, and non-metallized transduction in Chapters 4 and 5.

1.4.1 Structural Asymmetry Compensation

Achieving mode degeneracy in the DSG as-fabricated is complicated by asymmetry sources specific to wafer-level glassblowing, which differ in origin from the lithography- and etch-defined asymmetries of conventional MEMS fabrication [99]. Three geometric asymmetries have been identified in the DSG through optical metrology and modeled using Finite Element Analysis (FEA): deviation of the inner shell rim from circular, tilt of the junction between the inner and outer shells, and lateral shift of the stem [99]. For the $n = 2$ wineglass mode, radial deviation of the rim is the dominant contributor to frequency split, followed by junction tilt, while stem shift has a minimal effect on degeneracy, though it remains relevant to anchor loss [100, 99].

While DSGs with an as-fabricated frequency split of a few Hz have been demonstrated [3], achieving this with high yield remains a significant challenge [99]. The scaling of these tolerances with device size has been analyzed for the micro-glassblowing process by Senkal et al. [96], albeit for a comparatively simplified fabrication process. That analysis found that as the shell becomes thinner the susceptibility to frequency asymmetries increases by orders of magnitude, corresponding to a three order-of-magnitude tightening of the absolute fabrication tolerance for a microfabricated HRG relative to a macro-scale HRG. While process improvements reduce the magnitude of asymmetry, a post-processing trimming step is desired to achieve optimal scale factor.

Correcting as-fabricated mechanical frequency split is essential for non-metallized gyroscope operation, since electrostatic tuning is no longer available. Electrostatic tuning is in any case imperfect, as it introduces nonlinearities when high biasing voltages are used and can degrade gyroscope performance [101]. Trimming perturbs the stiffness and mass of the principal axes by selectively removing material, performed by laser ablation [102, 103, 104, 105], FIB etching [104, 106, 107], chemical etching [108], and directional lapping [109]. The specific advantages of laser trimming include fast and precise material removal, potential for excellent surface quality, and accuracy. Laser trimming has been demonstrated to achieve frequency splits of 0.5 Hz [102], 0.3 Hz [103], 0.1 Hz [104], and 0.032 Hz [105].

1.4.2 Reduction of Energy Dissipation

In WA mode, high Q-factor reduces angle drift, while in Force-to-Rebalance (FRB) mode, Q-factor determines scale factor, bias drift, and noise floor [110]. Increasing the Q-factor also reduces energy consumption, increases the thermomechanical SNR, and improves in-run noise performance [110, 111].

The dominant dissipation mechanism among those introduced in Section 2.2.4 is highly dependent on the specific fabrication process; for precision-milled macro-scale HRGs, for example, surface loss due to subsurface damage introduced by milling was found to be dominant [112]. Utilizing the ultra-low CTE of the fused-silica structural material and a dynamically balanced degenerate mode pair, such as the $n = 2$ or $n = 3$ mode, Q_{TED} and Q_{anchor} of fused-silica HRGs are typically on the order of a few million to a few tens of millions, to the point where the contribution of surface loss can become dominant [113, 97].

Material selection The ultimate limit of Q-factor for a fused-silica DSG is set by ThermoElastic Damping (TED) [114]. For fused-silica, with a CTE on the order of $0.52 \text{ ppm } ^\circ\text{C}^{-1}$ [115], this limit corresponds to a Q-factor on the order of tens of millions [87]; a Q-factor of 25×10^6

has been reported for macro-scale HRGs, and values as high as a hundred million are predicted [87].

Surface post-processing Surface loss originates from SSD accumulated during fabrication processes such as lapping and polishing [116, 117]. SSD in the form of microcracks has been linked to surface loss [118, 119], and the relationship between lapping-film grit size and microcrack depth has been characterized [117], with more detailed models relating SSD depth to material properties and surface roughness [120]. Surface defects can also form during cooling of the glassblowing and annealing processes due to residual stress [121]. SSD can be removed by chemical etching of fused-silica with hydrofluoric acid, followed by cleaning including sulfuric acid to remove water-insoluble fluorides and fluorosilicates [122]. Water adsorbed onto the fused-silica surface has also been identified as a surface-loss mechanism, with uncoated shell resonators exhibiting a Q-factor reduction of up to 70 % after exposure to aqueous environments [123].

Annealing Annealing of the DSG after glassblowing has been studied previously [124, 125, 89, 3]. The resulting increase in Q-factor was attributed to improved surface roughness and reduced thermal stress originating from non-isothermal cooling after micro-glassblowing.

Anchor loss Anchor loss is strongly geometry-dependent, and optimal geometries for the DSG have been studied [100, 6]. The γ ratio, defined as $\gamma = D/L$ where D is the distance from the stem to the junction between the inner and outer shells and L is the inner shell rim length, has been identified as indicative of anchor loss [100, 6]. In particular, a hemitoroidal design has been shown to reduce anchor loss by balancing the stress at the junction between the inner and outer shells in the out-of-plane direction, and Q-factor decreases with increasing junction radius [6].

Metallization Traditional mechanization coats the electrically non-conductive fused-silica HRG with a thin metal film to provide electrostatic bias, drive, and sense. A film on the order of 10 nm is deposited by sputtering or e-beam evaporation [3, 126, 127]. A sufficiently thick coating is necessary to ensure adequate conductivity of the proof mass, but such a coating has also been shown to decrease the Q-factor from 740×10^3 to approximately 410×10^3 [126]. Notably, sputtered metal stacks and conformal Atomic Layer Deposition (ALD) platinum coatings have been observed to reduce Q-factor by 30 % to 45 % and up to 70 %, respectively, on fused-silica shell resonators, with metals whose Young’s modulus closely matches fused-silica producing a comparable reduction [128]. An optimized fabrication process can mitigate anchor loss, internal loss, and surface loss, but cannot recover the Q-factor degradation introduced by surface coating of the resonating element. Therefore, a transduction method that preserves the pristine fused-silica inner shell eliminates this loss mechanism and opens a pathway toward the theoretical Q-factor limit for microfabricated DSGs, predicted to reach tens to even a hundred million [87].

1.4.3 Femtosecond Laser Induced Chemical Etching

Recent advances in Femtosecond Laser Induced Chemical Etching (FLICE) have established it as a uniquely capable tool for fabrication of high-aspect-ratio features in fused-silica. Femtosecond laser irradiation below the ablation threshold forms sub-wavelength nanogratings that locally increase the chemical etching rate by several orders of magnitude, and when aligned via polarization control these reach etch selectivity greater than 900:1 [129, 130, 131]. Because the modification is confined to the laser focal volume, it can be inscribed deep within the transparent substrate, allowing arbitrarily shaped channels, planes, or volumes to be defined [132]. The modified material is subsequently removed by wet etching in aqueous HF or KOH solution [129, 132].

The capability of FLICE addresses a long-standing gap in fused-silica micromachining. Conventional wet etching of fused-silica is isotropic and lacks the precision, controllability, and repeatability required for MEMS resonators, while plasma etching is complex and, due to the chemical inertness of the material, typically yields aspect ratios below 10:1 [132]. FLICE has conventionally been used for microfluidic [133] and photonic [134] applications, and has great potential for fabrication of MEMS resonators [132, 135, 136, 137].

The first application of FLICE to fabrication of capacitive MEMS resonators was demonstrated by Vatanparvar and Shkel [132], in which a fused-silica DRG was digitally patterned and released from a single-layer fused-silica substrate, eliminating the wafer bonding and lithography steps of plasma-based fabrication. Through optimization of pulse energy, marking speed, and polarization, an etch selectivity on the order of 855:1 was achieved, and channels with aspect ratios as high as 55:1 were demonstrated while preserving surface quality. The released, metallized DRG exhibited a Q-factor on the order of 6×10^5 at a resonant frequency near 100 kHz, approximately five times higher than previously reported for plasma-etched fused-silica DRGs [132].

1.4.4 Transduction of Hemispherical Resonator Gyroscopes

Electrostatic transduction Conventionally, HRGs are driven electrostatically [138], where the vibrating element, typically formed of glass, is coated with metal and aligned with electrodes to form a capacitor. In prior art, driving and sensing have been performed using an electrode array that is typically fabricated separately and then assembled together with the resonating element. Precise assembly is critical to create micrometer-level capacitive gaps, which are necessary to generate sufficient capacitance to drive and sense the resonator.

The industry standard involves separate fabrication of the electrode array, followed by precise assembly to establish capacitive gaps with the resonating element. The first record of this electrode-assembly approach with an HRG dates to 1969 [98]. Generally, the approaches

with the greatest capacitance use precisely manufactured 3-D electrode substrates tailored to individual HRGs, with shims for precise capacitive gaps; however, this is a costly and time-intensive approach [139, 140]. More recent methods have used 3-D electrode substrates fabricated out of silicon using DRIE with precise alignment [141, 142]. A further approach uses a flat, separately fabricated electrode substrate with capacitive gaps formed out-of-plane via precise bonding [143]. There are several disadvantages to fabricating a separate electrode assembly: the additional cost of electrode fabrication, the necessity of precise assembly, and the assembly error, which typically must be modeled or characterized and which limits gyroscope performance. Assembly misalignment produces uneven capacitances that impact actuation efficiency, detection sensitivity, and stiffness-adjustment flexibility, and this tilt-induced capacitance variation has been modeled and characterized in prior work [144, 145].

Several approaches for self-aligned electrodes have been proposed. Some approaches blow additional shells in close alignment with a primary resonating element to form 3-D electrodes with large overlap areas [146, 147]. The primary limitation of these approaches is the achievable capacitive gap: blowing two structures in close proximity is challenging and may achieve a small gap in some locations, but the overall capacitance typically remains small due to insufficient overlap across the entire structure. A more recent variant blows two structures with a sacrificial ALD platinum layer sandwiched between them, which is subsequently removed to leave conformal electrodes [148]. One proposed approach can achieve small capacitive gaps [149], but in that case the resonator is formed of polysilicon, which yields a substantially lower Q-factor than fused-silica glass. Once a resonator is metallized, conformal coating of the resulting high-aspect-ratio capacitive gaps is itself challenging: magnetron sputtering struggles to conformally coat even 5:1 aspect-ratio structures [150], whereas plasma-enhanced ALD of platinum has demonstrated full conformality on SiO_2 trenches with aspect ratios as high as 34 [151], at the cost of Q-factor-degrading film thickness [143], and sputtering shadowing effects on high-aspect-ratio structures further limit continuity of the

deposited film [152]. These limitations motivate a fabrication-integrated electrode approach that eliminates separate assembly (Chapter 3).

Non-metallized transduction Non-metallized transduction avoids a conductive film on the resonator but remains further from fielded use. First, it is necessary to detect displacement at specific locations on the rim of the DSG without metallization. This can be achieved using advances from the rapidly evolving field of MOEMS [153]. Multiple studies have published how evanescent coupling can be used for non-metallized sensing of the micro-fabricated HRG [154, 155]; however, no experimental results have been reported. MOEMS interference-based displacement sensors have displayed $123 \text{ pm}/\sqrt{\text{Hz}}$ displacement sensitivity with a range of 40 nm [156] and $0.14 \text{ pm}/\sqrt{\text{Hz}}$ sensitivity already integrated into a MEMS gyroscope [157]. A Mach-Zehnder Interferometer (MZI) scheme has also been proposed to detect the movement of the DSG at the die-level, and successful detection of rim movement has been reported [158, 159].

For non-metallized driving, in a CVG the precession of a standing wave under an applied rotation rate is the basis of angular measurement [110]; therefore, generating and sustaining this standing wave in the resonating element with maximum amplitude is essential. Non-metallized driving of a microfabricated HRG is generally a more challenging problem, because it must drive along a specific axis to enable rate-gyroscope operation, so simple shaker excitation, commonly used for initial characterization [89], is sufficient only for WA operation. A few methods to directionally excite a microfabricated HRG without metallization have been proposed. Opto-mechanical coupling has been proposed for driving via the whispering gallery mode [154, 155]; however, the efficacy of driving is heavily dependent on the coupling distance between the optical waveguide and the resonator, which can be challenging to co-fabricate. Alternatively, Honeywell filed a patent in 2020 regarding optical excitation of an HRG using selective laser heating [160], which has recently been demonstrated on an HRG [161, 162].

Indirect piezoelectric actuation To preserve the pristine fused-silica inner shell and eliminate Q-factor degradation from surface coating, IPA of the DSG was developed, first presented by Wang et al. [163, 164] and Shkel et al. [165]. IPA exploits the dual-shell architecture, in which the inner shell functions as the vibratory element and the outer shell, originally intended for shock resistance and die-level vacuum packaging [3], is also used to drive the inner shell. In IPA, a piezoelectric film is deposited exclusively on the 3-D outer shell [166] using either ALD or PHT [163] or sputtered PZT [164], and electrodes are patterned to selectively excite the $n = 2$ wineglass mode of the inner shell without metallizing the resonating element. When voltage is applied to these electrodes, the piezoelectric film stresses the outer shell; this stress propagates through the junction to excite the inner shell. Gyroscope operation using IPA has been demonstrated with proof-of-concept electrostatic detection [164].

1.5 Navigation Accuracy and Gyroscope Grade for ZUPT-aided INS

1.5.1 Summary of Commercially Available IMUs

IMUs are commonly categorized into performance grades that reflect characteristic navigation capabilities, fabrication technologies, and cost. These grades span consumer/industrial, tactical, and navigation tiers based on pricing and stability, with strategic grade forming the highest tier. Table 1.1 summarizes ten commercially available IMUs spanning these grades, alongside two MEMS HRGs representing recent research progress. The comparison is framed around a combined SWaP+C metric, motivated by the central goal of this dissertation: attaining high-end gyroscope performance while retaining the SWaP+C advantages that MEMS fabrication affords through parallel, wafer-level batch production.

Consumer-grade IMUs occupy the lowest performance tier. Most IMUs used in mobile robots and similar platforms are consumer-grade and based on MEMS technology, fabricated at low cost in silicon for high-volume applications. Their gyroscope bias stability typically falls between tens and hundreds of degrees per hour, and because they were never designed to hold a steady pose for long periods, they drift rapidly when used without an external reference such as GPS. Typical applications include motion sensing and user interaction in consumer electronics such as smartphones, game controllers, and virtual-reality devices.

Tactical-grade IMUs offer improved short-term navigation. A widely used demarcation defines a gyroscope in-run bias stability below 5°h^{-1} as tactical grade, with higher values classified as industrial grade, in contrast to consumer devices, which typically exhibit bias stability of 30°h^{-1} or more. In practical terms, tactical-grade units can navigate in GPS-denied environments for up to roughly ten minutes and carry price points ranging from approximately \$5,000 to \$25,000. They are most often used in military applications, ground robots, and vehicles navigating GPS-denied environments.

Navigation-grade IMUs provide high performance for demanding platforms. Their gyroscopes achieve bias stability on the order of $0.1^{\circ} \text{h}^{-1}$ or better, and they are generally no longer implemented in MEMS technology, relying instead on ring laser gyroscopes, fiber-optic gyroscopes, or macro-scale HRGs. These units are correspondingly expensive and are applied to platforms such as satellites and marine vessels.

Strategic-grade IMUs represent the highest-performance tier. Built on HRG and related technologies, they enable autonomous navigation for several days with virtually no drift, and are reserved for applications with extremely high reliability requirements, such as aerospace vehicles, submarines, and missile guidance. Such units can cost on the order of millions of dollars.

This progression in performance is accompanied by a dramatic escalation in resource cost. Fig. 1.2 plots the angle random walk of each device in Table 1.1 against the SWaP+C metric. Across the four commercial grades, the SWaP+C metric spans more than twenty orders of magnitude, increasing from approximately $10^{-15} \text{ m}^3 \cdot \text{kg} \cdot \text{W} \cdot \$$ for consumer-grade chips to roughly $10^6 \text{ m}^3 \cdot \text{kg} \cdot \text{W} \cdot \$$ for the strategic-grade unit. This span quantifies the trade-off motivating this dissertation: the orders-of-magnitude improvement in gyroscope performance from consumer to strategic grade is presently obtained only at a correspondingly steep cost in size, weight, power, and price. The MEMS HRGs listed in Table 1.1 point toward a way to break this trade-off. Because the HRG is the same core technology already used at the navigation and strategic grades, realizing it in a batch-fabricated MEMS form promises to decouple performance from resource cost. The two reported MEMS HRGs illustrate this progress: their gyroscope performance is approaching, and in the best case entering, the navigation-grade regime while retaining the low-SWaP+C character of MEMS. Attaining navigation- and strategic-grade performance in a MEMS device is the objective that motivates this work, and the following subsection establishes why that performance level matters for pedestrian inertial navigation.

1.5.2 Navigation Error in ZUPT-Aided INS

Having established the performance grades of available gyroscopes, this subsection links gyroscope performance to navigation error in ZUPT-aided pedestrian INS, and shows why gyroscope grade, in particular that of the yaw-axis gyroscope, remains the limiting factor even when zero-velocity updates are applied. Conventionally this link is made by modeling the motion of the foot during walking [178]; the closed-form analytical model of Wang et al. [171] is adopted here for simplicity. The model expresses the growth of the foot-mounted position-estimate uncertainty, resolved into components parallel and perpendicular to the direction of travel, in terms of the IMU noise coefficients and a small set of gait parameters. Under the assumptions of a straight, northward trajectory, two-dimensional foot motion,

Table 1.1: Comparison of some commercial gyroscopes by performance grade. A comparison of the gyroscope performance of reported MEMS fabricated HRGs is provided below as well.

#	Manufacturer	Part	Gyroscope			SWaP+C				Source
			ARW _g	BI _g	FS _g	V	m	P	Cost	
			$\left[\frac{^\circ}{\sqrt{\text{hr}}}\right]$	$\left[\frac{^\circ}{\text{hr}}\right]$	$\left[\frac{^\circ}{\text{s}}\right]$	$[\text{mm}^3]$	$[\text{kg}]$	$[\text{W}]$	$[\$]$	
<i>Consumer grade (chip-scale)</i>										
1	TDK InvenSense	ICM-42605	0.228 [§]	—	2000	6.8	$3 \times 10^{-5\ddagger}$	1.6×10^{-3}	3.6	[167]
2	STMicro	LSM6DSV	0.420 [§]	—	2000	6.2	$3 \times 10^{-5\ddagger}$	1.2×10^{-3}	5.5	[168]
3	Bosch	BMI323	0.420 [§]	—	2000	5.3	$3 \times 10^{-5\ddagger}$	1.4×10^{-3}	3.2	[169]
<i>Industrial / tactical grade (module-level)</i>										
4	Analog Devices	ADIS16505-2	0.15	2.7	500	1.3×10^3	0.011	0.14	864	[170]
5	VectorNav	VN-200	0.21	5	2000	1.6×10^3	4×10^{-3}	0.34	8×10^3	[171, 7]
6	Xsens	MTi-3	0.18 [§]	6	2000	1.1×10^4	6×10^{-4}	0.50	425	[172]
7	Gladiator	G200	0.18 [§]	10	300	1.83×10^4	0.018	0.055	$2 \times 10^{3\ddagger}$	[173]
<i>Navigation grade</i>										
8	Safran	Iconyx HP-30	0.001	0.15	4000	5.3×10^5	0.75	10	$5 \times 10^{4\ddagger}$	[174]
9	Honeywell	HG9900	0.0021	0.0021	4000	3.1×10^6	2.9	10	$2 \times 10^{5\ddagger}$	[175]
<i>Strategic grade</i>										
10	N. Grumman	SSIRU	5×10^{-5}	5×10^{-4}	7	7.75×10^6	7.1	43	$1 \times 10^{6\ddagger}$	[176]
<i>MEMS HRG</i>										
	NUDT	MSRG	0.035	0.877	—	—	—	—	—	[177]
	UC Irvine	DSG	0.037	0.3	—	—	—	—	—	[144]

[§] Derived ARW from the published gyroscope noise density.

[‡] Order-of-magnitude estimate.

moderate IMU performance, and a high sampling rate, the model reproduces Monte-Carlo results to within 10% [171].

The model predicts position-error covariance parallel ($P_{\parallel}$) and perpendicular ($P_{\perp}$) to the direction of travel. The along-track covariance grows linearly in the total navigation time t [171]:

$$P_{\parallel} = \left(2 - \frac{t_{\text{stride}}}{4}\right) \frac{r^2 t_{\text{stride}}}{f_s t_{\text{stance}}} t, \quad (1.1)$$




Figure 1.2: Comparison of commercially available IMUs by performance grade, showing ARW against SWaP+C. Numeric labels correspond to the entries in Table 1.1. Inspired by Lutwak [1].

so that $\sigma_{\parallel} = \sqrt{P_{\parallel}} \propto \sqrt{t}$. This component depends only on the zero-velocity measurement uncertainty r , the sampling frequency f_s , and the gait timing; it is independent of the gyroscope noise and is therefore common to all gyroscope grades [171]. The dominant and grade-dependent error is the cross-track covariance $P_{\perp}$. This is the crux of the problem: the yaw (heading) angle is unobservable during ZUPT-aided INS, so heading error accumulates without bound and couples directly into cross-track position error, making the yaw-axis gyroscope performance decisive. The cross-track covariance accumulates as [171]:

$$P_{\perp} = \left(2 - \frac{t_{\text{stride}}}{4}\right) \frac{t_{\text{stride}} r^2}{t_{\text{stance}} f_s} t + \frac{1}{3} \text{ARW}^2 s_N^2 t^3 + \frac{a_s^2}{12} \sigma_{gN}^2 s_N^2 t^4 + \left(\frac{1}{30} + \frac{a_c^2}{60}\right) \text{RRW}^2 s_N^2 t^5, \quad (1.2)$$

with $\sigma_{\perp} = \sqrt{P_{\perp}}$. Here ARW and Rate Random Walk (RRW) are the ARW and RRW of the gyroscope, s_N is the stride length, and a_c and a_s are the gait-cycle averages of the cosine

and sine of the foot pitch angle. The gyroscope-bias uncertainty along the north axis, σ_{gN} , which enters Eq. (1.2) through the fourth-order term, is given in closed form in [171].

The cross-track error is thus governed almost entirely by the two gyroscope noise coefficients, the ARW and the RRW. Both enter through the heading channel, and the lower the noise of the yaw-axis gyroscope, the slower the cross-track error grows; higher grades suppress $\sigma_{\perp}$ by orders of magnitude over a given interval, consistent with [171]. By contrast, a stand-alone (unaided) strapdown INS provides no bound on error growth at all: gyroscope errors couple a component of gravity onto the horizontal axes, which is then double-integrated, so that horizontal position error grows as t^3 and reaches hundreds of meters over the intervals considered here [46, 179]. ZUPT suppresses this growth, but only along axes it can observe; because it cannot observe yaw, the residual cross-track drift in Eq. (1.2) can only be reduced by an external heading reference or, where none is available in mission-critical scenarios, by a sufficiently precise yaw-axis gyroscope. This is the requirement that motivates the development of a low-SWaP+C, high-performance MEMS HRG as the z-axis gyroscope for pedestrian navigation. Fig. 1.3 evaluates Eq. (1.2) for representative gyroscope grades, showing that the cross-track uncertainty falls by orders of magnitude as the yaw-axis gyroscope advances from consumer grade toward navigation-grade performance.

1.5.3 Idealizations of the Position-Uncertainty Model

The model of Eqs. (1.1)–(1.2) isolates the contribution of gyroscope noise to position uncertainty by assuming that every other element of the system behaves ideally, and three idealizations are worth making explicit. First, the zero-velocity measurement is treated as a zero-mean random variable of standard deviation r , whereas in reality the sensor is not stationary during the declared stance, so the residual velocity is neither zero-mean nor uncorrelated and depends strongly on gait and mounting position, as measured directly in Section 7.2. Second, stance detection is assumed perfect and to recur identically, whereas in




Figure 1.3: Perpendicular position uncertainty $\sigma_{\perp}$ over time for ZUPT-aided pedestrian INS, evaluated from Eq. (1.2) for three gyroscope grades (consumer, tactical, and near-navigation), together with the parallel uncertainty $\sigma_{\parallel}$ from Eq. (1.1), which is common to all grades. Simulation parameters are listed in Table 1.2.

Table 1.2: Parameters used in the position-uncertainty model of Eqs. (1.1)–(1.2). Values marked “measured (this work)” were characterized by the author.

Symbol	Description	Value	Source
<i>Gait and filter parameters</i>			
t_{stride}	gait-cycle duration	1.53 s	[180]
t_{stance}	stance-phase duration	0.92 s ^a	[181]
s_N	stride length	1.32 m	[178]
a_c	$\langle \cos \theta \rangle$ over gait cycle	0.838	measured (this work) ^b
a_s	$\langle \sin \theta \rangle$ over gait cycle	0.271	measured (this work) ^b
r	zero-velocity meas. uncertainty	0.02 m s ^{−1}	[171]
f_s	IMU sampling frequency	800 Hz	matches this work [182, 183]
<i>Consumer grade: InvenSense ICM-20948</i>			
ARW	angle random walk	0.96 °/√h	measured (this work)
RRW	rate random walk	8.94 × 10 ^{−3} °/s/√h	measured (this work)
<i>Tactical grade: Analog Devices ADIS16490</i>			
ARW	angle random walk	0.09 °/√h	[184]
RRW	rate random walk	1.03 × 10 ^{−3} °/s/√h ^c	[184]
<i>Near-navigation grade: Innalabs KVG25 (CVG)</i>			
ARW	angle random walk	3 × 10 ^{−3} °/√h	[185]
RRW	rate random walk	8.32 × 10 ^{−5} °/s/√h	[185]

^a Taken as 60 % of t_{stride} following the standard stance-fraction convention for normal walking [181].

^b Values are averaged over 12 straight-line walking trials with a foot-mounted IMU, comprising ≈ 6 min of walking data from three subjects.

^c Order-of-magnitude estimate from the Allan deviation.

practice the interval must be inferred from the IMU signals through a detector and threshold, Section 2.1.8; the model further assumes straight, two-dimensional travel with constant gait parameters, so the gaits of Chapter 7, which differ from walking in stance duration and stability, are not represented, and the parameters of Table 1.2 are representative of one activity rather than general. Third, the error model retains only the four random-walk terms and neglects misalignment, scale factor, nonlinearity, and the thermal-transient bias of a self-heating platform [62, 63], as well as finite bandwidth and dynamic range [56, 57, 186]. Figure 1.3 should therefore be interpreted as a lower bound on the achievable uncertainty, and improving the gyroscope and improving the fidelity of the zero-velocity measurement are complementary rather than alternative routes to accuracy, both of which are pursued in the chapters that follow.

1.6 Problem Statement

Pedestrian inertial navigation is well suited to mission-critical operations because it is self-contained: it requires no satellite signal, no prior infrastructure, and no map of the environment, and it therefore remains available in the indoor, underground, and obscured settings where other modalities degrade. Its limitation is the unbounded growth of integration error, and the established remedy is ZUPT aiding, which exploits the stationary interval of the gait cycle to bound velocity error and to estimate the inertial sensor biases [47, 51]. ZUPT-aided inertial navigation has been demonstrated to deliver meter-level accuracy over trajectories of several minutes using low-cost IMUs [48].

Two limitations bound that performance. The first is structural: the zero-velocity pseudo-measurement constrains velocity and, through the gravity vector, roll and pitch, but it carries no information about heading. The yaw angle and the bias of the yaw-axis gyroscope are therefore unobservable, and heading error accumulates without bound, driven by the drift of that gyroscope. The second is behavioral: ZUPT presumes that the sensor is briefly stationary once per gait cycle, and the gaits encountered in emergency response may not reliably provide such an interval. This dissertation therefore asks how the accuracy of pedestrian inertial navigation can be improved in mission-critical scenarios where external aiding is unavailable, by extending ZUPT-aided inertial navigation to the non-standard gaits encountered in the field, and by developing a wafer-fabricated gyroscope capable of constraining the heading drift that ZUPT cannot observe.

1.6.1 Transduction of the Dual-Shell Hemispherical Resonator Gyroscope

Constraining heading drift requires a better gyroscope. As established in Section 1.5, instruments capable of the required performance exist, but at a SWaP+C that a body-worn system cannot accommodate and firefighters cannot afford, while consumer-grade devices

that satisfy those constraints do not provide the required performance. The DSG is a candidate for bridging this gap: it is batch fabricated at the wafer level [3], and has the potential to approach the performance of macro-scale HRGs. Several parameters stand between it and that performance, among them the Q-factor, the amplitude of motion, the modal mass, and the frequency split, Eqs. 2.31 and 2.32. This dissertation addresses those governed by transduction, the means by which the resonator is driven, sensed, and tuned, which jointly bounds the achievable Q-factor, the amplitude that can be sustained, and the means available for correcting the frequency split. Three challenges follow.

Conventional electrostatic transduction of a micro-scale HRG requires an electrode array fabricated separately from the resonator. This array must then be aligned to form a sufficiently small and uniform capacitive gap, and the tilt and misalignment introduced during this step produce a capacitance non-uniformity that must be compensated and, if left uncompensated, can limit gyroscope performance [3, 144]. The approach also requires a conductive film on the resonating element, and metallization of a fused silica DSG has been shown to reduce the Q-factor substantially [126, 128]. An optimized design and fabrication process can mitigate anchor, internal, and surface loss but cannot recover the loss introduced by the film, so metallization bounds the Q-factor regardless of how well the remaining loss channels are controlled. Each constraint defines a path toward improved transduction. Along the first, the assembly step is eliminated while electrostatic transduction is retained; along the second, the metallization is removed, which requires abandoning electrostatic transduction of the resonating shell altogether, and with it the driving, tuning, and sensing that electrostatics provides.

IPA drives the inner shell of the DSG by driving a piezoelectric film on the outer shell and transmitting the resulting stress through the junction where the two shells meet. The approach has been demonstrated [164, 163], but the design parameters governing how efficiently stress is transferred between the shells have not been investigated. Anchor loss, the radiation

of acoustic energy from the resonator into the supporting substrate, has been shown to dominate the Q-factor of DSG, and the geometric ratio governing it has been identified [100, 6]. That same ratio is expected to govern stress transfer through the junction, placing the two in tension: the geometry that delivers the most stress to the resonating element also radiates the most energy away from it. The amplitude of motion depends on the product of Q-factor and transduction efficiency, so this tension sets the achievable ARW and noise performance (Eq. 2.32).

Research Question

- **RQ1.** How can the transduction of the fused silica DSG be improved?
 - **RQ1a.** How can capacitive electrodes be defined for a DSG without an assembly step, and what governs the minimum capacitive gap and maximum electrode aspect ratio achievable by such a process?
 - **RQ1b.** How can a DSG be driven, sensed, and tuned without a conductive coating on the resonating element?
 - **RQ1c.** Which dual-shell geometry maximizes the displacement amplitude achievable by indirect piezoelectric actuation, given that the geometric parameter governing transduction efficiency also governs anchor loss?

RQ1a is addressed in Chapter 3, which develops laser-patterned self-aligned electrodes and the process optimization that sets the achievable capacitive gap. RQ1b is addressed in Chapter 4, which combines non-metallized driving, optical sensing, and laser ablation trimming on a single device. RQ1c is addressed in Chapter 5, which models and experimentally evaluates the geometric dependence of indirect piezoelectric actuation.

1.6.2 Realization of a Compact, Self-Contained Navigation Platform

An INS requires more than one gyroscope. It requires a full six-axis IMU, aiding sensors, a processor capable of executing the navigation filter hundreds of times per second, a power source, and a means of communicating the navigation solution. A ZUPT-aided INS depends on a stationary interval of the gait cycle, and the foot is the body segment where that interval is most reliably available [47], which makes the sole of the shoe the target location for the assembly. Placing the platform there is challenging, however, as the solution needs to be compact and able to withstand the stress from body weight during the stance-phase. Three challenges follow.

The volume available in a shoe sole must be traded against computational capability, which in turn bounds the states the estimation filter can afford to carry at the required update rate. Prior platforms have omitted sensor bias states to preserve that rate [42], a compromise that directly limits the drift the system can correct. Two further effects appear only once the platform is built and worn, and neither is captured by component-level specification. Power dissipated by the processor raises the temperature of the IMU, a source of sensor bias while the platform approaches steady state; and mechanical loading of the shoe sole during the stance-phase transmits stress to the sensor package.

Research Question

- **RQ2.** What system-level constraints arise when pedestrian inertial navigation is implemented in a compact, self-contained, shoe-integrated platform?
 - **RQ2a.** What navigation accuracy is achievable from a self-contained inertial platform small enough to be embedded in the sole of a shoe?

- **RQ2b.** What computational, thermal, and mechanical constraints does that integration impose?

RQ2a and RQ2b are addressed in Chapter 6, which presents the design, integration, and experimental evaluation of a meso-scale pedestrian navigation platform.

1.6.3 Navigation Across Firefighter Walking and Crawling Gaits

The accuracy reported for ZUPT-aided systems is obtained while walking with a shoe-mounted IMU, but firefighters performing smoke diving are forced to the floor by dense smoke and elevated temperatures, making crawling a gait that any firefighter localization system must support. During crawling the foot does not reach a clean, sustained zero-velocity interval, so the assumption underlying ZUPT is violated at the placement on which walking accuracy depends [182, 25]. Residual velocity during the detected stance-phase degrades the bias estimates that the correction is intended to produce (Eq. 1.2), and the resulting error accumulates over the trajectory. Relocating the sensor does not resolve ZUPT for all firefighter gaits: the knee provides a stable stance during common crawling gaits but no comparable anchor during walking, and firefighters alternate between the two within a single deployment. The stance-phase detector, its window length and threshold, and the measurement noise assigned to the zero-velocity update are all gait dependent, so configuring them correctly requires identifying the current gait, and the transitions between gaits, at run time.

Research Question

- **RQ3.** How can ZUPT-aided pedestrian navigation be extended from conventional walking to firefighter crawling and transitions between gait modes?
 - **RQ3a.** How should the stance-phase be detected during firefighter crawling, where on the body should the IMU be mounted, and what residual velocity does

that placement incur during stance? These parameters have been established for walking [51, 76], but not for the crawling gaits used during smoke diving.

- **RQ3b.** How can the navigation accuracy that ZUPT provides during walking be maintained across the crawling gaits used in firefighter operations, given that no single sensor placement provides a reliable stance-phase in both?

RQ3a is addressed in Sections 7.1 and 7.2 of Chapter 7, which conduct a parameter search over detector type, window length, and threshold, and evaluate four mounting positions across two crawling methods with the residual stance-phase velocity measured directly. RQ3b is addressed in Chapter 8, which develops a dual-IMU framework spanning both gaits.

Table 1.3: Summary of the research questions and the chapters in which they are addressed.

Research Question	Chapter
RQ1. How can the transduction and tuning of the fused silica DSG be improved?	Part I
RQ1a. Assembly-free capacitive electrodes and the achievable gap and aspect ratio.	Ch. 3
RQ1b. Driving, sensing, and tuning without a conductive coating on the resonating element.	Ch. 4
RQ1c. Dual-shell geometry maximizing indirect piezoelectric actuation amplitude against anchor loss.	Ch. 5
RQ2. What system-level constraints arise when pedestrian inertial navigation is implemented in a compact, self-contained, shoe-integrated platform?	Part II
RQ2a. Navigation accuracy achievable from an in-sole self-contained platform.	Ch. 6
RQ2b. Computational, thermal, and mechanical constraints that in-sole integration imposes.	Ch. 6
RQ3. How can ZUPT-aided pedestrian navigation be extended from conventional walking to firefighter crawling and transitions between gait modes?	Part II
RQ3a. Stance-phase detection, mounting position, and residual stance velocity during crawling.	Ch. 7
RQ3b. Maintaining walking accuracy across crawling when no single placement serves both gaits.	Ch. 8

1.7 Dissertation Overview

The accuracy of ZUPT-aided INS is bounded both by sensor-level drift and by the system-level challenges of packaging, aiding, and gait-robust estimation. This dissertation addresses these complementary limitations as parallel advances toward a future integrated system rather than by building that system itself, and is organized in two parts accordingly: Part I on sensor development and Part II on system-level navigation. The sensor contributions of Part I are component-level: two transduction methods for the wafer-level DSG were developed, characterized, and each operated as a rate gyroscope, but were not incorporated into a navigation platform, which remains future work. The navigation contributions of Part II were experimentally demonstrated: a meso-scale self-contained platform and a set of gait-robust estimation algorithms were built and evaluated on commercial IMUs, across walking and the crawling gaits of firefighter operation. The two are joined by the noise of the yaw-axis gyroscope, which bounds the accuracy a ZUPT-aided INS can reach. The technical background common to both parts, comprising strapdown inertial navigation and zero-velocity aiding together with the fundamentals of CVGs and the fabrication of the DSG, is presented in Chapter 2. The contributions of each part are described chapter by chapter in the overview that follows and stated formally in Chapter 9.

In Chapter 3, we report on a fabrication approach that uses FLICE to digitally define high-aspect-ratio capacitive electrodes directly within the bulk of fused-silica, eliminating both the separately fabricated electrode substrate and the assembly step it requires. The process is demonstrated on the DSG, where the region between the inner and outer shells is patterned to form self-aligned, IPE. On planar test samples, gaps as small as $2\mu\text{m}$ with aspect ratios exceeding 150:1 were achieved, and applied to the DSG the process yielded bulk capacitive gaps of $5.5\mu\text{m}$ to $6\mu\text{m}$. A maximum as-fabricated Q-factor of 4.57 million indicates that FLICE patterning does not substantially increase surface-related energy loss, and a metallization study further increases total capacitance and reduces inner-shell resis-

tance relative to sputtered Cr-Au. Gyroscope operation is demonstrated in open-loop rate mode, and a projection based on the optimized FLICE parameters indicates a potential $14\times$ improvement in scale factor.

In Chapter 4, we report on non-metallized operation of the fused-silica DSG, addressing the three functions conventionally provided by metallization: driving, tuning, and sensing. Driving is achieved using IPA, in which PHT on the outer shell transfers stress through the junction to excite the inner shell without coating the resonating element. Tuning is achieved using laser ablation trimming. On the PHT-coated samples, trimming directly into the rim reduced the as-fabricated frequency split from 50 Hz to 60 Hz to as low as 0.39 Hz, at the cost of Q-factor, since the ablation damages the resonating shell. Defining teeth-like-tines on the rim by FLICE enabled trimming to a frequency split as low as 0.10 Hz while reducing the trimming-induced quality-factor loss from 73 % to 21 %. Sensing is demonstrated using a custom dual LDV setup, serving as a benchtop prototype of future interferometric detection implemented in a PIC. Open-loop gyroscope operation is demonstrated with a scale factor of $0.02 \text{ mV } ^\circ\text{-1 s}^{-1}$.

In Chapter 5, we report on a study of how DSG shell geometry governs IPA transduction efficiency, motivated by the observation that greater transduction increases the amplitude of motion sustained at the antinode and therefore improves SNR. The geometry cannot be optimized for transduction alone, since the shell profile that sets the stress path through the junction also sets the radiation of acoustic energy into the substrate and therefore the Q-factor. Ten DSG designs spanning a range of shell profiles are generated using a glassblowing model and evaluated by adapting the RAFT, with the modeling predicting that displacement is maximized at an intermediate geometry rather than increasing monotonically. Five PZT-coated DSGs are used for experimental validation, demonstrating a $6.5\times$ improvement in displacement amplitude and a $12.3\times$ improvement in Q-normalized transduction efficiency

relative to the least efficient design fabricated, and establishing geometry as a key design lever for IPA.

In Chapter 6, we report on the design, implementation, and demonstration of a miniaturized inertial navigation platform for pedestrian navigation. The platform occupies a total volume of 7.1 cm^3 , excluding battery and case, and integrates an IMU, a chip-scale barometer, and a magnetometer with a computational engine capable of executing the navigation algorithm in real time. A barometer-aided ZUPT-aided INS is implemented in an EKF framework at an update rate of 500 Hz, and the form factor permits integration into the sole of a shoe. A series of indoor walking experiments demonstrated a CEP of 0.66 m over a 125 m, 2.75 min trajectory, or 0.53 % of the trajectory length, with a vertical RMSE of 0.16 m.

In Chapter 7, we report on ZUPT-aided pedestrian inertial navigation for firefighters, where smoke-diving operations force the subject to the ground and the walking gait assumption underlying conventional foot-mounted implementations no longer holds. First, a parameter search over stance-phase detector type, window size, and detection threshold is performed for hand/knee crawling with heel- and knee-mounted IMUs. Second, four candidate mounting positions and two crawling methods used by firefighters are evaluated using an electromagnetic position sensor to measure residual stance-phase velocity directly. Together these studies identify the center of the knee pad as the best placement, reducing crawling navigation error by up to $3.7\times$ relative to a shoe-mounted IMU.

In Chapter 8, we report on a dual-IMU framework that extends ZUPT-aided navigation across the walking, crawling, and transitional gaits of firefighter operation, since no single mounting position provides a reliable stance-phase across all of them. The shoe- and knee-mounted solutions are fused in a joint error-state EKF, coupled through a ZPD constraint on the rigid lower leg that requires no additional ranging hardware, with a random forest gait classifier configuring the filter for each gait. Evaluated on three subjects in full firefighter

equipment, the approach reduced the horizontal CEP from 3.86 m to 0.58 m and the vertical RMSE from 3.06 m to 1.26 m relative to a conventional foot-mounted solution.

Finally, Chapter 9 concludes the dissertation with a summary of contributions and future research directions.

Chapter 2

Background

This chapter consolidates the technical background on which the remainder of the dissertation builds. It is organized around the two threads of the work: the pedestrian inertial navigation algorithms of Part II, and the gyroscope of Part I that constrains their long-duration accuracy. Section 2.1 reviews strapdown inertial navigation, the zero-velocity update, and the estimation framework used throughout the navigation chapters. Section 2.2 then develops the fundamentals of CVGs, including the Allan deviation, the lumped-parameter model, the available mechanizations, the mechanisms of energy dissipation, and the fabrication of the DSG.

2.1 Pedestrian Inertial Navigation

2.1.1 Strapdown Inertial Navigation

Inertial navigation is based on accelerometer and gyroscope measurements. Inertial navigation algorithms are perfect in the sense that if perfect measurements were provided by the inertial sensors, then the navigation accuracy would, in turn, be perfect. However, due to sensor noise and stochastic time-varying biases, in reality the navigation solution will drift

on the order of a few meters within seconds for low-cost sensors without aiding. This section provides only a brief overview of dead reckoning and related concepts; more detailed reviews are provided in [2, 12, 179, 187], for example.

2.1.2 Representation of Navigation Information

First, in order to provide meaning to position, velocity, and attitude information it is necessary to define coordinate reference frames. A coordinate frame is defined by an origin and three axes which are usually mutually perpendicular. All navigation information is described in one coordinate frame relative to another coordinate frame. There are several common types of coordinate frames used in navigation; some of the most common are [179, 12]:

- **Inertial (i-frame):** Generally, an inertial coordinate frame does not accelerate or rotate with respect to the universe. For navigation it is useful to define the origin of the inertial frame at the Earth’s center of mass with its axes fixed to the stars. This is not strictly an inertial frame as this frame accelerates with the Earth’s orbit around the sun ($\sim 0.006 \text{ m/s}^2$), however for navigation on the surface of the Earth it is an adequate approximation of a true inertial frame.
- **Earth (e-frame):** This coordinate frame also defines its origin at the Earth’s center of mass, however the axes are aligned to and rotate with the Earth.
- **Navigation (n-frame):** The navigation frame is a local geographic frame with its origin at the center of the navigation system used. There are several common axes definitions used in literature; this dissertation will adopt the X-Y-Z axes along the directions of North-East-Down respectively. Note that the North, East, and Down directions change depending upon the system’s location on Earth.

- **Body (b-frame):** The body frame is a local coordinate frame also originating at the center of the navigation system used. However, the axes are fixed with respect to the defined roll, pitch, and yaw of the navigation system.




Figure 2.1: Illustration of i-, e-, and n- coordinate frames (adapted from [2]), where P denotes the position of the navigation system, O is the center of mass of the Earth, and ω_{ie} denotes the rotation of the Earth.

It is convention to express velocity as follows: v_β^α where the velocity of the system is expressed in the α -frame relative to the β -frame. Angular velocity is expressed as follows: $\omega_{\alpha\beta}^\gamma$ specifies the rotation of the β -frame with respect to the α -frame expressed in the γ -frame. For example ω_{nb}^b specifies the rotation of the body frame with respect to the navigation frame expressed in the body frame. The rotation of the Earth in the navigation frame can be expressed as:

$$\omega_{ie}^n = \begin{bmatrix} \Omega \cos L & 0 & -\Omega \sin L \end{bmatrix}^T \quad (2.1)$$

where L is latitude. Additionally, the rotation of the navigation frame due to the transport rate of the system is:

$$\omega_{en}^n = \begin{bmatrix} \frac{V_E}{R_P+h} & \frac{-V_N}{R_M+h} & \frac{-V_E \tan L}{R_P+h} \end{bmatrix}^T \quad (2.2)$$

where R_M and R_P are the radii of curvature of the Earth in the north-south and east-west directions respectively, V_E and V_N are the East and North components of v_e^n , and h is the altitude of the navigation system above the surface of the Earth. Finally, the rotation rate of the navigation frame relative to the inertial frame is:

$$\omega_{in}^n = \omega_{ie}^n + \omega_{en}^n \quad (2.3)$$

2.1.3 Coordinate Transformations

There are several methods available to convert attitude from one coordinate reference frame to another including: DCM, Euler Angles, and quaternion. Quaternion representation is typically used in navigation systems because attitude propagation is computationally efficient [12]. However, quaternion manipulation is not intuitive and this introduction will focus on Euler angles and DCM. A vector expressed in the α frame (r^α) can be converted from the α frame to the β frame as follows:

$$r^\beta = C_\alpha^\beta r^\alpha \quad (2.4)$$

where C_α^β is the DCM from the α frame to the β frame. Calculation of the DCM is well documented in literature. One useful coordinate transformation, for example, is from the e-frame to the n-frame, given as [12]:

$$C_e^m = \begin{bmatrix} -\sin L \cos \lambda & -\sin L \sin \lambda & \cos L \\ -\sin \lambda & \cos \lambda & 0 \\ -\cos L \cos \lambda & -\cos L \sin \lambda & -\sin L \end{bmatrix} \quad (2.5)$$

where L and λ denote the latitude and longitude of the navigation system. Note that DCMs are by definition orthonormal and therefore $(C_\alpha^\beta)^{-1} \triangleq C_\beta^\alpha = (C_\alpha^\beta)^T$ simplifying coordinate transformations.

2.1.4 Inertial Measurements and Gravity

Inertial sensors provide specific force (f_{ib}^b) and angular rate (ω_{ib}^b) information in the body frame relative to a true inertial reference frame. The specific force measurement can be expressed as:

$$\dot{v}_i^b = f_{ib}^b + g^b \quad (2.6)$$

where g is gravity. From this equation it is clear that accurate knowledge of gravity is essential for inertial navigation. Several gravity models have been developed for this purpose with the most popular having accuracy on the order of 5–10 μg [188, 187].

2.1.5 Velocity Update

The goal of this derivation is to find the velocity update ($\dot{v}_e^n$). First, the velocity update for the inertial frame (i-frame) with respect to the Earth ($\dot{v}_e^i$) is derived:

$$\left. \frac{dr}{dt} \right|_i = \left. \frac{dr}{dt} \right|_e + \omega_{ie} \times r \quad (2.7)$$

where r is the position vector. Note that this derivation is true in any coordinate frame, thus the frame is not defined here. Differentiating this expression, denoting $\left.\frac{dr}{dt}\right|_e = v_e$, and assuming the rotation rate of the Earth is constant ($\dot{\omega}_{ie} = 0$) yields:

$$\left.\frac{d^2r}{dt^2}\right|_i = \left.\frac{dv_e}{dt}\right|_e + \left.\frac{d}{dt}[\omega_{ie} \times r]\right|_i = \left.\frac{dv_e}{dt}\right|_e + 2\omega_{ie} \times v_e + \omega_{ie} \times (\omega_{ie} \times r) \quad (2.8)$$

The second part of the equation was derived by substituting the expression from Eq. 2.7. Note the $\omega_{ie} \times (\omega_{ie} \times r)$ term is the centripetal acceleration term due to the rotation of the Earth and is indistinguishable from gravity. It is usually combined as the “local” gravity:

$$g_l = g - \omega_{ie} \times (\omega_{ie} \times r) \quad (2.9)$$

Finally, substituting the expressions for specific force and local gravity yields:

$$\left.\frac{dv_e}{dt}\right|_e = f - 2\omega_{ie} \times v_e + g_l \quad (2.10)$$

Next, the velocity update is derived in the n-frame ($\dot{v}_e^n$):

$$\left.\frac{dv_e}{dt}\right|_n = \left.\frac{dv_e}{dt}\right|_e - (\omega_{in}) \times v_e \quad (2.11)$$

Substituting the expressions from Eqs. 2.10 and 2.3 together with the expression for specific force yields:

$$\left.\frac{dv_e}{dt}\right|_n = f + g_l - (2\omega_{ie} + \omega_{en}) \times v_e \quad (2.12)$$

Converting this into the navigation frame renders:

$$\left.\frac{dv_e}{dt}\right|_n = C_b^m f_{ib}^b + g_l^n - (2\omega_{ie}^n + \omega_{en}^n) \times v_e^n \quad (2.13)$$

Finally, this is the expression for the velocity update in the n-frame.

2.1.6 Attitude Update

Updating the attitude requires an expression for $\dot{C}_b^n$ in terms of the sensor measurements $(f_{ib}^b, \omega_{ib}^b)$. The attitude is updated using:

$$\dot{C}_b^n = C_b^n \Omega_{nb}^b \quad (2.14)$$

where Ω_{nb}^b is the skew-symmetric form of ω_{nb}^b and:

$$\omega_{nb}^b = \omega_{ib}^b - C_n^b \omega_{in}^n = \omega_{ib}^b - C_n^b (\omega_{ie}^n + \omega_{en}^n) \quad (2.15)$$

The second part of the equation was derived by substituting Eq. 2.3 for ω_{in}^n .

2.1.7 Error Model

A state space error model of strapdown inertial navigation can be constructed as follows with the following states [171]:

$$\delta \vec{x} = \left[\delta \vec{\theta}, \delta \vec{v}_n, \delta \vec{p}_n, \delta \vec{b}_g, \delta \vec{b}_a \right]^T \quad (2.16)$$

where $\delta \vec{\theta}$ is the three-axis attitude vector in the navigation frame, $\delta \vec{v}_n$ and $\delta \vec{p}_n$ are the velocity and position vectors along the North, East, and Down directions, and $\delta \vec{b}_g$ and $\delta \vec{b}_a$ are the gyroscope and accelerometer bias. An error propagation model can be constructed as follows [171]:

$$\delta\dot{\vec{x}} = \begin{bmatrix} 0 & 0 & 0 & -C_b^n & 0 \\ \begin{bmatrix} \vec{f}^n \times \end{bmatrix} & 0 & 0 & 0 & C_b^n \\ 0 & I & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \delta\vec{\theta} \\ \delta\vec{v}_n \\ \delta\vec{p}_n \\ \delta\vec{b}_g \\ \delta\vec{b}_a \end{bmatrix} + \begin{bmatrix} C_b^n \epsilon_{ARW} \\ C_b^n \epsilon_{VRW} \\ 0 \\ \epsilon_{RRW} \\ \epsilon_{AcRW} \end{bmatrix} \triangleq A(t)\delta\vec{x} + B(t) \quad (2.17)$$

where $\begin{bmatrix} \vec{f}^n \times \end{bmatrix}$ is the skew-symmetric matrix of the accelerometer output in the navigation frame; ϵ_{ARW} is the angle random walk (ARW) of the gyroscope, ϵ_{VRW} is the velocity random walk (VRW) of the accelerometer, ϵ_{RRW} is the rate random walk (RRW) of the gyroscope, and ϵ_{AcRW} is the acceleration random walk (AcRW) of the accelerometer. This model is developed in the context of pedestrian navigation, where the deterministic biases of the inertial sensors are assumed to vary slowly. Misalignment and scale factor errors are neglected. This model also assumes navigation will be performed for limited amounts of time and therefore neglects the rotation of the Earth and the transport rate. The sampling rates of INSs are typically at least a few hundred Hz; all pedestrian inertial navigation data collected in this dissertation was collected at 800 Hz. This rate is high enough to assume the step size is small, and therefore the continuous-time system $(A(t), B(t))$ can be approximated as a discrete-time system $(F(k), B(k))$ modeled as follows:

$$\delta\vec{x}(k+1) = F(k)\delta\vec{x}(k) + B(k) \quad (2.18)$$

where $F = \exp(A(t_k)\Delta t) \approx I + A(t_k)\Delta t$ and $B(k) = B(t_k)\Delta t$.

2.1.8 Zero-Velocity Update (ZUPT)

ZUPT has been an essential aiding element in pedestrian inertial navigation since the seminal paper by E. Foxlin [47, 48]. ZUPT is typically used only in pedestrian navigation, specifically

walking, since it leverages the brief stationary period during the gait cycle where the foot is on the ground, termed the stance-phase.

Detection of the stance-phase has been proposed through a variety of methods [51, 65, 189]. Most of these methods are based on the relative stability of the IMU as detected through the accelerometer and gyroscope. The most commonly used detector is the SHOE detector [51]:

$$T(\vec{z}_n) = \frac{1}{N} \sum_{k=n-N/2}^{n+N/2-1} \left(\frac{1}{\sigma_a^2} \left\| \vec{a}_k - g \frac{\vec{a}_n}{\|\vec{a}_n\|} \right\|^2 + \frac{1}{\sigma_\omega^2} \|\vec{\omega}_k\|^2 \right) < \gamma \quad (2.19)$$

where $\vec{a}_k$ and $\vec{\omega}_k$ are the accelerometer and gyroscope measurements at sample k , $\vec{a}_n$ is the accelerometer mean over the detection window of length N , g is the magnitude of gravity, σ_a^2 and σ_ω^2 are the accelerometer and gyroscope noise variances, respectively, and γ is the detection threshold. Stance phase is declared when the test statistic satisfies the inequality. When the stance-phase is detected, the velocity of the IMU is assumed to be zero, and this zero-velocity is used as a pseudo-measurement. If an EKF is applied to the navigation error model, the ZUPT pseudo-measurements can be integrated using the following measurement matrices:

$$\begin{cases} \delta \vec{z} = \begin{bmatrix} \delta v_N & \delta v_E & \delta v_D \end{bmatrix}^T \\ \mathbf{H} = \begin{bmatrix} \mathbf{O}_{3 \times 3} & \mathbf{I}_{3 \times 3} & \mathbf{O}_{3 \times 9} \end{bmatrix} \end{cases} \quad (2.20)$$

In this case $\delta \vec{z}$ is simply the current velocity estimate, since ZUPT assumes zero-velocity during stance. The measurement covariance ($\mathbf{R}$) is based on the velocity uncertainty of the foot during stance-phase. This is highly dependent on the type of locomotion (e.g., walking, crawling, or running); for the purposes of walking at a regular pace the velocity uncertainty is typically set to $\sigma_{\text{ZUPT}} = 0.02$ m/s [58]. The uncertainty is typically assumed equal in the North, East, and Down directions; therefore $\mathbf{R}(k) = \sigma_{\text{ZUPT}}^2 \mathbf{I}_{3 \times 3}$.

2.1.9 Extended Kalman Filter Implementation

The Kalman filter is one method of combining measurements from multiple sources [190]. However, to generalize this method to non-linear systems, the EKF linearizes the system model around the current state estimate and applies the standard Kalman filter. The prediction equations of the error-state Kalman filter are:

$$\begin{aligned} \text{A-priori state estimation:} \quad & \delta \bar{\vec{x}}(k+1) = \mathbf{F}(k) \delta \vec{x}(k) \\ \text{A-priori error covariance:} \quad & \bar{\mathbf{P}}(k+1) = \mathbf{F}(k) \mathbf{P}(k) \mathbf{F}^T(k) + \mathbf{Q}(k) \end{aligned} \quad (2.21)$$

where the accent $\bar{}$ denotes the a-priori or prediction estimate, $\mathbf{F}(k)$ is the state transition matrix at time k , and $\mathbf{Q}(k)$ is the process noise covariance. If an external measurement is available, the measurement is modeled as:

$$\vec{z}(k) = \mathbf{H}(k) \delta \vec{x}(k) + \vec{n}(k) \quad (2.22)$$

where $\mathbf{H}(k)$ is the measurement matrix and $\vec{n}(k)$ is the measurement noise, a zero-mean random vector with covariance $\mathbf{R}(k) = E[\vec{n}(k)\vec{n}^T(k)]$. The update step of the error-state Kalman filter is performed as follows:

$$\begin{aligned} \text{Innovation:} \quad & \vec{\nu}(k+1) = \delta \vec{z}(k+1) - \mathbf{H}(k+1) \delta \bar{\vec{x}}(k+1) \\ \text{Innovation covariance:} \quad & \mathbf{S}(k+1) = \mathbf{H}(k+1) \bar{\mathbf{P}}(k+1) \mathbf{H}^T(k+1) + \mathbf{R}(k+1) \\ \text{Kalman gain:} \quad & \mathbf{W}(k+1) = \bar{\mathbf{P}}(k+1) \mathbf{H}^T(k+1) \mathbf{S}^{-1}(k+1) \\ \text{A-posteriori state estimation:} \quad & \delta \hat{\vec{x}}(k+1) = \delta \bar{\vec{x}}(k+1) + \mathbf{W}(k+1) \vec{\nu}(k+1) \\ \text{A-posteriori error covariance:} \quad & \mathbf{P}(k+1) = (\mathbf{I} - \mathbf{W}(k+1) \mathbf{H}(k+1)) \bar{\mathbf{P}}(k+1) \end{aligned} \quad (2.23)$$

where $\vec{\nu}$ is the innovation, $\mathbf{S}$ is the innovation covariance, $\mathbf{W}$ is the Kalman gain, $\delta \hat{\vec{x}}$ is the a-posteriori state estimate, and $\mathbf{P}$ is the a-posteriori error covariance estimate.

Applying the EKF to the strapdown inertial navigation error model introduced in Section 2.1.7 requires defining how each system state is corrected given measurement information. The position, velocity, gyroscope bias, and accelerometer bias are corrected by subtracting the error estimate, $\vec{x}_c = \hat{\vec{x}} - \delta\vec{x}$, where $\vec{x}_c$ is the corrected state and $\hat{\vec{x}}$ is the estimated state. The attitude is corrected as follows [179]:

$$\mathbf{C}_{b,c}^n = (\mathbf{I} + \mathbf{\Psi}) \hat{\mathbf{C}}_b^n \quad (2.24)$$

where $\mathbf{C}_{b,c}^n$ is the corrected DCM from the body frame to the navigation frame, and $\mathbf{\Psi} = \vec{\psi} \times$ corresponding to $\vec{\psi} = \begin{bmatrix} \delta\alpha & \delta\beta & \delta\gamma \end{bmatrix}^T$. Finally, the latitude, longitude, and altitude are corrected as follows [179]:

$$\begin{cases} L_c = \hat{L} - \frac{\delta p_N}{R_0 + h} \\ \lambda_c = \hat{\lambda} - \frac{\delta p_E \sec L_c}{R_0 + h} \\ h_c = \hat{h} + \delta p_D \end{cases} \quad (2.25)$$

where L is latitude, λ is longitude, and h is the altitude. Note that $\delta\vec{p}$ was previously defined as the position error vector.

2.2 Coriolis Vibratory Gyroscopes

This section develops the background required to interpret the gyroscope results reported in Chapters 3 through 5. The Allan deviation, used throughout this dissertation to characterize gyroscope noise, is introduced first. The lumped-parameter model of a CVG is then developed, together with WA operation, to illustrate the importance of minimizing anisoeasticity and anisodamping and of maximizing the Q-factor for MEMS gyroscopes. The three mechanizations available to a CVG are compared next. The section closes with a review of

energy dissipation mechanisms in the context of the HRG and a description of the wafer-level fabrication process used to produce the DSG.

2.2.1 Allan Deviation

The Allan deviation represents the root-mean-square random drift error of a sensor as a function of averaging time and is commonly used to characterize noise in inertial sensors [191]. The analysis proceeds by dividing a stationary output record into K clusters of width τ , averaging the signal within each cluster, and computing the Allan variance as one half of the mean squared difference between successive cluster averages; repeating this computation over a sequence of τ values yields the $\sigma(\tau)$ curve [191]. When plotted on logarithmic axes, each noise mode appears as a region of characteristic slope, allowing the individual contributions to be identified and quantified, as illustrated in Fig. 2.2. The coefficients extracted in this way are the noise terms of the error model of Eq. 2.17 and the gyroscope parameters tabulated in Table 1.1.

Angle random walk ARW originates in white mechanical-thermal and thermo-electrical noise at frequencies well above the sampling rate, which appears as a flat power spectral density and a slope of $\tau^{-1/2}$ on the Allan deviation. When the rate signal is integrated to obtain angle, this white rate noise becomes a random walk of angle, conventionally expressed in $^{\circ}/\sqrt{\text{h}}$. Mechanical-thermal noise sets the theoretical noise floor of a rate gyroscope, making ARW directly influenced by the Q-factor of the resonator. This dependence is quantified in Section 2.2.3, and its consequences for resonator design are discussed in Section 1.4.2.

Bias stability Bias stability, termed bias instability in the Allan variance literature [191], describes fluctuation of the sensor bias and is generally attributed to electronics susceptible to random flicker. It exhibits a $1/f$ power spectral density below the cutoff frequency and appears as a region of zero slope on the Allan deviation, so the parameter is read as the minimum of the $\sigma(\tau)$ curve, in $^{\circ}/\text{h}$.

Rate random walk RRW exhibits a $1/f^2$ power spectral density and appears with a slope of $\tau^{+1/2}$ on the Allan deviation. Its effect can be understood as a white noise applied to the bias of the sensor, so that the bias itself executes a random walk. The physical origin of RRW is generally not well established.




Figure 2.2: Schematic of log-log plot of Allan Deviation.

2.2.2 Lumped-Parameter Model of a Vibratory Gyroscope




Figure 2.3: Illustration of the lumped 2-DoF mass resonator used to represent gyroscopes.

The equations of motion for a two degree-of-freedom lumped mass resonator, also termed the equations of motion for the generic vibratory gyroscope, are [110, 93]:

$$\begin{bmatrix} \ddot{x} \\ \ddot{y} \end{bmatrix} + Z \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} + K \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} F_x \\ F_y \end{bmatrix} \quad (2.26)$$

where x and y denote displacement of the proof mass along the reference X and Y axes as illustrated in Figure 2.3, and F_x and F_y are the control forces along the respective axes. The Z matrix includes the impact of damping and Coriolis forces:

$$Z = \begin{bmatrix} \frac{2}{\tau} + \Delta(\frac{1}{\tau}) \cos 2\theta_\tau & \Delta(\frac{1}{\tau}) \sin 2\theta_\tau - 2\kappa\Omega \\ \Delta(\frac{1}{\tau}) \sin 2\theta_\tau + 2\kappa\Omega & \frac{2}{\tau} - \Delta(\frac{1}{\tau}) \cos 2\theta_\tau \end{bmatrix} \quad (2.27)$$

where $\frac{1}{\tau} = \frac{1}{2} \left(\frac{1}{\tau_1} + \frac{1}{\tau_2} \right)$ and $\Delta(\frac{1}{\tau}) = \frac{1}{\tau_1} - \frac{1}{\tau_2}$, and τ_1 and τ_2 are the damping time constants along the respective principal axes. These axes do not necessarily align with the reference X and Y axes and are rotated by an angle θ_τ as shown in Figure 2.3. Ω is the input angular rate to the gyroscope and κ is the angular gain factor, also termed Bryan's factor [192]. The K matrix includes the elasticity components:

$$K = \begin{bmatrix} \omega^2 - \omega\Delta\omega \cos 2\theta_\omega & -\omega\Delta\omega \sin 2\theta_\omega \\ -\omega\Delta\omega \sin 2\theta_\omega & \omega^2 + \omega\Delta\omega \cos 2\theta_\omega \end{bmatrix} \quad (2.28)$$

where $\omega^2 = \frac{\omega_1^2 + \omega_2^2}{2}$, $\omega\Delta\omega = \frac{\omega_1^2 - \omega_2^2}{2}$, and ω_1 and ω_2 are the natural frequencies along their respective directions. Again, these axes do not necessarily align with the reference X and Y directions and are rotated by an angle θ_ω . Note that the sensitivity of the gyroscope to linear acceleration (g-sensitivity) is neglected in this formulation.

In an ideal case ($\omega_1 = \omega_2$ and $\tau_1 = \tau_2$) this model reduces to:

$$Z = \begin{bmatrix} \frac{2}{\tau} & -2\kappa\Omega \\ 2\kappa\Omega & \frac{2}{\tau} \end{bmatrix} \quad K = \begin{bmatrix} \omega^2 & 0 \\ 0 & \omega^2 \end{bmatrix} \quad (2.29)$$

In this ideal scenario, WA implementation with open-loop control would simply track the precession angle of the elliptical trajectory of the mass. This angle is exactly the negative of the angle of rotation [92]:

$$\theta = \theta_0 - \kappa \int \Omega dt \quad (2.30)$$

where θ is the pattern angle, θ_0 is the initial pattern angle, and Ω is the input angular rate. However, as illustrated by Lynch [110], the pattern angle of the pendulum variables in reality can be approximated by:

$$\dot{\theta} = -\kappa\Omega + \frac{1}{2}\Delta \left(\frac{1}{\tau} \right) \sin 2(\theta - \theta_\tau) + \frac{1}{2}\Delta\omega \cos 2(\theta - \theta_\omega) \frac{\mathcal{Q}}{E} \quad (2.31)$$

where E and $\mathcal{Q}$ represent the energy and quadrature in the system, relating to the elliptical trajectory of the gyroscope. Equation 2.31 illustrates how critical a uniform resonator is to WA operation. The drift from anisodamping ($\Delta(\frac{1}{\tau})$) and anisoelasticity ($\Delta\omega$) is indistinguishable from input angular rate. While closed-loop control can mitigate the impact of non-uniformity, the control electronics also inject noise and error indistinguishable from the input angular rate [93]. In addition, the design of closed-loop control is difficult because anisoelasticity and anisodamping change as a function of temperature.

2.2.3 Mechanization of the Coriolis Vibratory Gyroscope

Three gyroscope mechanizations are commonly implemented. In each, the drive axis is sustained at constant amplitude by an AGC loop while a phase-locked loop tracks its resonant frequency. In the two rate mechanizations, the fundamental limit on resolution is set by mechanical-thermal, or Brownian, noise of the resonator itself [193].

In open-loop rate mode, no control force is applied to the sense axis, and the sense-axis displacement is demodulated directly to recover the input rate. Requiring no sense-axis feedback, this is the mode used for the demonstrations of Chapters 3 and 4. Its primary limitation follows from the same absence of feedback: the sense axis responds on its own ringdown timescale, so the measurement bandwidth is inversely proportional to the Q-factor of that axis. The frequency split additionally attenuates the open-loop scale factor.

In FRB mode, the sense axis is instead nulled by a feedback force, and the force required to hold it at zero is taken as the output [110, 193]. In this implementation, the measurement bandwidth is set only by the update rate of the control loop rather than by the ringdown time. The ultimate limit of ARW for this mode is set by Brownian motion and is described as [193]:

$$\text{ARW}_{\text{FRB}} \approx \sqrt{\frac{k_B T \omega_y}{A^2 M \omega_x^2 Q_y} \left(1 + \frac{Q_y^2 (\omega_y^2 - \omega_x^2)^2}{\omega_x^2 \omega_y^2} \right)} \times 3437.7. \quad (2.32)$$

where k_B is the Boltzmann constant, T is the absolute temperature, A is the sustained amplitude of motion on the drive axis, M is the effective modal mass, ω_x and ω_y are the drive- and sense-axis natural frequencies, Q_y is the Q-factor of the sense axis, and the factor 3437.7 converts from $\text{rad}/\sqrt{\text{s}}$ to $^\circ/\sqrt{\text{h}}$.

Equation 2.32 is derived for a lumped two degree-of-freedom resonator in which the Coriolis coupling is assumed to be unity; for a shell resonator the coupling carries the angular gain

factor κ of Eq. 2.27, so the noise floor is inversely proportional to κ [192]. Because the angular gain is a property of the mode shape rather than of the fabrication process, it enters as a fixed multiplier for a given mode and does not alter the scaling with the remaining parameters.

In the WA mode, developed in Section 2.2.2, no nulling force is applied and the pattern angle is allowed to precess freely, which yields effectively unlimited dynamic range and an output that does not integrate white rate noise, at the cost of the stringent symmetry requirements expressed by Eq. 2.31.

For a degenerate resonator, $\omega_x = \omega_y = \omega_0$, the noise floor scales as $Q^{-1/2}$, so an order-of-magnitude increase in Q-factor improves ARW by a factor of approximately three, and the dissipation mechanisms enumerated in Section 2.2.4 therefore act directly on the noise performance of the sensor. The noise floor also scales as A^{-1} , so it improves in proportion to the displacement amplitude that the transduction method can sustain at the antinode, which is the quantity maximized in Chapter 5.

2.2.4 Energy Loss Mechanisms

The Q-factor of a resonator is defined by the ratio of energy lost per cycle to the total energy of the system. The total Q-factor of a resonator (Q_{tot}) is determined by the damping mechanism with the greatest energy loss. In a resonator the mechanisms of energy loss contributing to the overall Q-factor are as follows:

$$\frac{1}{Q_{tot}} = \frac{1}{Q_{visc}} + \frac{1}{Q_{TED}} + \frac{1}{Q_{anchor}} + \frac{1}{Q_{surface}} + \frac{1}{Q_{int}} + \frac{1}{Q_{mode}} + \frac{1}{Q_{other}} \quad (2.33)$$

- Viscous damping (Q_{visc}) is due to vibration of the resonator in a fluid, such as air, and can be effectively minimized by packaging the resonator in vacuum [194].

- TED (Q_{TED}) can be thought of as internal friction in a solid [195, 196]. It occurs when displacement of a resonator causes strain gradients in the material; this energy can be dissipated by irreversible flow of heat driven by local temperature gradients [114]. The coupling of strain and temperature fields is governed by the CTE of the material, and smaller values of CTE minimize TED. As noted previously, this is why fused silica or ultra low expansion titania silica glass (ULE TSG) is used for the HRG in preference to silicon.
- Anchor loss (Q_{anchor}) is caused by energy loss of acoustic waves radiated into the supporting substrate via the anchors [197, 198]. Methods to minimize anchor loss include placing anchors at nodes of zero displacement of the primary mode and balancing the loads and torques at the anchor [111]. Analytical modeling of anchor loss is challenging for complex structures and is generally evaluated using FEA with perfectly matched layer (PML) substrates, which are designed to absorb energy without reflecting it back into the resonator.
- Surface loss ($Q_{surface}$) is generally caused by surface defects and has been linked to surface roughness [125], coating with a metal layer for capacitive actuation and detection [126], and subsurface damage [116, 117]. In addition, contamination by adsorption of water or organic compounds can lead to significant energy loss. Surface loss is of particular importance in microstructures because the surface-to-volume ratio is large.
- Internal energy loss (Q_{int}) is related to imperfections of the resonator material such as defects, impurities, and residual stress, the last of which is particularly important in glassblowing due to large and rapid temperature changes. Thermal annealing has been demonstrated to alleviate some of the internal stress in the HRG after glassblowing [124, 125, 199].

- Mode coupling energy loss (Q_{mode}) is prominent when the difference in frequency between the desired actuation mode and parasitic modes is small. In this case, some of the energy from the desired mode is transferred to a parasitic mode and dissipated. Designing a large frequency split between the desired and parasitic modes is generally challenging for flat structures, as out-of-plane modes are typically low in frequency without proper care. 3-D structures are better able to mitigate parasitic out-of-plane modes as they generally have greater out-of-plane stiffness.
- Other sources of damping (Q_{other}) include Akhiezer damping, caused by the scattering of phonons, which is independent of geometry. However, the resonant frequency of HRGs is typically less than 100 kHz, where Akhiezer damping is not prominent and can generally be neglected [200].

2.2.5 Fabrication of the Dual-Shell Hemispherical Resonator Gyroscope

Wafer-level glassblowing of fused silica shells was first demonstrated by Eklund and Shkel [95], in which a sealed cavity beneath a glass layer is brought above the softening point of the glass so that the trapped gas expands and inflates the softened layer into a shell. Because the shape is set by surface tension rather than by lithographic or etch tolerances, the resulting structure achieves a degree of symmetry and surface quality not attainable with conventional glass microfabrication [97, 92]. The DSG extends this concept to a pair of concentric shells: the fabrication process was disclosed in [165] and first demonstrated in [3].

The fabrication process is illustrated in Fig. 2.4. Three high-purity fused silica wafers (Corning HPFS 7980) are used: a 2 mm thick substrate wafer, etched in step (b); a 150 μm thick device wafer, which forms the inner shell after glassblowing; and a 500 μm thick cap wafer, etched in step (a) [3]. On the substrate and cap wafers, a 2 μm undoped polysilicon film is deposited by low-pressure chemical vapor deposition to serve as a hard mask and patterned

by reactive ion etching, after which the cavities are formed by isotropic wet etching in hydrofluoric acid (Fig. 2.4(a,b)). The hard mask is then removed and the wafers chemically cleaned prior to bonding. The substrate, device, and cap wafers are joined by plasma-assisted direct bonding (Fig. 2.4(c,d)), and the completed triple stack is annealed to complete bond formation, sealing an array of aligned concentric cavities [3].

The bonded stack is then glassblown in a rapid thermal processing furnace at 1700 °C, Fig. 2.4(e), above the softening point of fused silica, where the viscosity drops sharply and the material behaves as a viscous fluid. As the temperature rises, the gas sealed within the cavities builds pressure, and the combined effect of this pressure and the reduced viscosity expands the cavities into a three-dimensional dual-shell structure with a self-centered stem. During deformation, surface tension minimizes the surface area and drives the structure toward an axisymmetric geometry [3].

The inner shell is released by parallel back-lapping of the substrate (Fig. 2.4(f)), which removes the material below the shell rim and defines the free edge of the resonator [3, 109]. The bulk of the substrate is removed using diamond lapping films; before the substrate is fully removed, the cavities are filled with a temporary adhesive to support the thin shell through the remainder of the release. The rim is then polished, the adhesive dissolved in a solvent bath, and the resonator chemically cleaned, yielding polished surfaces with a measured roughness on the order of 0.3 nm [3, 201].

The result is a monolithic fused silica structure in which the inner shell serves as the vibrating element while the outer shell acts simultaneously as the anchor and as a protective casing, shielding the resonator from the external environment [99]. This double-ended anchor architecture increases the robustness of the structure to mechanical shock; DSGs of this configuration have demonstrated continuous operation under 50 000 g shock [202], and the outer shell additionally provides a path to die-level vacuum packaging [3]. These advantages come at the cost of increased process complexity relative to a single-shell resonator, since

each additional wafer and bonding step introduces a further opportunity for misalignment. For conventional electrostatic transduction, the resonator is then metallized and assembled to a planar electrode substrate, shown at right in Fig. 2.4.




Figure 2.4: Fabrication process of the DSG: (a,b) cavity pre-etching, (c,d) plasma-assisted bonding of the triple wafer stack, (e) high-temperature glassblowing, (f) release of the inner shell by lapping. At right, the released resonator assembled to a planar electrode substrate for electrostatic transduction.

Part I

Fused-Silica Dual-Shell Gyroscope Development

Chapter 3

Laser-Defined Electrodes for Fused Silica Resonators

This chapter presents a fabrication approach in which the capacitive electrodes of the DSG are defined directly within the bulk fused silica of the sensor, eliminating the separately fabricated electrode substrate and the assembly step traditional methods require.¹ In the conventional process, the substrate remaining after glassblowing is removed entirely by back-lapping to release the inner shell, and the released resonator is then assembled to a separately fabricated electrode substrate, Fig. 2.4 [3, 126]. This approach, shared by most HRG architectures, requires a second fabrication process and alignment at micrometer-level tolerance, and it introduces capacitance non-uniformity through tilt and misalignment [204]. In the

¹This chapter is based on the following peer-reviewed publications and patent application:

- A. R. Parrish, E. Wong, M. H. Asadian, L. Meira-Lopez, and A. M. Shkel, “Self-Aligned, High-Aspect Ratio Electrodes for Fused Silica Resonators Digitally Defined by Laser Modification of Material,” *IEEE/ASME Journal of Microelectromechanical Systems*, 2026. Early Access [203].
- A. R. Parrish, L. Meira-Lopez, M. H. Asadian, E. Wong, and A. M. Shkel, “Laser Patterned Self-Aligned Electrodes for Dual-Shell Hemispherical Resonator Gyroscope,” In *Solid-State Sensor Actuators and Microsystems Workshop (TRANSDUCERS)*, Orlando, FL, USA, Jun. 2025 [138].
- A. R. Parrish, M. H. Asadian, L. Meira-Lopez, and A. M. Shkel, “Laser patterned self-aligned electrodes for hemispherical resonator gyroscope,” U.S. Patent Application, UC Case No. 2025-819-1, 2025 [4].

approach reported here, back-lapping is instead stopped before the inner shell is released, and the capacitive gaps are digitally patterned into the bulk fused silica remaining between the two shells using FLICE [205]. Because the gaps are formed within the resonator material itself, the electrodes are inherently self-aligned and no assembly is required. This architecture is termed the IPE-DSG, and the two configurations are compared in Fig. 3.1. The IPE-DSG eliminates the assembly step, but the resonating element is still metallized for capacitive transduction, and metallization bounds the achievable Q-factor regardless of how well the remaining loss channels are controlled; removing that coating is the subject of Chapter 4. In doing so, the chapter addresses **RQ1a**, how capacitive electrodes can be defined for a DSG without an assembly step and what governs the minimum capacitive gap and maximum electrode aspect ratio achievable by the process. Section 3.1 describes the fabrication process, including the two FLICE patterning steps and the metallization approaches evaluated. Section 3.2 reports an experimental study of the FLICE parameters governing the achievable capacitive gap, and Section 3.3 develops an analytical model of the resulting capacitance. Characterization of fabricated devices is presented in Section 3.4, followed by demonstration of open-loop rate gyroscope operation in Section 3.5. Section 3.6 discusses the merits and remaining limitations of the approach. Finally, the chapter is concluded with Section 3.8.




Figure 3.1: Comparison of DSG assembly methods: (a) eutectic bonding to separate electrode substrate (prior art) [3], and (b) FLICE-patterned in-plane electrodes in the bulk of the fused silica (this work).




Figure 3.2: Fabrication process of IPE-DSG [4]: (a/b) cavity pre-etching of cap and substrate wafers, (c/d) triple-stack plasma-assisted wafer bonding and annealing, (e) glassblowing in a high-temperature furnace at 1700 °C, (f) substrate removal by lapping until 300–400 µm of material remains before the second cavity is released, (g) FLICE patterning of KOH release trenches and electrical isolation with KOH etching duration of >10 hours, (h) FLICE patterning of capacitive gap followed by KOH etching for approximately 6.5 hours, (i) metallization with Pt via PE-ALD, (j) electrical isolation using laser ablation of PE-ALD, and (k) wire bonding to actuation PCB. Note: steps (g-j) are a zoomed-in view of the red box in step (f).




Figure 3.3: Cross-section of partially released dual-shell resonator illustrating the region where the electrodes are formed with the FLICE process.

3.1 Fabrication Process

The IPE-DSG fabrication process is illustrated in Fig. 3.2. Steps (a) through (e) follow the wafer-level glassblowing sequence described in Section 2.2.5, performed here on Corning 7980 high-purity fused silica with substrate, device, and cap wafers of 2 mm, 150 μm , and 500 μm thickness, bonded at 400 $^{\circ}\text{C}$ for 48 hours and glassblown at 1700 $^{\circ}\text{C}$. In step (f), the release and electrode defining process began. The shell was lapped until $\sim 300\text{--}400$ μm of material remained between the inner and outer shells, which defined the electrode thickness. At this stage, the region between the inner and outer shells was accessible for electrode patterning, Fig. 3.3. The shells were then annealed at 920 $^{\circ}\text{C}$ for 30 hours before FLICE patterning to relieve residual thermal stress from the glassblowing process [206]. Next, two FLICE patterning steps were performed (g,h). These two steps were separated to minimize the KOH etching time when the capacitive gap was formed via FLICE, thereby maximizing the capacitance. In our implementation, the KOH etching of step (g) was ~ 10 hours and step (h) was 6.5 hours, but other recipes are possible.

Next, the IPE-DSG was trimmed to correct as-fabricated asymmetry by selectively removing material on the rim with laser ablation, performed using a 1030 nm wavelength Light Conversion CARBIDE femtosecond laser source, with a repetition rate of 60 kHz and a pulse energy of 29 μJ , discussed further in Section 4.3.2. The IPE-DSG was then metallized, step




Figure 3.4: Fully Released IPE-DSG with features labeled. A close-up view of the capacitive gap is provided on the right illustrating the inner shell, capacitive gap, and electrical isolation. The perpendicular lines for electrical isolation are patterned to ensure alignment between the two separate FLICE patterning steps (steps (g) and (h) of Fig. 3.2)

(i), and electrical isolation of electrodes was performed using laser ablation with a 1030 nm laser source, repetition rate of 60 kHz, and pulse energy of 8.3 μJ . To remove adsorbed moisture, the IPE-DSGs were baked at 400 $^{\circ}\text{C}$ for 24 hours in vacuum <1 μTorr . During the experimental evaluation, the electrodes were wire bonded to a PCB, step (k), for front-end driving and detection. For the final product, the structure would be bonded to an LCC package or other packaging options.

3.1.1 FLICE Process

FLICE, reviewed in Section 1.4.3, was implemented here as a two-step process [132]. First, a tightly focused femtosecond laser patterned the fused silica material to inscribe a cavity from bottom to top with a z-step size of 4 μm . Then, the IPE-DSG was placed in an 8 M KOH solution at 85 $^{\circ}\text{C}$ for etching. Laser modification was performed using a 1030 nm wavelength Light Conversion CARBIDE CB5 femtosecond laser source, with a repetition rate of 250 kHz and nominal pulse energy of 0.684 μJ for fine features and 0.936 μJ for bulk features, while the sample was moved horizontally at 5 mm/s feed rate. A Mitutoyo 50 $\times$ magnification lens (NA=0.42) focused the collimated beam deep in the bulk of the transparent material. Laser polarization was kept perpendicular to the patterning direction to maximize selectivity [134].

Although the feed rate used in our process (5 mm/s) was relatively slow for high-throughput manufacturing, this limitation is primarily hardware dependent rather than fundamental to the FLICE process. Studies have demonstrated that precision and selectivity can be maintained or improved with feed speeds of 200 mm/s [133], and the FLICE process has been successfully demonstrated at speeds of up to 3 m/s [207].

3.1.2 In-Plane Electrodes Patterning

The first FLICE patterning step (step (g), Fig. 3.2) defined the KOH release trenches and electrical isolation. The KOH etching time during this step was not required to be precisely controlled and took approximately 10–15 hours. The second FLICE step, (h), was used to form the capacitive gap. The KOH etching time depended on the electrode thickness and was typically 6.5 hours.

Two KOH release trenches (fabricated as a cylinder of radius 150 μm) were patterned with 180° separation (Figs. 3.2(g) and 3.4). The trenches served two functions. First, after KOH wet etching, the trench coupled with ultrasonic cleaning enabled efficient removal of KOH from the cavity. Second, in the subsequent wet etching of the capacitive gap, (h), the trenches enabled KOH to etch the capacitive gap simultaneously from above and below the glass material, minimizing etching time and the resulting capacitive gap. In order to displace the air in the cavity between the inner and outer shells, the sample was initially dipped in water with one trench facing down. Then, ultrasonic agitation paired with surfactants was used to lower surface tension to penetrate the 300 μm diameter trench. Once the cavity was filled, the IPE-DSG was moved to KOH etching solution.

3.1.3 Metallization

Two metallization approaches were investigated to complete step (i) of Fig. 3.2. The first employed magnetron sputtering to deposit 2–16 nm Cr-Au using a Denton DV-502M AC

Sputter Coating System operated in 20 W RF mode at a chamber pressure of 4.8 mTorr. However, achieving conformal coverage using magnetron sputtering is challenging for structures with aspect ratios as low as 5:1 [150]. To address this limitation, Plasma-Enhanced Atomic Layer Deposition (PE-ALD) was explored due to its superior conformality in high aspect ratio trenches. An Ultratech/Cambridge NanoTech Fiji PE-ALD system was used to deposit 10 nm thick Pt films. The process utilized (MeCp)PtMe₃ as the precursor with O₂ plasma as the co-reactant. Deposition was performed at a substrate temperature of 300 °C for 300 cycles.

3.2 Optimization of FLICE Capacitive Gap

An experimental study was conducted to investigate the parameters of FLICE in achieving small, high aspect-ratio capacitive gaps. Two-dimensional channels were patterned using FLICE through flat, 300 µm thick fused silica samples. This structure was selected to closely mimic the etching conditions of the IPE-DSG capacitive gap, where KOH release trenches enabled simultaneous etching from both the top and bottom surfaces.

All samples were cross-sectioned using lapping and fine polishing to enable direct observation of the channel width within the bulk of the fused silica. After etching, the samples were cleaned using Crest Ultrasonic P360D-HF-132 at 70 °C for 20 minutes to remove residual KOH, followed by piranha cleaning. Channel widths were measured using SEM images of the cross sections at the top, bottom, and middle of the channel, with three measurements taken at each location per sample. In all cases, the parameters were kept constant: feed rate of 5 mm/s, z-stepping distance of 6 µm, 1030 nm femtosecond laser source with repetition rate of 250 kHz. The following subsections report on the effects of laser pulse energy and KOH etching duration on channel width to identify optimal FLICE parameters that balance complete release with minimal capacitive gap dimensions.




Figure 3.5: Illustration of FLICE patterning a 2D channel from bottom to top. The laser follows a serpentine pattern stepping in the z-direction in $6\text{ }\mu\text{m}$ increments. The cross-sections discussed in Section 3.2 were prepared by lapping the material, as indicated in green.

3.2.1 Pulse Energy Analysis

To determine optimal pulse energies for 2D channel formation, we tested a range of energies from $0.504\text{ }\mu\text{J}$ to $0.684\text{ }\mu\text{J}$ in $0.018\text{ }\mu\text{J}$ increments. It was observed that lowering the pulse energy any further did not produce channels that penetrated the entire $300\text{ }\mu\text{m}$ thick fused silica sample within the etching time of 6.5 hours, selected to be representative of the time to fully release the capacitive gap of the IPE-DSG. Three samples were prepared, each with 11 trenches patterned using the described pulse energies, and were simultaneously etched in KOH solution for 6.5 hours. A summary of the channel widths across all three samples is provided in Figure 3.7. The center, which constitutes the vast majority of the channel width, is observed to increase from $3.05\text{ }\mu\text{m}$ on average for $0.558\text{ }\mu\text{J}$ to $4.6\text{ }\mu\text{m}$ with $0.684\text{ }\mu\text{J}$ pulse energy. This suggests that using lower pulse energy will lead to smaller channel widths and larger capacitance. The top/bottom values were on average $1.15\text{ }\mu\text{m}$ larger than the center channel width. Selective laser etching relies on the formation of Type II laser modifications, self-ordered periodic nanogratings that facilitate etchant penetration into the bulk material [130, 129]. At pulse energies exceeding the Type II threshold, observed to be near $0.660\text{ }\mu\text{J}$ in [129], the modification transitions to Type III, characterized by structural disruption that disorders these nanogratings and inhibits etchant transport.




Figure 3.6: SEM cross-sectional images of 2D channels patterned in 300 μm -thick fused silica samples using varying laser pulse energies. All samples were etched in KOH for 6.5 hours. Cross sections were prepared as described in Fig. 3.5.

3.2.2 Etching Duration Analysis

An investigation of the effect of KOH etching duration on channel width is detailed in this section. Six cross-sectional samples were prepared, each FLICE patterned with two 2D channels for three laser pulse energies, 0.540 μJ , 0.612 μJ , and 0.684 μJ . The duration of etching was varied from 4 to 14 hours in increments of 2 hours, with each sample immersed in 8 M KOH at 85 $^{\circ}\text{C}$. An illustration of the channels for each etching duration with 0.540 μJ pulse energy and the average and standard deviation of the channel widths is shown in Figure 3.8. The mean etching rate of the center width was 0.76 $\mu\text{m/hr}$, averaged across all three laser pulse energies.

This study illustrates that the channel width is critically dependent on the etching duration, with etching duration having a more significant effect than laser pulse energy. Based on these results, the optimal FLICE patterning strategy involved balancing two competing objectives:




Figure 3.7: Measured gap widths for 2D channels formed in three 300 μm -thick fused silica samples as a function of laser pulse energy. Each sample was measured at three locations, resulting in a total of nine data points. The plot displays the mean and standard deviation for each pulse energy. Only one sample fully released for the lowest pulse energy and the data reported for the center corresponds to this sample.

achieving complete channel release and minimizing channel width (maximizing capacitance). While the lowest pulse energy (0.540 μJ) produced the smallest channel widths, it failed to fully release channels in two of the three samples within 6.5 hours. Notably, the highest pulse energy sample (0.684 μJ) also failed to fully release at the shortest etching duration of 4 hours, consistent with the proximity of this energy to the Type II/III transition threshold discussed in Section 3.2.1. Therefore, a pulse energy slightly higher than the minimum value tested, such as 0.558 μJ , could be preferred to ensure reliable release while maintaining small channel widths.

3.3 Capacitance Analysis

The total capacitance formed by the IPE-DSG is calculated as follows:

$$C = \frac{2\pi\epsilon RT}{g} \quad (3.1)$$




Figure 3.8: (Top) Cross section of six 300 μm -thick fused silica samples prepared with 0.54 μJ pulse energy with KOH etching time ranging from four to fourteen hours in two hour increments. (Bottom) Summary of gaps observed for three laser pulse energies: 0.540 μJ , 0.612 μJ , and 0.684 μJ . Each etching time is a separate sample. Note: no channels patterned with 0.684 μJ pulse energy and 4 hours etch time fully released.




Figure 3.9: Analytical capacitance of the IPE-DSG as a function of capacitive gap and electrode thickness.

where ϵ is the permittivity of free space (8.854×10^{-12} F/m), R is the radius of the inscribed capacitive gap, T is the thickness of the inscribed electrode, and g is the capacitive gap.

In principle, the total capacitance is limited only by the FLICE patterning capabilities and the metallization conformality. As demonstrated in Section 3.2, a capacitive gap width of <2 μm was achieved for optimized FLICE parameters with a height of 300 μm , as illustrated in Fig. 3.8. Additionally, the radius of the patterned capacitive gap used in this analysis was $R = 3.15$ mm. The capacitance that can be achieved by the proposed fabrication process is illustrated in Fig. 3.9. For example, using a conservative capacitive gap of 2.5 μm , slightly larger than the best demonstrated result of 2 μm from Section 3.2, with 300 μm electrode thickness, the total capacitance can exceed 20 pF.

3.4 Device Characterization

Table 3.1 summarizes the characteristics of two fabricated IPE-DSG samples, referred to as Samples A and B. The samples were characterized using a bulk piezo stack for excitation and a Laser Doppler Vibrometer (LDV) for detection. Given the frequency f , the ringdown time-constant τ was measured and used to determine the Q-factor, where $Q = \pi f \tau$. The as-fabricated Q-factor of all samples was measured to be greater than or equal to 2.8 million

Table 3.1: Experimentally measured parameters of IPE-DSG samples

Device	A	B
Resonant Freq [kHz]	20.85	29.15
As-Fabricated Δf [Hz]	11.6	137.2
Trimmed Δf [Hz]	3.7	1.1
As-Fabricated Q-Factor	4.6M/4.3M	2.8M/3.0M
Q-factor after Δf Trimming	1.1M/2.1M	788k/598k
Metallized (ALD) Q-Factor	980k/1.1M	964k/775k
Inner Shell Resistance (Cr-Au) [Ω]	168	-
Inner Shell Resistance (ALD) [Ω]	68.4	44.8
Total Usable Capacitance (Cr-Au) [pF]	2.55	-
Total Usable Capacitance (ALD) [pF]	4.68	3.41

with a maximum of 4.57 million. This Q-factor is the highest among previously reported DSGs [6], suggesting that FLICE patterning does not substantially increase surface energy losses. Ringdown time-constant measurement for sample A is reported in Fig. 3.10 with decay times of 69.8 s and 65.0 s for the n=2 wineglass mode. A significant reason for the high Q-factors is the incorporation of a hemitoroidal design which mitigates anchor loss by reducing the stress on the junction between the inner and outer shells in the out-of-plane direction [6]. After laser ablation trimming of frequency split to less than 6 Hz for each device, the Q-factors were observed to decrease by an average of 70%, likely due to material damage, Section 4.3.2. Furthermore, after metallization, it was observed that the Q-factors further decreased for sample A, but surprisingly increased for sample B. We hypothesize that sample B's Q-factor increased after metallization due to the substantial material damage from the laser ablation used to correct its 132 Hz frequency split: the deposited metal likely smoothed the damaged material sites, increasing the Q-factor.

3.4.1 Cross-Section of IPE-DSG

A cross-section of an IPE-DSG was examined using SEM. Fig. 3.11 illustrates a cross-section of the full IPE-DSG design. The inner shell can be seen in the bottom-left and the outer shell in the bottom-center. Figure 3.12 presents close-up SEM images of the left and right capacitive gaps, with further insets showing the top, middle, and bottom sections of each




Figure 3.10: Experimental Q-factor measurement of sample A to determine the Q-factor of the n=2 wineglass mode as-fabricated.

gap. The bulk capacitive gap was observed to be 5.5–6 μm with increased damage at the top of the capacitive gap increasing to 8.5–10 μm . The height of each capacitive gap was observed to be 331 μm and 384 μm . A maximum aspect ratio of 64:1 was achieved experimentally on an IPE-DSG, verifying the potential for achieving large capacitance using the proposed process.

3.4.2 Metallization

Sample A was first coated with 2–16 nm Cr-Au via AC sputtering to establish a baseline for comparison. Afterwards, Cr-Au was completely removed by wet etching followed by 10 nm PE-ALD of Pt to evaluate the improvement in conformality and resistance. It was observed that the inner shell resistance decreased from 168 Ω to 68.4 Ω with only 10 nm Pt as opposed to 2–16 nm Cr-Au. Lower inner shell resistance led to increased SNR and greatly affected gyroscope noise performance [143]. This improvement is attributed to the superior conformality of PE-ALD, which eliminated the shadowing effects of sputtering on high-aspect-ratio structures [152]. Despite Au having lower bulk resistivity than Pt, the uniform and continuous PE-ALD Pt coating achieved lower overall resistance than the discontinuous sputtered Au film.




Figure 3.11: Cross-section of a partially released dual-shell resonator illustrating the in-plane electrode design. The inner shell is shown on the left.

3.4.3 Capacitance

Measurement of capacitance of IPE-DSG inevitably measures both usable capacitance formed by FLICE patterning and capacitance due to the coplanar top-region of the inner shell and electrode. The coplanar capacitance was estimated as follows [208]:

$$C_{coplanar} = \frac{L \epsilon_r \ln(2 \frac{1+\sqrt{k'}}{1-\sqrt{k'}})}{377 \pi c} \quad (3.2)$$

where $k = \frac{g_{top}}{g_{top}+2W}$ and $k' = \sqrt{1-k^2}$, ϵ_r is the relative permittivity of free space, c is the speed of light in vacuum (2.998×10^8 m/s), g_{top} is the gap measured at the top of the capacitive gap (the average value of $9.25 \mu\text{m}$ was taken from Fig. 3.12), L is the length of the electrode pair (0.83 mm for one electrode) and W is the width, assumed to be $250 \mu\text{m}$ as this is the width of the inner shell. The estimated coplanar capacitance was determined to be 34.1 fF per electrode and was uniformly subtracted from all presented capacitance measurements.




Figure 3.12: SEM images of the left and right side cross-sections. The three images on each side illustrate the top, middle, and bottom of the capacitive gaps.

The total capacitance of sample A was measured to be 2.55 pF with Cr-Au and increased by 84% to 4.68 pF with ALD-Pt, as reported in Table 3.1. This confirms that the IPE-DSG fabrication method can produce large usable capacitance without assembly. The standard deviation of capacitance was measured to be 12% across electrodes for sputtering, but increased to 23.2% for PE-ALD.

3.5 Gyroscope Operation

The device was operated in open-loop rate mode, Section 2.2.3. Electrostatic actuation and detection was implemented using electromechanical amplitude modulation (EAM) to mitigate feed-through parasitic signals [209]. A schematic of the gyroscope operation is provided in Fig. 3.14. A DC bias of 5V and AC carrier at 150kHz with 5V amplitude was applied to the inner shell through the central stem. AC drive with 40mV was applied to X-Forcer electrodes. Quadrature tuning and frequency split tuning was performed with 72V applied to Y-Forcer electrodes and 7V/12V to off-diagonal tuning electrodes ($Q+/Q-$




Figure 3.13: Measured capacitance of each of the 24 electrodes of sample A. The estimated coplanar capacitance has been removed from these measurements. Red illustrates capacitance with Cr-Au sputtering and Blue with PE-ALD of Pt. The dashed lines represent the average values of 102fF (Cr-Au) and 187fF (PE-ALD of Pt).

) reducing the post-trimming frequency split from 3.7 Hz to 200mHz. The device was mechanized in open-loop mode of operation using Zurich Instrument HF2LI lock-in amplifier. After carrier demodulation, AGC and PLL control loops were implemented on the drive axis to maintain drive amplitude and track the frequency of the drive signal, Fig. 3.14.

The scale factor was characterized using 2102 Ideal Aerosmith precision rate table. A sinusoidal rotational input of frequency 0.05Hz and amplitude of $3.1[^\circ/s]$ was applied, and the gyroscope response is illustrated in Fig. 3.15. For a Coriolis vibratory gyroscope in open-loop mode of operation, the scale factor can be expressed as $SF = SF_{mech}SF_{ele}$, where each is defined as [210, 9]:

$$SF_{mech} = \frac{2X_a\kappa\omega_x}{\sqrt{(\omega_y^2 - \omega_x^2)^2 + (2\mu_y\omega_x)^2}} \quad (3.3)$$

$$SF_{ele} = \frac{1}{2}R\frac{C_{sn}}{g}V_c(\omega_c + \omega_x) \quad (3.4)$$




Figure 3.14: Schematic of control loops implemented on the Zurich HF2LI Lock-in Amplifier. AGC and PLL were established along the drive axis (indicated as X-axis) for gyroscope operation.




Figure 3.15: Characterization of the scale factor of sample A by applying a sinusoidal excitation at a frequency of 0.05Hz and an amplitude of 3.1[°/s]. The resulting scale factor is 0.07 [mV/(°/s)].

where X_a is the amplitude of motion along the drive axis, κ is Bryan's factor, R is the feedback resistance, C_{sn} is the sense capacitance, V_c and ω_c are the carrier signal amplitude and frequency, ω_x and ω_y are the drive/sense axis frequencies, and μ_y is the damping coefficient for the sense axis. This expression scales linearly with respect to capacitance, but exponentially with respect to capacitive gap, as described in Eq. 3.1. Projecting the capacitance based on the thicknesses and trench widths discussed in Section 3.2, the scale factor can increase by a factor of 14, as illustrated in Fig. 3.16.




Figure 3.16: Projection of scale factor improvement based on optimized FLICE parameters from Section 3.2. The star indicates the measured scale factor of sample A using dimensions from Fig. 3.12. All values are normalized to sample A’s measured scale factor.

3.6 Discussion

The FLICE-based IPE-DSG fabrication method offers several advantages over conventional approaches. In terms of achievable sensing capacitance, IPEs attain capacitive gaps approaching 2 μm , comparing favorably with reported out-of-plane electrode gaps of 17 μm [204] and 5 μm [145]. While smaller capacitive gaps reduce the vibration amplitude at which nonlinear effects emerge, representing a tradeoff between sensitivity and linearity, the IPE-DSG can be designed to minimize the ratio of in-plane to out-of-plane motion for the $n=2$ wineglass mode [211], mitigating this constraint. Two additional benefits of the architecture are enhanced electrostatic tuning through the IPE configuration [212] and the ability to perform online trimming of asymmetry without a specialized setup [213] or temporary attachment of an electrode substrate [214, 215].

This chapter presents an initial demonstration of the IPE-DSG; further research is necessary to establish this as a scalable fabrication process. Key directions for future work include:

1. Optimizing ALD for conformal coating of high-aspect-ratio structures is critical. Thermal ALD of Pt [151] or aluminum-doped zinc oxide (AZO), which offers lower CTE and less nucleation delay on fused silica [216, 217], are promising candidates.

2. FLICE has been demonstrated at speeds up to 3 m/s [207]; however, maintaining etch selectivity at high speeds requires managing polarization consistency. Solid-state polarization control via electro-optic modulators (EOM) [131] or the use of point-defect modification, which achieves a uniform etch rate independent of laser polarization [218] are two viable approaches.
3. While aspect ratios $>150:1$ were demonstrated here, systematic optimization of FLICE parameters including repetition rate, z-step size, pulse energy, numerical aperture, and beamforming to mitigate refraction-induced distortion [219, 217] will be necessary to reliably achieve the highest possible aspect ratio capacitive gaps.

3.7 High-Capacitance Electrode Design

The single-cut electrode experimentally demonstrated represents only the simplest realization of in-plane patterning. This section speculates on what higher-capacitance designs could look like, presenting a configuration that maximizes capacitance but has not yet been experimentally tested.

The single-cut capacitance, given by Equation 3.1, serves as the baseline. For representative DSG parameters of $R = 3.15\text{ mm}$, $g = 5.5\text{ }\mu\text{m}$, and $T = 400\text{ }\mu\text{m}$, the total capacitance formed between the inner and outer shell is 12.7 pF. If the FLICE process is optimized to the $g = 2\text{ }\mu\text{m}$ capacitive gap demonstrated on flat samples, the total capacitance increases to 35.1 pF.

Drawing inspiration from parallel plate electrodes in planar MEMS [209], one implementation could include a combination of teeth-like-tines and concentric, pseudo-parallel plate electrodes, Figure 3.17. The use of teeth-like-tines is inspired by the design in [220], and would enable decoupling of mass and stiffness when performing tuning of frequency split [104]. Furthermore, the capacitance is greatly enhanced, as the number of parallel plates

is limited only by the width of the second cavity, which has already been experimentally increased to over 1 mm. Consider the pseudo-parallel plate in-plane electrodes design illustrated in Figure 3.17. This configuration comprises 28 tines, each with several parallel plate pairs patterned on both sides and each of angular width 3° . The capacitance of this configuration is calculated as follows:

$$C = \frac{2 \epsilon \alpha T N_{\text{tines}} \sum_{n=1}^{N_{pp}} (R_0 + nP)}{g} \quad (3.5)$$

where $N_{\text{tines}} = 28$ is the number of tines, α is the angular width of each parallel plate (3°), P is the pitch of each parallel plate pair in the radial direction, N_{pp} is the number of parallel plates in the radial direction, and R_0 is the inner shell radius (3.15 mm). The leading factor of two accounts for the parallel plates patterned on both sides of each tine, and the summation term accounts for the differing radii of the parallel plate pairs.

The radial pitch, P , is a function of several parameters. A minimum feature size of approximately $15 \mu\text{m}$ is assumed, which constrains the width of the tines, w_t . Capacitive gaps have reliably been fabricated with around $5.5 \mu\text{m}$ width, and gaps as small as $2 \mu\text{m}$ have been demonstrated on flat $300 \mu\text{m}$ thick fused silica substrates [203]. The spacing between parallel-plate electrodes is typically five times the capacitive gap of the electrodes to mitigate coupling between parallel plate pairs. In this configuration, the radial pitch of pseudo-parallel plate electrodes is therefore:

$$P = 2w_t + g + 5g = 2w_t + 6g \quad (3.6)$$

which leads to a total number of parallel plates in the radial direction of:




Figure 3.17: Comparison of electrode configurations for the DSG: out-of-plane electrodes, single-cut in-plane electrodes, and pseudo-parallel plate in-plane electrodes. The top row illustrates the design, the middle row the corresponding capacitance model, and the bottom row the calculated capacitance as a function of capacitive gap, shown for electrode thicknesses of 100 μm , 200 μm , 300 μm and 400 μm in the two in-plane cases.

$$N_{pp} = \left\lfloor \frac{W_2}{P} \right\rfloor - 2 \quad (3.7)$$

where W_2 is the radial width of the second cavity. While this radial width has been experimentally demonstrated up to 1 mm in width, we assume a value of 600 μm for this analysis.

Evaluating Equations 3.5–3.7 at the same electrode thickness of $T = 400 \mu\text{m}$ used above, a capacitive gap of $g = 5.5 \mu\text{m}$ yields a radial pitch of $P = 63 \mu\text{m}$ and $N_{pp} = 7$ parallel plates, for a total capacitance of 45.0 pF, an improvement of 3.5 times over the single-cut design. Optimizing the capacitive gap to $g = 2 \mu\text{m}$ yields $P = 42 \mu\text{m}$ and $N_{pp} = 12$, for a total capacitance of 213 pF, an improvement of 6.1 times. The dependence of capacitance on capacitive gap and electrode thickness for both configurations is illustrated in Figure 3.17.

This design has not been fabricated. Its demonstration would require metallization of a multi-plate array at aspect ratios beyond those achieved for the single-cut electrode, which remains a subject of future work.

3.8 Conclusion

This chapter presents an alternative approach to defining electrodes in fused silica resonators. While demonstrated on DSGs, this electrode patterning approach is applicable to other fused silica resonator architectures as FLICE can pattern electrodes in any preserved fused silica region. Experimental optimization of FLICE parameters demonstrated capacitive gaps as small as 2 μm on 300 μm thick fused silica samples with aspect ratios $\geq 150:1$ achieved through optimized laser pulse energy (0.540 μJ) and minimized etching duration (4 hours). This technique was applied to a DSG to define self-aligned, in-plane electrodes fabricated without assembly. Cross-sectional analysis of a fabricated IPE-DSG confirmed bulk capacitive gaps of 5.5–6 μm with a maximum height of 384 μm corresponding to aspect ratios up to 64:1.

The maximum as-fabricated Q-factor was 4.57 million, suggesting that FLICE patterning does not substantially increase surface energy loss. Metallization studies showed that 10 nm PE-ALD platinum increased total capacitance by 84% (to 4.68 pF) while reducing inner shell resistance by 59% compared to sputtered Cr-Au, both of which were attributed to the superior conformality of ALD in high-aspect-ratio structures. Feasibility of gyroscope operation was demonstrated in open-loop rate mode with a scale factor of 0.07 mV/(°/s). Projections based on optimized FLICE parameters suggest potential for 14× scale factor improvement through reduced capacitive gaps and increased electrode thickness. This work establishes IPE-DSG as a promising platform for next-generation micro-HRGs with reduced manufacturing complexity and potential for increased capacitance. Increased capacitance and the elimination of assembly-induced non-uniformity both act on the electrical noise floor of the sensor, and therefore on the drift of the yaw-axis gyroscope that bounds heading accuracy in a ZUPT-aided pedestrian INS. This remains an initial demonstration: scaling the process requires conformal metallization of the high-aspect-ratio gaps, and high-performance gyroscope operation has yet to be demonstrated with an IPE-DSG.

RQ1a: Capacitive electrodes are defined for DSG without an assembly step by stopping back-lapping before release and patterning the gaps directly into the bulk fused silica with FLICE. The minimum achievable gap and maximum aspect ratio are governed by the laser pulse energy and the KOH etching duration.

Chapter 4

Non-Metallized DSG

This chapter demonstrates operation of the DSG without a conductive film on the resonating element, to the best of our knowledge the first such demonstration for a microfabricated HRG, addressing the three functions that metallization conventionally provides: driving, tuning, and sensing¹. Of the mechanisms that dissipate energy in a fused silica shell resonator, discussed in Section 1.4.2, metallization of the resonating element is the only one that cannot be recovered by subsequent process optimization and is typically much more significant than the TED of the material. Preserving the pristine fused silica inner shell therefore requires replacing each of the three functions the conductive film supplies. Driving must establish a standing wave along a designated axis, since the shaker excitation used for initial characterization is sufficient only for WA operation. Tuning must correct the as-fabricated frequency split permanently, since electrostatic tuning is unavailable once the electrodes on the shell are removed. Sensing must resolve displacement at the wineglass antinodes without a capacitive gap. Driving the inner shell without metallization through indirect actuation is effective, but its efficiency depends strongly on the shell geometry through which

¹This chapter is based on the following peer-reviewed publication: A. R. Parrish, N. A. Strnad, Z. J. Klutts, L. Meira-Lopez, D. Wang, J. S. Pulskamp, and A. M. Shkel, “Non-Metallized Dual-Shell Hemispherical Resonator Gyroscope,” *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Salt Lake City, UT, USA, Apr. 2025 [221].

stress is transferred, a dependence quantified and optimized in Chapter 5. This chapter addresses **RQ1b**, how a DSG can be driven, sensed, and tuned without a conductive coating on the resonating element. Section 4.1 addresses the first, introducing IPA, in which a piezoelectric film deposited on the outer shell drives the inner shell through the junction between them, together with the electrode shaping that makes the excitation directional. Section 4.2 then describes the fabrication process, including deposition of the piezoelectric film and shadow-mask patterning of the outer electrodes. Section 4.3 presents the mode-shape characterization procedure and the laser ablation trimming it enables, including a teeth-like-tines approach that relocates ablation damage off the resonating shell to better preserve Q-factor. Section 4.4 describes the dual LDV setup developed for non-metallized detection. Section 4.5 then demonstrates open-loop rate gyroscope operation of the resulting device. Finally, the chapter is concluded with Section 4.6.

4.1 Non-Metallized Driving

The DSG comprises two concentric fused silica shells joined at a junction and sharing a common substrate, Section 2.2.5. The inner shell is the resonating element, while the outer shell was originally introduced to act as the second anchor of the double-ended architecture and as a protective casing for the resonator [3]. IPA exploits this arrangement by using the outer shell as the actuator, so that the piezoelectric film and its electrodes are confined to a structure that is not itself resonating and the inner shell is never coated. The approach was first presented by Wang et al. [163, 164] and Shkel et al. [165].

The film is deposited conformally over the three-dimensional outer shell [166] and sandwiched between a continuous bottom electrode and a patterned top electrode, Fig. 4.1. When a voltage is applied across the film, the converse piezoelectric effect generates an in-plane stress within it. Because the film is bonded to the outer shell, this stress is transferred into the shell, and because the two shells are continuous through the junction, it propagates from

the outer shell into the inner shell, exciting it into resonance. The junction is the region where the inner and outer shells meet and is the only mechanical path between them, so it carries the entire transduction. Driving the film at the resonant frequency of the $n = 2$ wineglass mode therefore sustains that mode in the inner shell without any electrode on the resonator. The film and electrode materials used in this chapter, and the process by which they are deposited, are described in Section 4.2.

Sustaining a mode is not sufficient for rate operation. In a CVG, angular rate is measured from the response of the standing wave established along a designated drive axis [110], so the excitation must be directional. A uniform electrode covering the outer shell would excite the shell without preference for orientation, in the manner of the shaker excitation used for initial characterization, which is adequate only for WA operation. Directionality is instead obtained by shaping the top electrode to the geometry of the resonator, following the electrode-shaping principles of Pulskamp et al. [5]. The modal stress of the $n = 2$ mode varies over the surface of the outer shell, and an electrode is placed only where the stress the film applies is compatible with it. Regions in which the two would oppose are excluded, so that the contributions of the remaining regions reinforce rather than cancel. This shaping analysis is developed in detail, and its dependence on shell geometry evaluated, in Chapter 5.

Because the resulting pattern is fixed relative to the mode shape rather than to the device outline, it must be aligned to the principal axes of elasticity of the individual resonator when the electrodes are deposited, Section 4.2. The two degenerate $n = 2$ modes of a trimmed resonator differ in Q-factor, and either may serve as the drive axis; here the lower Q-factor mode was selected and the electrodes aligned to it, following [164].




Figure 4.1: Schematic of the IPA-driven DSG. The outer shell is coated with a piezoelectric film and selectively excited using patterned electrodes. Stress applied to the outer shell is transferred to the vibrating element, the inner shell, through the junction, introducing a directional standing wave in the $n = 2$ wineglass mode. The film and electrode materials used in this chapter are given in Section 4.2; those used in Chapter 5 differ.

4.2 Fabrication Process

The fabrication process is illustrated in Fig. 4.2. Steps (a) through (e) follow the wafer-level glassblowing sequence described in Section 2.2.5, performed here on high-purity fused silica wafers of 500 μm (cap), 150 μm (device), and 2 mm (substrate) thickness, isotropically etched in 48 % wt. HF solution through a polysilicon hard mask. The hard mask is removed in KOH, the cap wafer is bonded to the device wafer and the resulting pair to the substrate wafer by plasma-assisted bonding, and the stack is optionally diced before glassblowing at 1700 $^{\circ}\text{C}$ in a custom-configured KSL-1700X-RTP MTI furnace with $\pm 2^{\circ}\text{C}$ uniformity over an 80 mm range. In contrast to prior DSG implementations, the pre-etched cavity geometry used here produces a structure with no central stem. The stem is not required for non-metallized operation, since the inner shell is never electrostatically biased.

The remaining steps depart from the conventional sequence in order to accommodate the piezoelectric film. In step (f), the backside of the substrate is flattened by lapping, which provides greater contact area for temperature control during the subsequent film deposition. The sample is then annealed at 920 $^{\circ}\text{C}$ for 30 h to relax the internal stress introduced by thermal shock during glassblowing [206].

In step (g), 20/100 nm Ti/Pt is sputtered onto the outer shell to serve as the bottom electrode, and 100 nm of $\text{PbHf}_x\text{Ti}_{1-x}\text{O}_3$ (PHT) is deposited by ALD in the SEMASC cleanroom at the DEVCOM Army Research Laboratory, following the process described in [166]. The films were deposited on a Kurt J. Lesker Company, Inc. ALD 150-LX platform at a nominal substrate temperature of 250 °C. An amorphous mixed-oxide film with polycrystalline PbO_x domains is first deposited using a super-cycle of the constituent binary oxide processes until the target thickness is reached, then crystallized into the perovskite phase by annealing at 525 °C for 30 min in an oxygen environment.

The inner shell is released in step (h) by continued lapping, concluding with fine polishing using cerium oxide slurry. The released resonator was cleaned in solvent only, using acetone followed by isopropyl alcohol, as the cleaning process compatible with a piezoelectric film had not yet been developed at the time of this study; that process is presented in Chapter 5. Laser ablation trimming of the as-fabricated frequency split is then performed, as described in Section 4.3, before the outer electrodes are deposited.

The electrode pattern is aligned to the principal axes of elasticity of the individual resonator. These axes are identified by the mode-shape characterization procedure of Section 4.3, Fig. 4.3(a). A custom shadow mask is then fabricated by mounting cleanroom tape on a 3D-printed holder and patterning it by laser ablation. The mask is aligned to the measured mode shape using an iDonus shadow mask aligner, which brings it into conformal contact with the three-dimensional outer shell, Fig. 4.3(b). In step (i), 25/75 nm Cr/Au is sputtered through the mask to form the patterned top electrodes, Fig. 4.3(c).

4.3 Non-Metallized Tuning

Trimming perturbs the stiffness and mass of the resonator at locations defined relative to the mode shape, so the principal axes of elasticity must be identified before any material




Figure 4.2: Fabrication process for the IPA-driven DSG: (a,b) cavity pre-etching of cap and substrate wafers, (c,d) triple-stack plasma-assisted wafer bonding and annealing, (e) glassblowing in a high-temperature furnace at 1700 °C to form the dual-shell structure, (f) backside flattening via lapping, (g) deposition of the Ti/Pt bottom electrode and the piezoelectric film, (h) continued lapping to fully release the inner shell, and (i) selective sputtering of the outer shell electrodes through a shadow mask. The film and electrode materials differ between this chapter and Chapter 5; those used here are given in the text. Laser ablation trimming, Section 4.3, is performed between steps (h) and (i).




Figure 4.3: Shadow masking method for deposition of the outer electrodes, step (i) of Fig. 4.2. (a) Characterization of the $n = 2$ wineglass mode principal axis orientation using an LDV with shaker excitation. (b) Application of the laser-ablated shadow mask, fabricated from cleanroom tape on a 3D-printed holder, aligned and applied conformally to the DSG outer shell. (c) Completed DSG with outer electrodes selectively deposited through the shadow mask.

is removed. Section 4.3.1 describes the characterization procedure developed for the DSG geometry, and Section 4.3.2 the laser ablation trimming it enables.

4.3.1 Mode-Shape Characterization

Identifying the principal axes of elasticity of a single shell by manual probing requires several hours, dominated by positioning the DSG with micrometers to points around the rim and re-tuning the PLL at each point. A semi-automated setup was therefore developed to reduce this time and to make the procedure repeatable across samples.

Setups reported in prior art for characterizing the Principal Axis of Elasticity (PAEs) of an HRG typically mount the resonator by its central stem on a rotary stage and probe the rim normal to the curved surface [213]. Two features of the DSG preclude this arrangement. The absence of a central stem, Section 4.2, requires that the device be mounted by its flat substrate, and the presence of the outer shell obstructs any optical path normal to the curved surface of the inner shell, requiring that the rim be probed from above. The setup adopted therefore mounts the DSG on a three-axis stage with the rim facing upward and probes it with a LDV, Fig. 4.4.




Figure 4.4: Conceptual illustration of the mode-shape characterization setup. The DSG is mounted flat by its substrate with the rim facing upward, driven into resonance by a bulk piezoelectric stack, and probed point by point along the rim by an LDV on a three-axis positioning stage, all inside a vacuum chamber.




Figure 4.5: Summary of direct-rim laser-ablation trimming iterations performed on the two piezoelectrically coated samples, sample 1 (top) and sample 2 (bottom).

A Micronix PPS-20 piezoelectrically driven three-axis stage, with precision on the order of tens of nanometers, provides X-Y-Z movement inside the vacuum chamber, and a PI Ceramic piezoelectric actuator drives the device into resonance. The DSG is mounted flat on carbon tape, which holds the device stationary while the stages move. The tape was selected deliberately to increase the anchor loss of the mount, reducing the observed Q-factor from millions to hundreds of thousands and correspondingly reducing the ringdown time and therefore the time required for the PLL to lock at each point. The rim is probed normal to its surface using a Polytec LDV, Fig. 4.6.




Figure 4.6: Experimental setup for semi-automated mode-shape characterization inside the vacuum chamber, realizing the arrangement of Fig. 4.4, with the DSG mounted on the three-axis stage.

The procedure used to identify the PAEs is as follows:

1. Focus the LDV at the center of the device and record the rim location at four points. The stage moves radially outward by the nominal rim radius to each estimated point, and the operator adjusts the position in small increments, using the LDV camera view and signal amplitude as feedback, until the probe is centered on the rim.
2. Fit a circle to the four recorded points to establish the rim geometry and plan the trajectory of the stage.

3. Move the stage to the next point on the fitted trajectory and record the amplitude and phase of the response.
4. Verify that the PLL remains locked, by testing for a 180° phase shift, and adjust if necessary.
5. Repeat for the remainder of the rim.

Measurement of the amplitude at each point was automated using the Zurich lock-in amplifier API to MATLAB. The PLL is first locked using the built-in PLL error estimation, the phase is then adjusted to maximize the amplitude of motion, and the maximum amplitude is recorded. The piezoelectric drive can optionally be stopped at each point and the resulting decay measured and filtered, so that the slope yields an estimate of Q without waiting for the full 63.2% decay required to measure the time constant directly.

Manual setup, comprising the four-point rim mapping and the initial PLL lock, requires 5 min to 10 min, after which acquisition proceeds without operator input. The duration of the automated acquisition is set by the ringdown time of the shell and is therefore reduced by the deliberately lossy mount, requiring a few minutes for a low Q -factor sample. Points are acquired at 6° spacing. Because the PAEs follow from the location of a single antinode, characterization over 90° of the rim is sufficient to orient the electrode pattern, while a full mode shape is obtained from 60 points spanning the complete rim. The mode shapes reported in Section 4.3.2, Fig. 4.7, were acquired with this setup, and the setup has since been used to characterize approximately 300 mode shapes. The full procedure is documented in Appendix B.4.

4.3.2 Direct-Rim Trimming

Of the trimming methods reviewed in Section 1.4.1, laser ablation was selected for this study, since it removes material quickly and precisely while preserving surface quality. The

principal axes of elasticity are first identified by characterizing the mode shape of the lower-frequency $n = 2$ wineglass mode, Section 4.3.1, using an LDV with shaker piezoelectric excitation, Fig. 4.7. Material is then removed by ablating holes directly into the rim of the DSG at the antinodes of that mode, raising its frequency toward that of the higher-frequency mode. The ablation was performed using a 1030 nm wavelength Light Conversion CARBIDE femtosecond laser source with a repetition rate of 60 kHz and a pulse energy of 29 μ J.

Two DSG samples were trimmed as part of this study, with characteristics summarized in Table 4.1. For both samples, the piezoelectric film was already deposited on the outer shell and remained in place throughout the direct-rim trimming. The as-fabricated frequency split of both resonators was in the range of 50 Hz to 60 Hz and was reduced to below 1 Hz in each case, the lowest being 0.39 Hz; the final frequency response of both resonators is shown in Fig. 4.8. The Q-factor declined over the course of trimming, which is attributed to material damage at the ablation sites, from 215k/810k to 124k/73k for sample 1 and from 80k/43k to 59k/49k for sample 2 along the two axes. The Q-factor and frequency split history of each sample over successive trimming iterations is shown in Fig. 4.5. Because ablation alters the mode shape, the principal axes are re-identified after trimming, and the outer electrodes are then aligned to the post-trim orientation and deposited, Section 4.2.

Table 4.1: Characteristics of the two piezoelectrically coated DSGs, trimmed directly on the rim.

Shell #	$(n = 2)$ Freq [kHz]	As-Fabricated		Post-Trimming	
		Δf [Hz]	Q-Factor X/Y Axes	Δf [Hz]	Q-Factor X/Y Axes
1	19.24	50.1	215K/810K	0.39	124K/73K
2	14.12	60.5	80K/43K	0.87	59K/49K

4.3.3 Teeth-Like-Tines Trimming

Direct-rim trimming corrects the frequency split but was also observed to significantly degrade Q-factor, since material is ablated into the resonating element itself, Section 4.3.2. To




Figure 4.7: Illustration of the direct-rim laser ablation trimming approach. The mode shape of the lower-frequency mode is obtained and laser ablation is performed at the antinodes.




Figure 4.8: Frequency response characterization of the two piezoelectrically coated DSG samples, trimmed directly on the rim, using shaker piezoelectric stack excitation and LDV for detection. The top/bottom figures show the frequency response of samples 1/2, respectively.

relocate this damage away from the vibrating shell, teeth-like-tines were defined on the rim by FLICE, following the release process detailed in Chapter 3. The rim was patterned with 32 tines with rounded corners to limit stress concentration, Fig. 4.9.




Figure 4.9: (Left) Cross-sectional illustration of the teeth-like-tines geometry. The nominal dimensions used in this dissertation are as follows: the thickness of the tines (T) was approximately $400\text{ }\mu\text{m}$, the angular width of the tine (α) was approximately 5.25° , and the radial width of the tine (W) was approximately $220\text{ }\mu\text{m}$. (Right) Illustration of tines undergoing ablation-based material removal; the trimming locations were aligned to the antinodes as detailed in Fig. 4.7.

During the $n = 2$ wineglass motion the rim itself is stressed while the tines are not, so ablating a tine perturbs the modal mass with negligible change in the stiffness of the resonating shell. This decoupling of mass and stiffness both preserves Q -factor and makes trimming more predictable [222, 104]. FLICE release of teeth-like-tines has similarly been reported to preserve Q -factors in excess of six million with minimal subsurface damage [223].

The teeth-like-tines approach also inverts the fabrication sequence used for the demonstrated device. In the process of Section 4.2, the piezoelectric film is deposited before the inner shell is released and trimmed, so the high-temperature anneal that recovers Q -factor after ablation cannot subsequently be applied. Depositing the film only after release instead allows each shell to be released, trimmed, and annealed before deposition. Owing to limited sample availability, the trimming and Q -factor study reported here was performed on bare shells rather than on piezoelectric-integrated devices, so the retention improvement is demonstrated on bare resonators and is expected to transfer through the post-release deposition sequence.

All shells in this study, with and without tines, were trimmed, cleaned, and annealed before the final Q-factor measurement, so that the presence of tines is the only variable between the two groups.

Across 17 shells, teeth-like-tines trimming reduced the frequency split from an average of 36.5 Hz to 1.0 Hz, with a lowest achieved split of 0.10 Hz, Fig. 4.10. For comparison, the seven bare shells trimmed directly on the rim were reduced from an average of 89.5 Hz to 3.6 Hz, with a lowest of 0.36 Hz. The shell reaching the record 0.10 Hz split was not the shell with the highest Q-factor retention, so the minimum split and the maximum retention were obtained on different devices.




Figure 4.10: (Left) Summary of trimming iterations across 17 teeth-like-tines DSGs, reducing the frequency split from an average of 36.5 Hz to 1.0 Hz. (Right) Frequency response of the DSG with the lowest achieved frequency split of 0.10 Hz.

The Q-factor retention of the two approaches is compared in Fig. 4.11. For the seven bare shells trimmed directly on the rim, the mean Q-factor fell from 2.17 million to 0.58 million, a reduction of 73 %. For the 17 shells trimmed with teeth-like-tines, the mean Q-factor fell from 1.82 million to 1.44 million, a reduction of 21 %. The substantially smaller reduction confirms that relocating ablation damage onto the tines preserves Q-factor, consistent with the mass-stiffness decoupling argument.




Figure 4.11: Mean quality factor before and after laser-ablation trimming, for shells trimmed directly on the rim and for shells trimmed with teeth-like-tines. Error bars denote one standard deviation. All shells were cleaned and annealed before the post-trimming measurement.

4.4 Non-Metallized Sensing

Detecting the response of the resonator without a capacitive gap requires that displacement be resolved directly at the wineglass antinodes. Two LDVs were used for this purpose, one probing the drive axis and one the sense axis, serving as a benchtop proxy for the integrated optical readout discussed in Section 4.6. Because gyroscope operation requires the drive and sense axes to be observed simultaneously while the device is rotated, both instruments must be mounted on the rate table itself and must hold their alignment throughout its motion.

A custom setup was developed that attaches and precisely positions two Polytec LDVs on a 2102 Ideal Aerosmith precision rate table, Fig. 4.12. In order to probe the ~ 6 mm diameter rim of the DSG at multiple locations, a 90° deflection unit with laser output 4.5 mm from the end of the tube was used on one LDV. The setup enables both positioning and tilt correction of both instruments, and uses 14.5 mm thick machined aluminum plates to ensure their stability during rate-table rotation, with counterbalances to uniformly distribute forces about the rate-table axis. Across both LDVs, the worst-case resolution was $0.124 \text{ pm}/\sqrt{\text{Hz}}$ at the 19.3 kHz operating frequency of the sample characterized in this chapter.

The resulting configuration is used for the drive-axis frequency response reported in Section 4.5 and for the open-loop gyroscope demonstration presented there, in which one LDV provides the feedback signal for the drive-axis control loops and the other is demodulated for the gyroscope response, Fig. 4.13.




Figure 4.12: Experimental setup for optical probing of the DSG sample. Both the vertical and horizontal LDVs were precisely positioned and corrected for tilt to maximize the reflected signal from the sample.

4.5 Gyroscope Operation

Gyroscope operation was demonstrated on sample 1 of Table 4.1, which after trimming exhibited a frequency split of 0.39 Hz at an $n = 2$ resonant frequency of 19.24 kHz. Mounting the sample introduced a further asymmetry: the tape used to secure the DSG to the PCB increased the trimmed frequency split from 0.39 Hz to 2 Hz, indicating that the packaging of




Figure 4.13: Schematic of the IPA-driven DSG in open-loop gyroscope configuration. One LDV is aligned to the drive axis, which is used to establish the AGC and PLL control loops. The other LDV is aligned to the sense axis and demodulated for the gyroscope response.

a trimmed resonator is itself a source of frequency split and warrants attention in a fielded implementation.

Testing was performed using an FPGA-based Zurich Instruments HF2LI lock-in amplifier. Electromechanical amplitude modulation, required in Chapter 3 to suppress feedthrough in the electrostatic device, was neither possible nor necessary here, since no capacitive path exists between drive and sense in a non-metallized implementation. IPA was used to selectively drive the $n = 2$ wineglass mode, with 240 mV AC applied across the PHT film at 19.3 kHz. The resulting frequency response, probed by the drive-axis LDV, is shown in Fig. 4.14. The intended drive axis was excited 25.6 dB above the noise floor and exceeded the response observed on the sense axis by 9.9 dB, confirming that the shaped electrodes of Section 4.1 excite the intended mode preferentially rather than driving the resonator without directional preference.

The open-loop configuration is summarized in Fig. 4.13. AGC and a PLL were implemented using feedback from the drive-axis LDV, maintaining the amplitude of the standing wave and tracking the resonant frequency. The sense LDV probed the rim 135° from the drive LDV, aligned to an antinode of the sense axis. A sinusoidal angular stimulus of $15.7^\circ \text{ s}^{-1}$

magnitude and 4 s period was applied with the rate table, and the demodulated sense-axis response is shown in Fig. 4.15. The observed scale factor was $\sim 0.02 \text{ mV}/(^{\circ}/\text{s})$, corresponding to $\sim 1.2 \text{ nm}/(^{\circ}/\text{s})$ of displacement at the sense-axis antinodes, since the LDV reports displacement directly.




Figure 4.14: Frequency response probed by the drive-axis LDV for the IPA-driven DSG. The intended drive axis is observed to have $\approx 25.6 \text{ dB}$ SNR and is 9.9 dB larger than the observed sense-axis displacement.




Figure 4.15: Open-loop gyroscope response of the PHT-coated DSG to sinusoidal input rotation with a period of 4 s and an amplitude of $15.7^{\circ} \text{ s}^{-1}$.

4.6 Conclusion

In this chapter, a microfabricated HRG was operated as a rate gyroscope with no conductive film on the resonating element, to the best of our knowledge for the first time, replacing each of the three functions that metallization conventionally provides. Driving was achieved by IPA, in which 100 nm of atomic layer deposited PHT on the outer shell, patterned with

electrodes shaped to the modal stress of the $n = 2$ wineglass mode, transfers stress through the junction to excite the inner shell; the intended drive axis was excited 25.6 dB above the noise floor and 9.9 dB above the sense axis, confirming directional excitation. Tuning was achieved by laser ablation at the antinodes of the lower-frequency mode, following a semi-automated mode-shape characterization procedure developed for the stemless dual-shell geometry, and reduced the as-fabricated frequency split from 50 Hz to 60 Hz to below 1 Hz on two piezoelectric coated samples, the lowest being 0.39 Hz. A teeth-like-tines variant of the trimming, introduced to preserve Q-factor and decouple mass and stiffness, was demonstrated on bare resonators to reach a frequency split as low as 0.10 Hz while Q-factor was reduced on average by 21 % (n=17), compared with 73 % for direct-rim trimming (n=7). Sensing was achieved with a dual LDV setup mounted on a precision rate table, probing the drive and sense antinodes at a worst-case resolution of $0.124 \text{ pm}/\sqrt{\text{Hz}}$. Open-loop rate operation was demonstrated with a scale factor of $\sim 0.02 \text{ mV}/(^{\circ}/\text{s})$, corresponding to $\sim 1.2 \text{ nm}/(^{\circ}/\text{s})$ at the sense-axis antinodes.

Demonstrating all three functions without metallization removes the one energy loss mechanism in a fused silica shell resonator that no subsequent process optimization recovers, and therefore removes the constraint that has bounded the achievable Q-factor of microfabricated HRGs below the limit set by the structural material. That limit governs the noise floor of the sensor, and through it the drift of the yaw-axis gyroscope on which heading accuracy in a ZUPT-aided pedestrian INS depends. Two elements of the demonstration remain to be carried further. The benchtop vibrometers stand in for an interferometric readout integrated alongside the resonator, toward which our group has reported progress[158, 224, 159], and the trimmed frequency split increased from 0.39 Hz to 2 Hz on mounting, so packaging of a trimmed resonator is itself a source of asymmetry once electrostatic tuning is unavailable.

RQ1b: Rate operation of a DSG with no conductive film on the resonating element is achieved by replacing each function of metallization independently: IPA of the outer shell

for directional driving, dual-LDV optical detection for sensing, and laser ablation trimming for permanent correction of the frequency split.

Chapter 5

Analysis of DSG Geometry for Indirect Piezoelectric Actuation

Chapter 4 demonstrated non-metallized driving but left its efficiency dependent on shell geometry; this chapter examines how the shell geometry of the DSG governs the efficiency with which IPA, introduced in Section 4.1, transfers stress from the outer shell into the inner shell¹. Greater transduction efficiency produces a larger displacement amplitude at the antinodes for a given drive voltage, and the noise floor of a rate gyroscope scales inversely with that amplitude, Eq. 2.32. The mechanism governing this transfer between the shells, and the design parameters that control it, have not previously been investigated. The geometry cannot be optimized for transduction alone: the shell profile that sets the stress path through the junction also sets how acoustic energy radiates into the substrate, and the dimensionless γ ratio previously identified as governing anchor loss in dual-shell resonators [100, 6] is also the ratio expected to govern transduction efficiency. Because resonant displacement depends on

¹This chapter is based on the following publication under review (revise and resubmit): A. R. Parrish, N. A. Strnad, J. S. Pulskamp, and A. M. Shkel, “Analysis of Dual Shell Hemispherical Resonator Gyroscope Geometry for Indirect Piezoelectric Actuation,” *IEEE/ASME Journal of Microelectromechanical Systems*, Accepted, 2026.

the product of Q-factor and transduction efficiency, the two must be evaluated together. This chapter addresses **RQ1c**, which dual-shell geometry maximizes the displacement amplitude achievable by IPA, given that the geometric parameter governing transduction efficiency also governs anchor loss. Section 5.1 defines the γ ratio and advances the hypothesis that it is the dominant geometric parameter governing IPA efficiency. Section 5.2 describes the modeling approach, including generation of candidate designs from a glassblowing model, identification of optimal electrode regions on the outer shell, and adaptation of the RAFT technique to the three-dimensional geometry of the DSG. Section 5.3 applies this modeling to the influence of geometry on transduction efficiency. Section 5.4 describes fabrication of the samples evaluated, including the cleaning process developed for this study, and Section 5.6 reports their experimental characterization. Section 5.7 discusses the results, and finally the chapter is concluded with Section 5.8.

5.1 Transduction Efficiency and the γ Ratio

In IPA, stress generated in the piezoelectric film is transferred into the outer shell and propagates through the junction into the inner shell, Section 4.1. The junction is the intersection region between the two shells, and as the only mechanical path between them it carries the entire transduction. Its position relative to the resonating element is therefore expected to govern the efficiency with which stress is transferred, which motivates describing candidate geometries by a parameter that captures that position.

The γ ratio [100] is defined as $\gamma = D/L$, where D is the distance from the stem to the junction and L is the inner shell rim length, Fig. 5.1. This ratio was previously identified as critical for anchor loss in dual-shell resonators [100, 6], and is applied here to analyze IPA transduction efficiency.




Figure 5.1: Overview of DSG geometry (cross-sectional view of right half of the device as shown in Fig. 4.1). The dimensionless γ ratio used to characterize DSG designs in this study is defined as $\gamma = D/L$, where D is the distance from the base of the central stem (or the equivalent location for stemless designs) to the junction between the inner and outer shells and L is the length of the inner shell.

It is hypothesized that increasing the γ ratio will increase transduction efficiency. This is motivated in part by Guan et al. [100], who showed that anchor loss is more significant through the flat substrate than through the central stem at larger values of γ ratio, potentially indicating larger energy coupling from the inner shell to the outer shell and through the flat substrate. This hypothesis is first explored through theoretical modeling, Sections 5.2 and 5.3, then evaluated experimentally on five PZT-coated DSGs spanning a range of γ ratio, Section 5.6.

5.2 Modeling Approach

To test the hypothesis that the γ ratio governs IPA transduction efficiency, a modeling procedure was developed to evaluate IPA performance across a range of DSG geometries. Ten DSG designs with varying γ ratio were generated, the optimal electrode placement on the outer shell was determined using the electrode shaping method of Pulskamp et al. [5], and transduction efficiency was evaluated by adapting the RAFT technique [225] to treat the DSG as a lumped mass-spring-damper system.

5.2.1 DSG Designs Considered

Ten DSG designs were evaluated for IPA transduction. The geometries were generated using the glassblowing model developed by Asadian et al. [3], implemented in COMSOL Multiphysics as a 2D axisymmetric isothermal fluid flow model. Fused silica was modeled as a generalized Newtonian incompressible fluid with temperature-dependent viscosity, and cavity pressures were calculated at each time step according to the ideal gas law until equilibrium of the glassblowing is reached. For fixed wafer thickness and glassblowing temperature, the final DSG profile was determined solely by the geometry and depth of the pre-etched cavities, as illustrated in Fig. 4.2(a–e). By varying the cavity dimensions, designs spanning low to high γ ratios were generated, Fig. 5.2.

As described in Section 4.2, the central stem is not required once the resonating element is no longer electrostatically biased. Here, its presence is treated as a design variable rather than removed outright, since stem removal was hypothesized to reduce anchor loss. Accordingly, two groups of five designs were evaluated: stemmed designs (1s–5s) and stemless designs (1n–5n), each indexed 1–5 by progressively increasing γ ratio, Fig. 5.2.




Figure 5.2: Ten DSG designs considered for IPA transduction analysis, organized into two groups of five by increasing γ ratio: stemmed designs (1s–5s) and stemless designs (1n–5n). Designs 2s/2n and 4s/4n were fabricated with PZT deposited on the outer shell and experimentally characterized.

5.2.2 Electrode Shaping

To evaluate the various DSG geometries, optimal regions for electrode placement need to be identified; this problem is referred to as electrode shaping. The electrodes need to be




Figure 5.3: Two adjacent differential volume elements taken from (right) DSG outer surface and a piezoelectric layer: (top left) the excited piezoelectric material, (bottom left) DSG outer shell experiencing deformation dictated by a given mode shape. The traction forces illustrated correspond to a state of biaxial stress in the piezoelectric element and plane stress in the modal element. The coordinate frames from global X-Y-Z frame to local piezoelectric poling frame (1-2-3) are illustrated. The local 3-direction is normal to the outer surface at each point. 1- and 2-directions are defined to align with the invariant principal strains at each point. Adapted from [5].

placed in regions where the piezoelectric stress is compatible with the modal stress, the principle applied without derivation in Section 4.1. The method used in this study is based on Pulskamp et al. [5], but extended to account for the 3-D geometry of the DSG. First, the $n = 2$ wineglass mode shape of the DSG was computed via eigenanalysis in COMSOL Multiphysics. The phase of the mode shape was selected to maximize displacement and the piezoelectric material stress (σ_p) was arbitrarily defined to be tensile in the following analysis ($\sigma_p > 0$). This analysis assumed small piezoelectricity and the effects of piezoelectricity on modal results were neglected. This is justified by the large thickness ratio of fused silica to PZT, approximately 1000:1 for a typical 500 μm outer shell.

To evaluate the compatibility of the local stress, it is necessary to rotate the stress tensor from the global frame (X-Y-Z) to the local piezoelectric poling frame (1-2-3), Fig. 5.3. In the local poling frame, it is assumed that the 3-direction is perpendicular to the outer-shell surface, and 1/2 directions are in-plane aligned to the local invariant principal stresses. The stress tensor was transformed into the piezoelectric poling frame using a two-step process. First, a

vector normal to the outer shell surface at each mesh element was exported from COMSOL in the global coordinate frame. An intermediate coordinate system was constructed at each mesh element with the 3-axis aligned with the outward surface normal and two orthogonal in-plane directions arbitrarily selected to complete the orthonormal basis. The stress tensor was rotated into the surface-normal coordinate system using the standard second-order tensor transformation [226]. The second transformation, following Pulskamp et al. [5], rotated the in-plane coordinate system to align with the principal stress directions. This eliminates the in-plane shear stress τ_{xy} and yields the two in-plane principal stresses σ_1 and σ_2 , which can be computed from the post-transformation normal stresses σ_x , σ_y , and shear stress τ_{xy} as:

$$\sigma_{1,2} = \frac{\sigma_x + \sigma_y}{2} \pm \sqrt{\left(\frac{\sigma_x - \sigma_y}{2}\right)^2 + \tau_{xy}^2} \quad (5.1)$$

This study employs PZT as the piezoelectric material. Therefore, the piezoelectric coefficients were assumed to satisfy $e_{31} = e_{32}$ and $\text{sign}(e_{31}) = -\text{sign}(e_{33})$ [227]. A parameter, β , was defined to identify regions where stresses do not cancel out [5]:

$$\begin{cases} \beta = \text{sign}(\sigma_1), & \text{if } \text{sign}(\sigma_1) = \text{sign}(\sigma_2) \\ \beta = 0, & \text{if } \text{sign}(\sigma_1) \neq \text{sign}(\sigma_2) \end{cases} \quad (5.2)$$

where σ_1 and σ_2 are the in-plane principal stresses. With this definition, $\beta \in \{-1, 0, +1\}$. For $\beta = 0$ there is cancellation between the principal stresses; however, for $\beta = \pm 1$ there is no cancellation and electrodes are placed in these locations, illustrated in Fig. 5.4(a). The stress profile $\sigma_1 + \sigma_2 - \sigma_3$, where σ_3 is the out-of-plane stress component perpendicular to the shell surface, was also used to guide electrode placement, capturing regions where partial stress cancellation occurs but net positive contributions to piezoelectric transduction remain, Fig. 5.4(b). The final electrode placement for the DSG is illustrated in Fig. 5.4(c).




Figure 5.4: (a) Distribution of β values on the outer shell of DSG 4n for the $n = 2$ wineglass mode. Red regions ($\beta = +1$) and blue regions ($\beta = -1$) indicate locations where electrodes are driven with opposite polarities. Green regions ($\beta = 0$) indicate regions where some in-plane stress cancellation occurs. (b) Illustration of $\sigma_1 + \sigma_2 - \sigma_3$. This metric identifies regions where, despite partial cancellation indicated by $\beta = 0$, significant modal force contribution remains. (c) Final electrode placement for IPA of the $n = 2$ wineglass mode. Red and blue regions denote electrodes requiring opposite voltage polarity to generate compatible stress states for the $n = 2$ wineglass mode.

5.2.3 Rapid Analytical Finite Element Analysis Technique

A modified version of RAFT developed by Puder et al. [225] was employed to evaluate IPA of DSG geometries. RAFT was selected for its ability to predict resonant displacement amplitude directly from FEA eigenanalysis results, without requiring computationally intensive harmonic analysis, enabling rapid evaluation of all ten DSG designs. Relative to the full RAFT implementation, the analysis was simplified in two ways: the parallel plate assumption eliminates the need for a separate electrostatic simulation, and only the resonant displacement amplitude is evaluated rather than the full frequency response. The analysis treats the $n=2$ wineglass mode as a lumped mass-spring-damper system, with the goal of extracting the modal stiffness (k_m) and modal force (F_m) to predict displacement amplitude at resonance. Other modes such as tilt or the $n=3$ wineglass mode were neglected and assumed to be sufficiently separated in frequency for the Q-factors considered.

Eigenanalysis was performed in COMSOL Multiphysics. The same coordinate transformations described in Section 5.2.2 were performed to align strains normal to the outer shell

surface and in-plane invariant strains. From this simulation, the modal strain (S_n), total mechanical energy (U_m), and maximum displacement (δ) were obtained. U_m was extracted from the model as follows:

$$U_m = \frac{1}{2} \delta^2 \int_V S_n^T c^E S_n dV, \quad (5.3)$$

where S_n^T denotes the transpose of S_n , c^E is the stiffness matrix at constant electric field, and V is the volume of the DSG. The modal stiffness is obtained from Hooke's law by equating the total mechanical energy of the system to the potential energy of an equivalent lumped spring, $U_m = \frac{1}{2} k_m \delta^2$, from which k_m is directly extracted [225]. The modal force was calculated as:

$$F_m = \Phi \int_{V_p} e^T \nabla \phi S_n dV, \quad (5.4)$$

where Φ is the applied voltage magnitude and $\nabla \phi$ is the unity-normalized electric potential field, such that $E_{app} = \Phi \nabla \phi$ is the real applied electric field, and V_p denotes the volume of only the piezoelectric layer. The modal force is a lumped scalar quantity that characterizes the effectiveness with which an applied electric field drives a specific vibrational mode through the piezoelectric effect. Finally, the displacement amplitude (u_{res}) evaluated for IPA is:

$$u_{res} = \left(\frac{F_m}{k_m} \right) \frac{e^{j\omega t}}{\left(1 - \frac{\omega^2}{\omega_n^2} \right) + \frac{j\omega}{Q_m \omega_n}}, \quad (5.5)$$

where ω is the excitation frequency, ω_n is the natural frequency of the mode, and Q_m is the assumed mechanical Q-factor. At resonance ($\omega = \omega_n$), Eq. 5.5 simplifies to $|u_{res}| = Q_m \frac{F_m}{k_m}$.

To simplify the electric field calculation, the outer shell surface was assumed to be locally flat, approximating each finite element as planar and enabling the piezoelectric film and electrode stack to be treated as a parallel plate. Under this assumption, the piezoelectric poling direction (3-direction) is always normal to the outer shell surface for the conformally deposited film. Consequently, the electric field magnitude is determined purely by the applied voltage and film thickness, eliminating the need for finite element analysis of the electric

field. Another consequence of this assumption is that the modal force depends only on the compatibility between the stress profile of the mode shape, the piezoelectric material coefficients, and electrode placement. With this assumption, the differential modal force can be expressed as:

$$\begin{aligned} dF &= \Phi e^T \nabla \phi S_n dV = e^T E_{app} S_n dV \\ &= E_{app} (e_{31} \epsilon_{xx} + e_{32} \epsilon_{yy} + e_{33} \epsilon_{zz}) dV, \end{aligned} \tag{5.6}$$

where E_{app} is the applied electric field magnitude, and ϵ_{xx} , ϵ_{yy} , ϵ_{zz} are the strain components from the normalized modal strain field S_n .

To quantify the error introduced by the locally flat assumption, a 2D-axisymmetric electrostatic simulation of the 500 nm PZT film between conformal electrode boundaries was performed on the cross-sectional profile of DSG 4n, resolving the field on the curved geometry. Across the E_1 electrode footprint, where the approximate radius of curvature is 2.4 mm, approximately 4800 times the film thickness, the simulated electric field in the poling direction differed from the parallel-plate value Φ/t_{film} by less than 0.11%. Propagating this difference through the modal force integral of Eq. 5.6 changed the predicted resonant displacement amplitude by 0.01%, and the locally flat assumption is therefore justified for the DSG geometries considered in this work.

5.2.4 Piezoelectric Analysis

This analysis adopts PZT thin-film properties from Ledermann et al. [228], who characterized {100}-textured PZT 53/47 thin films near the MPB, a composition closely matched to the films used in this work, described in Section 5.4. The $e_{33}:|e_{31}| = 0.83$ ratio from Ledermann et al. [228] was adopted and scaled by the $d_{33,f}^{\text{meas}}$ measurement described in Section 5.5, yielding $e_{33} = 11.47 \text{ C m}^{-2}$ and $e_{31} = e_{32} = -13.84 \text{ C m}^{-2}$.




Figure 5.5: Illustration of the contributions to the modal force from the e_{31} , e_{32} , and e_{33} components for DSG design 4n. The legend applies to all three component plots.

The differential modal force contributions from each piezoelectric coefficient are defined as $dF_{e_{31}} = E_{app}e_{31}\epsilon_{xx}$, $dF_{e_{32}} = E_{app}e_{32}\epsilon_{yy}$, and $dF_{e_{33}} = E_{app}e_{33}\epsilon_{zz}$. These contributions are evaluated over the surface of DSG 4n and illustrated in Fig. 5.5. The spatial distribution reveals that the most significant contributions to modal force occur directly above the junction region, where stress propagates from the outer shell to the inner shell.

Across all ten DSG designs, the integrated modal force contributions on average satisfy $F_{e_{31}} + F_{e_{32}} \approx 6.63F_{e_{33}}$ with a standard deviation of 18%, indicating that the modal force is dominated by transverse piezoelectric coupling terms. Because many piezoelectric materials with 4mm or 6mm crystal symmetry satisfy $e_{31} = e_{32}$ [227], the electrode shaping methodology applies directly to other materials such as PHT, BaTiO₃, AlN, and ZnO. Since the modal strain field and modal stiffness are determined by the DSG geometry alone, changing the piezoelectric material scales the modal force through its coupling with the modal strain, so the γ ratio at which F_m/k_m is optimized remains set by the DSG geometry.

5.3 Theoretical Analysis

The analysis of the ten different DSG geometries is presented to evaluate the hypothesis that increasing γ ratio improves IPA transduction efficiency. Two distinct electrode groups were identified, termed E_1 (closer to the axis of symmetry) and E_2 (further away). It

was observed that electrode group E_1 's contribution to the modal force is on average 9.5 times greater across the ten designs, Fig. 5.6. The F_{E1}/F_{E2} ratio increases monotonically with γ ratio within both families, from 0.8 at design 1s to 17.7 at 5n, exceeding 3.9 for all fabricated designs. E_2 electrodes were not patterned on the fabricated devices, and all subsequent analysis and reported modal forces therefore correspond exclusively to E_1 . Since E_1 occupies regions of high modal strain, increasing its surface area while maintaining stress compatibility was anticipated to increase modal force. Cross-sections of the first principal stress invariant, Fig. 5.7 (left), reveal that the transition from tensile to compressive stress on the outer shell is governed by the junction location, which is set by the γ ratio. Consequently, the outer radius of E_1 scales directly with γ ratio, Fig. 5.7, providing a mechanism by which increasing γ ratio increases modal force.

Consistent with the hypothesis, modal force increases strongly with γ ratio, Fig. 5.8; comparing designs 1s and 5s, modal force increases by $2,224\times$. Modal stiffness, however, also increases with γ ratio, by a factor of 14.3 for the same designs, Fig. 5.8. At higher γ values, the junction radius becomes larger, reducing the effective length of the inner shell and thereby increasing its modal stiffness for the $n = 2$ wineglass mode. Since the predicted resonant displacement is $|u_{res}| = Q_m(F_m/k_m)$, maximizing displacement requires optimizing the ratio F_m/k_m rather than F_m alone, Fig. 5.9. The predicted displacement is maximized at design 3n, corresponding to $\gamma = 0.96$, which represents an optimum balancing the competing effects of increasing modal force and increasing modal stiffness. These results partially confirm the hypothesis: increasing γ ratio does improve IPA transduction efficiency, but only up to an optimal value beyond which rising modal stiffness dominates.

5.4 Fabrication

To experimentally evaluate the modeling prediction that IPA transduction efficiency is maximized at an intermediate γ ratio, four DSGs spanning a range of γ were fabricated: designs




Figure 5.6: (Left) Electrode group E_1 is located closest to the axis of symmetry; E_2 is located further in the radial direction. (Right) The modal force of E_1 and E_2 for the ten DSGs considered. Assumes an applied electric field of $E = 200 \text{ kV cm}^{-1}$.




Figure 5.7: (Left) Cross-section view of the first principal invariant of stress for a low γ ratio design (top, DSG 1n) and a high γ ratio design (bottom, DSG 5n). (Right) The relationship between the γ ratio and r_{E1} , defined as the outer radius of electrode group E_1 , is shown for all DSG designs considered.

2s, 2n, 4s, and 4n. Of these, design 4n ($\gamma = 1.09$) is predicted to achieve the largest displacement, representing 97% of the maximum predicted at design 3n ($\gamma = 0.96$).

The process follows the sequence described in Section 4.2 through step (f) of Fig. 4.2, comprising cavity pre-etching, triple-stack plasma-assisted bonding, glassblowing, backside flattening, and the 920°C stress-relief anneal. The steps that follow differ from Chapter 4 in the piezoelectric film deposited and in the metal used for the outer electrodes.

In step (g), 500 nm of PZT was sputtered onto the outer shell of the DSG. In preparation for the PZT, a 20 nm Ti interlayer was first deposited to promote adhesion to the sub-




Figure 5.8: Modal force (blue) and modal stiffness (orange) versus design number for stemmed (1s–5s) and stemless (1n–5n) DSG geometries. Modal force increases by a factor of $2,224\times$ comparing stemmed designs 1s to 5s, while modal stiffness increases by a factor of 14.3. Modal force values assume an applied electric field of $E = 200 \text{ kV cm}^{-1}$.

strate, followed by the 100 nm Pt bottom electrode using pulsed DC magnetron sputtering, as described in [166]. Thin film PZT with a nominal composition near the MPB was then deposited by RF magnetron sputtering from a single PZT ceramic target. Both the Pt bottom electrode and the PZT thin films were deposited using a commercial Evatec (Trübbach, Switzerland) Clusterline sputter deposition platform at a nominal substrate temperature of 620°C .

The inner shell was released in step (h) by mechanical lapping. Before the outer electrodes were deposited, the sample was cleaned by the process developed for PZT-coated DSGs and described in Section 5.5. In step (i), 50/100 nm Ti/Pt outer electrodes were selectively sputtered onto the outer shell using the electrode geometry determined in Section 5.2.2. As described in Section 4.2, the orientation of the $n = 2$ wineglass mode is not deterministic, since the PAEs vary between devices as a result of fabrication imperfections, so electrode alignment was performed individually for each DSG using the mode-shape characterization procedure of Section 4.3.1 and a laser-ablated shadow mask. Conformal contact between the mask and the curved outer shell was essential to prevent sputtered material from depositing beneath the mask edges and creating electrode shorts.




Figure 5.9: Predicted displacement amplitude at the rim antinode for the $n = 2$ wineglass mode versus design number for stemmed (blue, 1s–5s) and stemless (orange, 1n–5n) DSG geometries. Displacement is maximized for intermediate γ ratios where the F_m/k_m ratio is optimized. Assumes $Q_m = 100\,000$ and an applied electric field of $E = 200\text{ kV cm}^{-1}$.

5.5 Cleaning of PZT-Coated Dual-Shell Gyroscopes

Following release of the inner shell, the DSG must be cleaned to remove residue from lapping. Conventional fused silica cleaning methods, such as RCA-1 or piranha solutions, are incompatible with PZT, while solvent-only cleaning with acetone and isopropyl alcohol was found insufficient for achieving high-Q-factor operation. A three-step cleaning process was therefore developed. First, ultrasonic cleaning was performed using a Chem-Crest P360D/HF-132 system at 70°C for 20 min in deionized water with added Chem-Crest 103 cleaner. Second, AZ 4620 photoresist was spun over the PZT-coated outer shell to protect the piezoelectric film, and O_2 plasma cleaning was performed using a PE-100 Benchtop Plasma System at 150 W power, 20 sccm O_2 flow, and 100 mTorr chamber pressure for 10 min. Third, to remove adsorbed water [123], the sample was baked at 350°C for 12 h in vacuum below 1 μTorr .

The effectiveness of this cleaning process was evaluated through Q-factor measurements on five PZT-coated DSGs, Table 5.1. Although this study focused on IPA of the $n = 2$ wineglass mode, Q-factors for both the $n = 2$ and $n = 3$ modes are reported to provide broader evidence of cleaning effectiveness. The Q-factor was determined from ringdown measurements as $Q = \pi f \tau$, where f is the resonant frequency and τ is the measured time

Table 5.1: Characterized Q-factor of PZT-coated DSG samples before and after cleaning, for the $n = 2$ and $n = 3$ wineglass modes.

Sample	Q (thousands)			
	$n = 2$		$n = 3$	
	Pre-Cleaning	Post-Cleaning	Pre-Cleaning	Post-Cleaning
2s	320	317	1490	1686
2n	21	1059	141	1372
4s	20	256	230	430
4s	82	421	60	425
4n	41	278	116	503

constant. Mean Q-factors increased from 97k to 466k for the $n = 2$ mode and from 407k to 883k for the $n = 3$ mode after cleaning, representing $4.8\times$ and $2.2\times$ improvements.

Film quality was further assessed through two independent measurements: leakage current density and the effective piezoelectric coefficient $d_{33,f}^{\text{meas}}$ measured by LDV. Leakage current density was measured at an applied field of $\pm 240 \text{ kV cm}^{-1}$ on the curved outer shell electrode of DSG 2n, where the cleaned sample exhibited $0.03 \mu\text{A cm}^{-2}$ compared to $0.38 \mu\text{A cm}^{-2}$ for the uncleaned flat reference. Separately, $d_{33,f}^{\text{meas}}$ was measured on the flat region of the cleaned DSG and on the uncleaned flat reference sample. Measurements were performed at frequencies away from mechanical resonance, 24 kHz and 10 kHz for the uncleaned and cleaned samples respectively, to ensure adequate SNR, yielding $d_{33,f}^{\text{meas}} \approx 235 \text{ pm V}^{-1}$ for the uncleaned reference and $d_{33,f}^{\text{meas}} \approx 98 \text{ pm V}^{-1}$ for the cleaned sample. The notation $d_{33,f}^{\text{meas}}$ is used rather than $d_{33,f}$, the standard notation for the effective longitudinal piezoelectric coefficient of a clamped piezoelectric thin film, because the electrode diameter of approximately 1 mm is large and may produce non-negligible substrate bending due to $e_{31,f}$ [229]. While the cleaning process reduced both the leakage current density and the measured piezoelectric coefficient, no recovery anneal was attempted, for example 600°C for 5 min in O_2 [166], which may recover some of the lost electromechanical response. These results indicate that the cleaning process represents a tradeoff between DSG Q-factor and PZT electromechanical response, with the cleaned film retaining functional piezoelectric properties.

5.6 Experimental Evaluation

To experimentally test the hypothesis that IPA transduction efficiency is maximized at an intermediate γ ratio, the fabricated DSG samples were characterized. Since the resonant amplitude depends on both Q_m and the transduction efficiency F_m/k_m as $|u_{res}| = Q_m(F_m/k_m)$, these effects were evaluated independently. First, Q-factor measurements established how anchor loss varies with γ ratio, revealing a competing effect not captured by transduction efficiency alone. Second, IPA displacement measurements tested whether the predicted dependence on γ ratio governs transduction efficiency, with Q-factor normalization used to isolate the two effects.

5.6.1 Quality Factor Evaluation

Prior to IPA characterization, the Q-factor of the fabricated DSG samples was measured to quantify the influence of γ ratio on anchor loss. The Q-factor was characterized using a piezoelectric shaker for excitation and an LDV for detection. The measured Q-factors are summarized in Fig. 5.10 and include 28 samples from prior work [6] and 11 samples from this study. Of the 11 samples fabricated in this work, five were coated with PZT and subjected to the cleaning process described in Section 5.5, while six remained uncoated. Samples are categorized by γ ratio: low ($\gamma = 0.37$), medium ($\gamma = 0.69$), and high ($\gamma = 1.11$). A clear trend of decreasing Q-factor with increasing γ ratio was observed, with mean Q-factors of 3.05M, 0.98M, and 0.38M for the low, medium, and high γ designs, respectively. This trend is consistent with the results of Meira et al. [6] and is attributed to increased anchor loss at higher γ ratios, where elevated junction stress in the out-of-plane direction degrades mechanical isolation [6]. These results establish that while designs with higher γ ratios were expected to benefit from higher transduction efficiency, they incur a Q-factor penalty which must be taken into account in a fair comparison of IPA performance.




Figure 5.10: Q-factor versus γ ratio for 39 DSG samples grouped into low, medium, and high γ categories, with maximum variation of $\gamma \pm 0.07$ within each group. The dataset includes 28 DSGs from [6] and 11 DSGs fabricated in this study, five of which were coated with PZT.

5.6.2 IPA Displacement Evaluation

The IPA displacement of the five PZT-coated DSG samples was measured to test the hypothesis directly. A custom sample holder was designed to support the cap shell while applying AC forcer signals to the inner electrodes via wire bonds, Fig. 5.11(a). As in the Q-factor characterization, an LDV was positioned outside the vacuum chamber for motion detection; in this case the LDV was explicitly aligned to the antinode of the $n = 2$ wineglass mode, where maximum displacement occurs. Each sample was characterized over a range of applied electric fields from 200 kV cm^{-1} to 0 kV cm^{-1} in decreasing increments, as illustrated in the time-domain displacement data of Fig. 5.11(b). The displacement of all samples is summarized in Fig. 5.12. Samples 2n and 4n exhibited the largest displacement amplitudes, a $6.5\times$ enhancement over the lowest-performing samples.

To isolate the effect of γ ratio on transduction efficiency from the Q-factor variations characterized above, the measured displacement amplitude was normalized by the measured Q-factor. Since $|u_{res}| = Q_m(F_m/k_m)$, dividing the measured displacement by Q_m yields a metric proportional to the transduction efficiency F_m/k_m alone. This normalized displacement was compared with RAFT predictions evaluated at $Q_m = 1$ in Fig. 5.13. The experimental results were in agreement with the RAFT predictions, with both identifying




Figure 5.11: (a) Experimental setup for IPA DSG displacement characterization. The DSG was mounted on a PCB in vacuum and the outer electrodes were driven via wire bonds. An LDV aligned with an antinode of the $n = 2$ wineglass mode measured displacement amplitude. (b) Time-domain displacement measurement of sample 4n under IPA with incrementally decreasing applied electric field.




Figure 5.12: Displacement amplitude at the $n = 2$ wineglass mode antinode versus applied electric field, 0 kV cm^{-1} to 200 kV cm^{-1} , for five PZT-coated DSG samples actuated via IPA.

DSG 4n, at $\gamma = 1.09$, as the design of greatest transduction efficiency among those fabricated. This supports the central finding of the modeling: IPA efficiency does not increase monotonically with γ ratio, but is instead maximized at an intermediate value. The variation in normalized displacement across designs was substantial, with DSG 4n exhibiting $12.3\times$ greater transduction efficiency than DSG 2s.




Figure 5.13: Comparison of predicted displacement from the RAFT method at $Q_m = 1$ with measured displacement divided by Q_m for all PZT-coated DSGs. The experimental measurement and the modeling were both evaluated at an applied electric field of 200 kV cm^{-1} .

5.7 Discussion

A reduction in Q-factor was observed during active IPA operation in this study. This reduction was attributed to additional anchor loss introduced by the experimental setup, Fig. 5.11(a), rather than reflecting an intrinsic limitation of IPA. This interpretation is supported by prior work [164], in which IPA induced Q-factor reductions of only 7% for the $n = 2$ mode and 9% for the $n = 3$ mode were measured. Future implementations with improved mechanical boundary conditions are expected to reduce this setup-induced anchor loss.

The effect of the central stem on Q-factor was also evaluated, addressing the hypothesis advanced in Section 5.2.1 that stem removal would reduce anchor loss [3]. In the medium- γ group, mean Q-factors of 1.02M and 986k were measured for the stemless and stemmed designs, respectively; in the high- γ group, 535k and 382k were observed. These results indicate comparable anchor loss between the two configurations, and the hypothesis is therefore not supported. The finding is consistent with the FEA predictions of [100, 6], which indicate that acoustic energy radiates into the anchor predominantly through the flat bonding interface rather than through the central stem. Since the stem is not required for IPA, its retention or

removal can therefore be decided on fabrication grounds rather than on the basis of anchor loss.

The ultimate goal of IPA is to implement die-level, non-metallized gyroscope operation. Prior demonstrations of IPA-based gyroscope operation have employed electrostatic detection [164] and benchtop optical detection via two LDVs [221]. MOEMS provide a compelling pathway to fully integrated non-contact sensing [153]. Evanescent optical coupling has been proposed for HRG displacement sensing [154, 155], though experimental validation remains unreported. Among demonstrated MOEMS approaches, interference-based sensors have achieved a displacement sensitivity of ~ 0.14 pm/ $\sqrt{\text{Hz}}$ [157], and MZIs integrated on photonic chips project a minimum detectable displacement of ~ 10 pm on a bare fused silica surface at 6 dB SNR [158, 159], compared with ~ 0.66 pm for benchtop LDV detection referred to the 46 kHz DSG 4n resonance, over a common 20 Hz measurement bandwidth [230].

The drive-axis amplitude measured for DSG 4n, 8.35 nm, exceeds the ~ 10 pm minimum detectable displacement projected for a bare fused silica surface [159] by nearly three orders of magnitude. The sense-axis amplitude, however, is considerably smaller and scales with the input rotation rate. For open-loop operation it is given by the mechanical scale factor [9]:

$$SF_{mech} = \frac{2X_a\kappa\omega_x}{\sqrt{(\omega_y^2 - \omega_x^2)^2 + (2\mu_y\omega_x)^2}}, \quad (5.7)$$

where X_a is the drive-axis amplitude, κ is Bryan's factor, ω_x and ω_y are the drive- and sense-mode frequencies, and $\mu_y = \omega_y/2Q$ is the sense-mode decay rate. Taking $\kappa = 0.277$ [231], the minimum detectable rotation rate for DSG 4n is 0.21 $^\circ/\text{s}$ using chip-scale detection at a frequency split of 0.1 Hz, or 0.013 $^\circ/\text{s}$ with benchtop LDV detection [230].

5.8 Conclusion

The displacement of the $n = 2$ wineglass mode at resonance is formed by the product of transduction efficiency and Q-factor. This displacement determines the amplitude of the gyroscope output at resonance and, through Q-factor, the noise floor. In this chapter, the hypothesis that the γ ratio governs the efficiency of IPA was evaluated through RAFT modeling of ten DSG designs and experimental characterization of five PZT-coated samples. The electrode shaping analysis extended to the three-dimensional geometry of the DSG identified two electrode groups, of which the inner group contributes $9.5\times$ the modal force of the outer. Modal force was found to increase by a factor of 2,224 across the range of γ considered, partially offset by a $14.3\times$ increase in modal stiffness, so that predicted displacement is maximized at an intermediate value of γ rather than increasing monotonically. Measured Q-factor normalized displacement agreed with the predictions, and two routes to large displacement amplitude were identified: a low- γ design reaching 9.1 nm at 200 kV cm^{-1} on a Q-factor of 1.06M, and a high- γ design reaching comparable displacement through superior transduction efficiency despite a Q-factor of 278k. The cleaning process developed for PZT-coated resonators improved Q-factor by $4.8\times$ and $2.2\times$ for the $n = 2$ and $n = 3$ modes while retaining a functional piezoelectric response. Integrating an optimized design with a PIC for die-level, non-metallized gyroscope operation remains future work.

RQ1c: The dual-shell geometry that maximizes IPA displacement amplitude is one of intermediate γ ratio, since transduction efficiency rises with γ while the Q-factor falls to anchor loss over the same range. Displacement, the product of the two, is therefore optimized by balancing these effects rather than by maximizing transduction efficiency alone.

Part II

Self-Contained Pedestrian Navigation

A self-contained pedestrian navigation system must carry a complete IMU and perform the signal processing that fuses its measurements in real time. As shown in Section 1.5, the accuracy attainable by a ZUPT-aided INS of this kind is limited by the yaw-axis gyroscope, since ZUPT leaves the yaw angle and the yaw-axis gyroscope bias unobservable, so that their errors accumulate without correction while the remaining states are bounded by the update [48]. To this point, only the fabrication of a low-cost, high-performance gyroscope has been treated, which improves the best-case accuracy of ZUPT-aided INS. The gyroscope, however, is only one component of the complete system. Reaching the best accuracy requires reliable, repeatable stance-phase detection at each step, through which the observable states, the accelerometer biases, and the pitch and roll gyroscope biases are estimated. This stance-phase is difficult to identify during firefighter operation, and its detection is the subject of this Part; without it, the accuracy of the navigation solution deteriorates regardless of the quality of the gyroscope.

The chapters that follow focus on the pedestrian inertial navigation system as a whole and ZUPT. First, a platform is developed that is small enough to be worn, but contains an IMU, aiding sensors, and real-time processing; this platform is developed and evaluated in Chapter 6. The zero-velocity update depends on a stationary interval at the foot during each step, which is readily detected in walking but is not reliably present in the crawling gaits of firefighter operation, where the foot may never come fully to rest [25]. The extension of zero-velocity aiding to the crawling gait is developed in Chapter 7. A single foot-mounted IMU, however, cannot reliably provide a stance-phase across the full range of these motions, and so the fusion of multiple mounting locations, through which a dependable stance-phase is recovered across gaits, is developed in Chapter 8.

Chapter 6

Meso-Scale Inertial Navigation Platform

Chapters 3 through 5 developed the sensing element itself. Navigation, however, requires more than a single sensor. It requires a platform carrying a complete IMU, the aiding sensors the environment permits, and the computation that fuses them. In this chapter, a self-contained inertial navigation platform small enough to be embedded in the sole of a shoe is developed and experimentally evaluated¹. The platform is built from consumer-grade components at a cost of approximately \$61, with each sensor carried on its own board and joined to the main board through castellated edges so that a single sensor can be substituted without a redesign, Fig. 6.1. The processing unit is selected with sufficient margin to retain the accelerometer and gyroscope biases as filter states, and the firmware is reprogrammable through the Arduino development environment. This chapter addresses **RQ2a**, the navigation accuracy achievable from a self-contained inertial platform small enough to be embedded in the sole of a shoe, and **RQ2b**, the computational, thermal, and mechanical constraints

¹This chapter is based on the following peer reviewed publication: A. R. Parrish, C. Jao, D. Wang, and A. M. Shkel, “‘Sugar-Cube’ Navigation Board: Pedestrian Hardware Platform that Fits in the Sole of a Shoe,” *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Kaua’i, HI, USA, Mar. 2023.

that such integration imposes. Section 6.1 describes the system architecture. Section 6.2 presents the barometer-aided ZUPT-aided INS implemented within an EKF framework. Section 6.3 then details the hardware implementation, together with the computational, power, thermal, and mechanical constraints the form factor imposes. Section 6.4 reports the navigation accuracy of the platform, evaluated over a repeated indoor walking trajectory spanning flat floors, stairs, and an elevator ride. Finally, the chapter is concluded with Section 6.5.




Figure 6.1: “Sugar-Cube” conceptual architecture including a microcontroller, separate sensor PCBs, a Bluetooth module, battery regulation, and a 3D-printed case for stress relief and environmental protection.

6.1 System Architecture

The “Sugar-Cube” embedded INS consists of an IMU as well as desired aiding sensors (e.g. barometer), a processing unit that reads data from external sensors and performs navigation filtering computation, a method of communicating the navigation solution to an external display, and finally a method of mounting the system. The following sections describe the implementation in detail.




Figure 6.2: Illustration of “Sugar-Cube” architecture platform integrated in the sole of a shoe.

6.2 Navigation Algorithm

The “Sugar-Cube” Board was programmed with C/C++ using the Teensyduino library in the Arduino Integrated Development Environment (IDE). The barometer-aided ZUPT-aided INS was implemented using an EKF framework, Section 2.1.7, following [171, 37]. The runtime framework of the platform is described in Fig. 6.3. During the initialization process, sensor and Bluetooth communication is established while the IMU biases are estimated by averaging over a 10-second interval in which the “Sugar-Cube” is assumed to be stationary. The heading angle of “Sugar-Cube” could also be determined via the magnetometer during the initialization step of the algorithm. In the present implementation, two sources of aiding for the INS are included. ZUPT was employed to correct for IMU velocity drift while walking, Section 2.1.8, using the SHOE detector of Eq. 2.19 to detect stance-phase [51]. Barometer aiding was also employed to constrain the vertical drift of the INS [37]. In this initial demonstration, the magnetometer was not included in the navigation algorithm.




Figure 6.3: Flowchart of the runtime framework. The framework runs in a real-time loop as new sensor data is collected.

6.3 Hardware Implementation and System Constraints

A prototype of the “Sugar-Cube” Navigation Board is shown in Fig. 6.2; the system PCB is $39\text{ mm} \times 18\text{ mm} \times 10\text{ mm}$. The reported system integrates off-the-shelf components including a Teensy 4.0 microcontroller with Invensense ICM-42605 IMU, Invensense ICP-20100 Barometer, and Asahi Kasei Microdevices (AKM) AK-09915 Magnetometer. A summary of these components is provided in Table 6.1. The sensors selected for this module are consumer grade ($< \$10$) and tailored for foot-mounted navigation. The utilized barometer and magnetometer are designed to operate in ambient atmospheric conditions. The IMU selection is critical to the navigation performance. The selected IMU had sufficient bandwidth, sampling rate, and full-scale range to capture the dynamics of walking with sufficient accuracy, particularly the shock when the foot contacts the ground [58]. For the same reason, it is critical to have a powerful computation engine capable of processing these measurements in

Table 6.1: Specifications of sensors and other components used on the “Sugar-Cube” Navigation Board.

Component		Specifications	Samp. Rate [Hz]	Power Draw [mW]
IMU ICM-42605		Gyro: noise floor $0.0035 \text{ dps}/\sqrt{\text{Hz}}$, bias instability 5.15 deg/hr , range 2000 dps . Accel: noise floor $0.07 \text{ mg}/\sqrt{\text{Hz}}$, bias instability 0.041 mg/hr , range 16g	500 (max 8,000)	8.5
Barometer ICP-20100		$30\text{--}110 \text{ kPa}$ range, 1 Pa resolution, 0.5 Pa-RMS	25	10.5
Magnetometer AK-09915		$0.15 \mu\text{T}$ resolution, 49.1 mT range	100	6.5
Microcontroller 4.0	Teensy	600 MHz clock rate (1 GHz overclock), 1024 kB RAM	—	490
Bluetooth BL653 μ		Bluetooth 5.0, up to 1 MBps	—	26
Lithium Ion Battery		500 mAh capacity ($30 \times 20 \times 9.5 \text{ mm}$ size)	—	—

real-time. In this implementation, all sensors are soldered on separate PCBs, with relevant pull-up resistors and filtering capacitors, while castellated edges provide the connection to the main PCB; this provides some margin of flexibility for the platform. For example, if a custom IMU is used, only the IMU PCB needs to be replaced.

Sensor communication is performed using SPI for the barometer with a 25 Hz sampling rate, while I²C was used for the IMU/magnetometer at sampling rates of $500/10 \text{ Hz}$, respectively. In order to minimize the time the microcontroller takes to receive each sample, each of the sensors is configured to provide continuous measurements at a specific data rate. When new data is available, a data ready bit triggers the corresponding interrupt service routine on the microcontroller. The sensor firmware was developed based on existing libraries (ICM-42605 [232] / AKM09915 [233]) or custom developed (ICP-20100). The data is transmitted via USB 2.0 connection at 500 Hz or via the Bluetooth module (BL653 μ) at $\sim 4 \text{ Hz}$ to an exter-

nal display. Only two bytes of the navigation solution are sent to the BL653 μ per navigation update, which uses its internal buffer to handle Bluetooth communication; this minimizes the impact of Bluetooth communication on the navigation update rate. The Laird BL654 Bluetooth USB Adapter acts as a Virtual Serial Port (VSP) in order to facilitate communication with an external display; any computer with Serial communication can receive the results. However, it is possible to circumvent the use of the BL654 with the development of a customized communication protocol.

6.3.1 Computational Constraints

Due to the high computational power of Teensy 4.0, the platform is capable of including accelerometer and gyroscope biases in the EKF states, in addition to the position, velocity, and attitude, at an update rate of 500 Hz. In prior implementations, these states were omitted in order to improve the navigation update rate [42]. The navigation update rate is limited by computation of the state transition matrix exponential, which is approximated by a seven-term power-series. In one update, it takes 0.58 ms to compute the matrix exponential, 0.46 ms for the remaining EKF calculations, and 0.2 ms to receive the next IMU sample. Currently, the update rate is constrained to 500 Hz in order to eliminate issues with timing of IMU samples. An update rate of ~ 760 Hz is possible without this limitation and even 1 kHz could be achieved if the Teensy 4.0 were overclocked to 816 MHz; however, a significant overclocking would require active cooling. The “Sugar-Cube” module can also be used purely for data collection, which can output sensor data at significantly higher data rates via USB (e.g. 2000 Hz for IMU).

6.3.2 Power and Thermal

The “Sugar-Cube” Board integrates a 3.7 V Lithium-Ion battery as the power source, which is connected to a voltage boosting module (PAM2401) for 5 V output. This system consumes roughly 550 mW during operation, with the majority of the power drawn from the Teensy 4.0

operating at 600 MHz (490 mW). The sensors and Bluetooth module only consume roughly 50 mW of power in total, and the efficiency of the voltage boosting module will marginally decrease as the battery voltage decreases, increasing the overall power consumption. The thermal profile of the “Sugar-Cube” is also shown in Fig. 6.4 after one hour of operation. The hottest parts are the ARM Cortex M7 microprocessor of the Teensy 4.0, which reached a temperature of 48.4 °C, the 3.3 V voltage regulator on Teensy, and the battery boosting circuit (42 °C). The temperature of the IMU was ~ 38 °C, which will likely cause sensor bias shift during the initial operation of the platform as it reaches steady-state temperature. The battery used in this demonstration is of dimension 30 mm $\times$ 20 mm $\times$ 9.5 mm with 3.7 V output and 500 mA h capacity yielding ~ 3 hours of operation.




Figure 6.4: The thermal profile of the “Sugar-Cube” after one hour of operation is shown with the corresponding color map.

6.3.3 Packaging and Mechanical Loading

A 3D-printed case was designed to minimize the stress exerted on the sensors due to mounting of the system in the sole of the shoe. In this implementation, a slot was drilled into the sole of a shoe in order to house the relevant battery and “Sugar-Cube” board.

6.4 Performance Evaluation

The navigation performance of the platform was evaluated in a real-world scenario, summarized in Fig. 6.5. The scenario was repeated 10 times spanning flat planes, stairs, and an elevator ride for one floor. The platform was mounted in the sole of the shoe near the heel, as shown in Fig. 6.2; however, it can be mounted anywhere on the shoe or the body.


Accuracy was evaluated using CEP, the median error in the North-East plane, together with RMSE in the vertical direction. The total distance traveled in the experiments was 125 m and the duration was 2.75 minutes. The CEP was 0.66 m, which is 0.53 % of the trajectory length, and the vertical RMSE was 0.16 m. The navigation solution appears to drift to the left. This is hypothesized to be due to stress on the IMU from compression of the shoe during the stance-phase.

6.5 Conclusion

In this chapter, a meso-scale inertial navigation platform occupying 7.1 cm^3 , excluding battery and case, was designed, implemented, and experimentally evaluated. Consumer-grade inertial, barometric, and magnetic sensors were mounted on separate boards with castellated edges, allowing an individual sensor to be substituted without redesigning the main board. A barometer-aided ZUPT-aided INS was implemented in an EKF framework at an update rate of 500 Hz, with sufficient computational headroom to retain the accelerometer and gyroscope biases as filter states. The platform drew approximately 550 mW, yielding roughly three hours of operation, and the IMU settled near 38°C after one hour of continuous use. Over ten repetitions of a 125 m, 2.75-minute indoor trajectory spanning flat floors, stairs, and an elevator ride, the platform achieved a CEP of 0.66 m, corresponding to 0.53 % of the trajectory length, with a vertical RMSE of 0.16 m.

Placing the platform in the sole of the shoe locates the IMU where the stance-phase is most reliably available for walking gait, and the computational power of the processing unit allows the sensor biases to be estimated rather than omitted to preserve the update rate, so that each zero-velocity measurement corrects accumulated drift rather than only bounding velocity error. Both directly impact the accuracy a pedestrian INS can sustain over a prolonged trajectory. The residual lateral drift, attributed to compression of the sole during stance, and the bias shift observed as the platform warms indicate that packaging stress and thermal transients remain to be compensated. The demonstration reported here was performed while walking; the gaits encountered in emergency response are the subject of the next chapter.

RQ2a: A platform small enough for the sole of a shoe attains meter-level accuracy, here a CEP of 0.66 m, or 0.53 % of the 125 m trajectory. **RQ2b:** Its integration is bounded by computation, which limits the filter states retainable at the update rate; by processor heat, which biases the IMU as it warms toward steady-state; and by mechanical loading of the sole during stance, which stresses the package and appears as residual lateral drift.



Circular Error Probable




Figure 6.5: Navigation Solution and Circular Error Probable for the “Sugar-Cube” board over a 125 m trajectory spanning flat planes, an elevator, and stairs. The experiment was performed 10 times.

Chapter 7

Stance-Phase Detection and IMU Placement for Firefighter Crawling

The platform of Chapter 6 was demonstrated while walking, but the stance-phase on which ZUPT depends is not reliably present in the crawling gaits of firefighter operation. This chapter addresses **RQ3a**, the detection of the stance-phase and the placement of the IMU for ZUPT-aided pedestrian inertial navigation during firefighter crawling, where the stationary interval assumed by conventional foot-mounted implementations is not reliably available¹. Section 7.1 reports a parameter search over stance-phase detector type, window length, and detection threshold for heel- and knee-mounted IMUs during crawling. Section 7.2 then evaluates four candidate mounting positions across two crawling methods used by firefighters, measuring the residual velocity during the detected stance-phase directly with an electromagnetic position sensor. Finally, the chapter is concluded with Section 7.3. This chapter

¹This chapter is based on the following peer-reviewed publications:

- A. R. Parrish, C. Jao, and A. M. Shkel, “Stance-Phase Detection for ZUPT-aided INS using Knee-Mounted IMU in Crawling Scenarios,” *IEEE Sensors Letters*, vol. 6, no. 5, 2022 [182].
- A. R. Parrish, C. Jao, D. Wang, and A. M. Shkel, “Study of IMU Mounting Position for ZUPT-Aided INS in the Case of Firefighter Crawling,” *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Kaua’i, HI, USA, Mar. 2023 [183].

establishes the knee as the placement that recovers a stance-phase during crawling, as the shoe does during walking. Because neither placement alone serves both gaits, and firefighters move between them within a single deployment, the two are fused in Chapter 8 to sustain navigation across walking, crawling, and the transitions between them.

7.1 Stance-Phase Detection of Crawling

When crawling, the foot may not reach a fully stationary state during ground contact, so ZUPT with a foot-mounted IMU becomes unreliable [25]. The knee is a candidate alternative, since it is a primary point of contact with the ground in many crawling gaits. The stance-phase detector, window length, and detection threshold are established for walking, but their appropriate values for a knee-mounted IMU during crawling are unknown. This section reports a parameter search over all three, conducted for heel- and knee-mounted IMUs, to identify the configuration minimizing navigation error.

7.1.1 Data Collection

A series of crawling experiments was performed with both knee-mounted and heel-mounted IMUs. The crawling was performed in roughly a straight line with a distance of 42.6 m, which was measured by a ruler. Ten trials were performed for each IMU mounting position. This experiment considers hand/knee crawling [69], which is similar to the knee-dragging crawling method described in Rantakokko et al. [25], although the knees left the ground during swing phase in this experiment. As shown in Fig. 7.1, the subject was supported by their hands and knees, moving the opposite hand/knee simultaneously with a stride length of approximately 61 cm. The subject of this experiment is a 183 cm, 77 kg male who crawled at a rate of approximately 60 strides per minute on a hard floor. The subject was stationary for roughly 10 s before crawling for calibration of the accelerometer and gyroscope biases and used a metronome to aid in maintaining a constant crawling pace. The subject wore

knee-pads which included a rigid outer-frame, padding for the knee, and straps to secure the knee-pad and minimize relative motion between the knee-pad and knee. An industrial-grade IMU, the VectorNav VN-200 [7], was used with a sampling rate of 800 Hz; the same IMU and sampling rate are used in all experiments reported in this chapter. The IMU was rigidly mounted to the heel via a custom 3D-printed mount, and mounted to the outside of the knee via the fabric of the knee-pad, Fig. 7.1. The setups of the heel-mounted and knee-mounted crawling experiments were identical, but were performed on separate trials.

7.1.2 Data Processing

The ZUPT-aided INS was implemented using an EKF framework. Details of the EKF are documented in [171]. In each trial, the initial yaw angle was aligned after 10 m of travel from the origin. A parameter search was conducted for the crawling data in order to minimize the CEP of the 10 trials for each mounting position. Three parameters were considered: the window size, stance-phase detector type, and detection threshold.




Figure 7.1: Diagram of the mounting positions used in the crawling experiments and crawling method. The IMU is attached to the side of the knee using the fabric of the knee-pad. The IMU is rigidly attached to the heel using a custom 3D-printed mount. The IMU used is the VectorNav VN-200 [7].

7.1.3 Stance-Phase Detectors

Four stance-phase detectors based purely on IMU measurements are considered [51]. Each detector produces a test statistic which can broadly be interpreted as the stability of the IMU at that moment [56]. The detectors are described as follows:

1. Acceleration Magnitude Detector (AM)

$$T(\vec{z}_n) = \frac{1}{\sigma_a^2 N} \sum_{k=n-N/2}^{n+N/2-1} (\|\vec{a}_k\| - g)^2 < \gamma \quad (7.1)$$

2. Acceleration Moving Variance Detector (AMV)

$$T(\vec{z}_n) = \frac{1}{\sigma_a^2 N} \sum_{k=n-N/2}^{n+N/2-1} \|\vec{a}_k - \bar{\vec{a}}_n\|^2 < \gamma \quad (7.2)$$

3. Angular Rate Energy Detector (ARE)

$$T(\vec{z}_n) = \frac{1}{\sigma_\omega^2 N} \sum_{k=n-N/2}^{n+N/2-1} \|\vec{\omega}_k\|^2 < \gamma \quad (7.3)$$

4. SHOE Detector, restated here from Eq. 2.19 for comparison with the three alternatives

$$T(\vec{z}_n) = \frac{1}{N} \sum_{k=n-N/2}^{n+N/2-1} \left(\frac{1}{\sigma_a^2} \left\| \vec{a}_k - g \frac{\bar{\vec{a}}_n}{\|\bar{\vec{a}}_n\|} \right\|^2 + \frac{1}{\sigma_\omega^2} \|\vec{\omega}_k\|^2 \right) < \gamma \quad (7.4)$$

In Eqs. (7.1)–(7.4), N is the number of consecutive measurements included in the window (discussed in Section 7.1.4), $\vec{a}_k$ and $\vec{\omega}_k$ are the accelerometer and gyroscope readouts at time index k , $\vec{z}_n = \{\vec{a}_k, \vec{\omega}_k\}_{k=n-N/2}^{n+N/2-1}$ is the set of IMU readouts over the window, $\bar{\vec{a}}_n$ is the average value of the accelerometer readouts over the window, σ_a and σ_ω are related to the white noise level of the accelerometer and gyroscope, g is the magnitude of gravity, and γ is the detection

threshold. In all the detectors, a stance-phase is declared when the test statistic $T(\mathbf{z}_n)$ is lower than the threshold γ . The white noise levels of the gyroscope and accelerometer used in these experiments are $3.5 \text{ m}^\circ/\text{s}/\sqrt{\text{Hz}}$ and $0.14 \text{ mg}/\sqrt{\text{Hz}}$, respectively. Therefore, the ratio of the noise levels squared is $\sigma_a^2/\sigma_\omega^2 = 505$, meaning the gyroscope measurements are heavily weighted in the SHOE detector.

7.1.4 Window Size

As described, the stance-phase detectors consider a series of consecutive IMU measurements when calculating the test statistic, termed the window size. A higher window size will smooth the test statistic and can make gait cycles more clear by mitigating false-alarms, which typically have a smaller duration than the stance-phase of the gait cycle [65]. However, if the window size is too large, the detected stance-phase duration will become shorter and potentially have a negative impact on the ZUPT algorithm. If the ZUPT-aided INS is implemented in real time, a large window size also incurs a delay of half the window, while waiting for the rest of the data. As shown in equations (7.1), (7.2), (7.3), and (7.4), this experiment considers the window centered around the current time, from $n - N/2$ to $n + N/2 - 1$. This is important when considering large window sizes in order to ensure the test-statistic is representative of the current time. One method of selecting the window size is to match it to the expected length of the stance-phases [51]. In this experiment, the window size is swept from 0.01 s (8 samples) to 1 s (800 samples).

7.1.5 Parameter Search Results

The results of the parameter search for crawling with the IMU mounted on the heel and knee are shown in Fig. 7.2. The optimal CEP for the heel-mounted IMU was 0.95 m with a corresponding RMSE of 1.05 m and occurs with (1) the ARE detector, (2) a window size of 0.07 s (56 samples), and (3) a detector threshold of $e^{5.1}$. Another parameter selection for heel-mounted IMU rendered a slightly smaller CEP of 0.93 m. This used (1) the AMV

detector, (2) a window size of 0.75 s (600 samples), and (3) a detector threshold of $e^{4.54}$. However, the RMSE for these parameters was considerably larger, 1.48 m, rendering this not the optimal parameter selection. It appears the SHOE and ARE detectors perform best at lower window sizes. The minimum CEP for knee-mounted IMU crawling was 0.22 m with a corresponding RMSE of 0.43 m. This used (1) the AMV detector, (2) a window size of 0.9 s (720 samples)², and (3) a detector threshold of $e^{0.87}$. In the experiments, the knee-mounted IMU appears to render improved navigation accuracy for most detector and window size selections. Using the optimal parameters for both, the knee-mounted IMU renders a 77.3 % / 59.1 % improvement in CEP / RMSE respectively compared to the heel-mounted IMU. It is also apparent that accelerometer-based ZUPT detectors (AM/AMV) outperform ARE and SHOE detectors for most window sizes, particularly at low window sizes. The navigation accuracy for all 10 trials is shown in Fig. 7.5 and confirms the knee-mounted IMU position for this type of crawling offers improved navigational accuracy, as compared to the heel-mounted IMU.

7.1.6 Detector Statistic Comparison

In order to better understand the parameter search results, an examination of representative detector statistics was performed. First, in Fig. 7.3 the SHOE detector statistic for walking, crawling with heel-mounted IMU, and crawling with knee-mounted IMU for three representative gait cycles is presented. This figure highlights the challenges associated with stance-phase detection of crawling. The use of the SHOE detector with a window size of 0.05 s typically yields good performance for walking, as shown in Fig. 7.3(a); however, this does not render the same results for crawling, Fig. 7.3(b),(c). In fact, for knee-mounted crawling, it is difficult to even observe periodic behavior under these detection parameters; this is reflected in Fig. 7.2(b), as the optimized CEP is 13.3 m. Second, the heel-mounted

²The minimum knee-mounted CEP at smaller window sizes was 0.26 m with a corresponding RMSE of 0.50 m. This used (1) the AMV detector, (2) a window size of 0.15 s (120 samples), and (3) set the detector threshold at $e^{-0.44}$.




Figure 7.2: Parameter search for heel-mounted (a) and knee-mounted (b) IMU crawling. The window size was swept along with detector threshold values (not shown) for all 4 detectors considered. The threshold that rendered the minimum CEP over 10 crawling trials for a given window size was selected and displayed.




Figure 7.3: SHOE Detector statistic comparison with a window size of 0.05 s (40 samples). Three typical gait cycles are shown for (a) walking with heel-mounted IMU, (b) crawling with heel-mounted IMU, and (c) crawling with knee-mounted IMU. The estimated stance-phase for all figures is highlighted in red; this was derived visually for (a). In order to improve accuracy the stance-phase was derived for (b) using AMV detector with $e^{4.51}$ and $W = 0.9$ s (720 samples) and (c) using AMV detector with $e^{-2.51}$ and $W = 0.05$ s (40 samples). Both thresholds are derived from the parameter sweep discussed in Section 7.1.5. Note the stance-phase for (b) and (c) has been smoothed to appear continuous during these periods.




Figure 7.4: Knee-mounted IMU crawling detector statistic comparison. Three typical gait cycles are shown for (a)/(b) ARE detector with a window size of $W = 0.07s$ (56 samples) and $W = 0.3s$ (240 samples) respectively and (c)/(d) AMV detector with a window size of $W = 0.05s$ (40 samples) and $W = 0.3s$ (240 samples) respectively. The thresholds selected are all derived from the parameter sweep. Finally (e) shows one gait cycle and highlights the detected stance-phase periods for all four detectors.

IMU presents more periodic behavior, however considering these parameters render near the best performance of the parameter sweep, this is not sufficient for reliable detection. Therefore, stance-phase detection in crawling can likely benefit from a different approach.

As discussed, accelerometer based ZUPT detectors appear to render improved navigation performance compared with gyroscope based ZUPT detectors for knee-mounted IMU crawling. A comparison of these detectors for knee-mounted IMU crawling is shown in Fig. 7.4. Similar to Fig. 7.3(c), it can be observed that stance-phase detection is difficult for the ARE detector, as it had a false alarm at low window sizes. The use of a larger window size smooths the data and effectively filters out the shorter, false alarm detections. A possible explanation for this behavior is that the IMU may not translate significantly when the knee is on the ground; however, the body still pivots around the knee, slightly rotating the IMU and rendering a less stationary stance-phase, as observed by the gyroscopes. Another possible explanation is that there appears to be a significant portion during the swing phase where little IMU rotation occurs, causing false alarm detection. These issues are not observed to be as pronounced with the AMV/AM detectors and render significantly more accurate results for all window sizes.

Under the conditions of this experiment, relocating the IMU from the heel to the knee reduced the CEP by more than a factor of four, with accelerometer-based detectors outperforming those weighting angular rate. Two limitations of the study motivate what follows. First, the knee was treated as a single location, although the knee pad affords several distinct mounting points that differ in their distance from the contact patch and in how much the limb rotates about them. Second, the stability of each placement was inferred from the resulting navigation error rather than measured, so the residual velocity incurred by the zero-velocity assumption was never observed directly. Both are addressed in Section 7.2.




Figure 7.5: Final position errors from the heel and knee-mounted IMUs are shown with the optimal parameters from the parameter sweep. The heel-mounted IMU trajectory uses the ARE detector with a window size of 0.07 s (56 samples) and renders a CEP of 0.95 m. The knee-mounted IMU trajectory uses AMV detector with a window size of 0.9 s (720 samples) and renders a CEP of 0.22 m.

7.2 Evaluation of Optimal Mounting Position for Firefighter Crawling

This section evaluates four mounting positions on the foot and knee across two crawling methods used by firefighters during smoke diving. Residual velocity during the detected stance-phase is measured directly with an electromagnetic position sensor, rather than inferred from the resulting navigation error, and is reported alongside stance-phase duration and navigation accuracy over a straight-line trajectory.

Three metrics govern the choice of mounting position. Most importantly, the residual velocity of the IMU during the detected stance-phase determines how accurately IMU biases can be estimated and corrected. Second, the duration of the stance-phase impacts the reliability of stance-phase detection. Finally, a high shock level necessitates high dynamic range,

bandwidth, and navigation update rate in order to accurately capture the motion [56]. In this study, all mounting positions were observed to have relatively low shock levels, generally less than 16 g and 2000° s^{-1} , as crawling is not a highly dynamic motion.




Figure 7.6: Illustration of hand/knee crawling, duck crawling, and four mounting positions for the IMU and electromagnetic position sensors used in this experiment. A close up of the IMU and electromagnetic position sensor and the corresponding magnetic source is provided on the right. The IMU is the VectorNav VN-200 [7], and the electromagnetic tracker is the Polhemus Patriot [8].

7.2.1 Crawling Methods

Two crawling methods commonly employed by firefighters when searching a building are investigated in this study, illustrated in Fig. 7.6. The first is referred to in literature as hand/knee crawling [69]. The subject is supported by their hands and knees, moving the opposite hand/knee simultaneously. The second style is referred to as the “duck crawling” method [70], in which the subject has one knee on the ground with the other leg in front. The palm is placed on the ground allowing the rear leg to swing forward until contact is made with the front foot. This method is commonly employed by firefighters to advance hose lines as it leaves one hand free. This method also allows firefighters to shift their weight onto their back leg so if they were to encounter a hole or ledge they would not fall in, and

the method enables greater awareness of their surroundings since firefighters are not facing the floor.

Table 7.1: Summary of velocity uncertainty and stance-phase duration across 50 gait cycles, as well as navigation accuracy for the 10 trials of length 42.6 m for the hand/knee crawling method.

Mounting Position	Hand/Knee Crawling					
	CEP	Vertical RMSE	Avg. Stance Phase Duration	Velocity Uncertainty [cm/s]		
	[m]	[m]	[s]	North	East	Down
Outside Knee	1.87	1.25	0.059	1.40	1.88	0.94
Center Knee	0.84	0.25	0.38	0.82	0.91	0.92
Inside Knee	2.5	1.68	0.076	3.4	2.19	4.81
Heel	1.62	1.69	0.037	12.33	22.62	19.49

Table 7.2: Summary of velocity uncertainty and stance-phase duration across 50 gait cycles, as well as navigation accuracy for the 10 trials of length 42.6 m for the duck crawling method.

Mounting Position	Duck Crawling					
	CEP	Vertical RMSE	Avg. Stance Phase Duration	Velocity Uncertainty [cm/s]		
	[m]	[m]	[s]	North	East	Down
Outside Knee	4.17	5.38	0.021	2.83	1.56	3.07
Center Knee	0.79	0.23	0.197	0.90	0.32	0.86
Inside Knee	3.42	4.83	0.163	22.58	8.45	17.72
Heel	3.12	1.11	0.059	11.37	15.88	21.28

7.2.2 Data Collection

A series of crawling experiments was performed with knee-mounted and heel-mounted IMUs, Fig. 7.6. The experiments were performed wearing firefighter PPE including bunker pants and coat, helmet, and air pack, which added weight of roughly 20 kg. Firefighter PPE was worn as it is known to change the gait dynamics including range of motion, center of gravity, and increase dynamic instability [71]. In addition, knee-pads were worn over the bunker pants, both to protect the knees and house the IMU. The experiments were performed by the same subject as Section 7.1; the subject is not a trained firefighter, but replicated the crawling motion as closely as possible. Note that when performing duck crawling, the




Figure 7.7: Final position error for all mounting positions: outside knee, center knee, inside knee, and heel. The final position is shown in red for hand/knee crawling and blue for duck crawling. Each series of experiments was performed 10 times and the resulting CEP and vertical RMSE are detailed in Tables 7.1 and 7.2.




Figure 7.8: Histogram of residual stance-phase velocity for the heel-mounted IMU performing hand/knee crawling. The standard deviation of the fitted normal distribution (in red) is defined as the velocity uncertainty.

IMUs were always located on the rear leg in this experiment. Both crawling methods were performed with a metronome to aid a uniform crawling pace of 60 and 30 strides per minute for hand/knee crawling and duck crawling, respectively.

A 3D-printed custom knee-pad insert, shown in Fig. 7.6, was designed for this experiment to enable mounting of the IMU in the center of the knee-pad. While this IMU mounting position was demonstrated to be the most stable in this study, placing an IMU in the selected location generates additional packaging requirements in order to minimize the compressive load of the firefighter on the knee as they crawl. Without properly attending to the load distribution, the load will stress MEMS resonators from the package to the sensing element via the anchor point of the sensor element and will be observed as sensor bias.

Two datasets were collected in this study to evaluate the navigation accuracy and residual velocity uncertainty of the mounting positions. The first dataset included 50 gait cycles collected for each mounting position and crawling method and is used to evaluate the residual velocity uncertainty during stance-phase. An alternating current electromagnetic position

sensor [8] was rigidly mounted to the IMU, as shown in Fig. 7.6, and was employed to estimate the velocity of the IMU during the detected stance-phase. The resolution of the position sensor decreases as the magnetic sensors move further from the source; within 30.5 cm of the source the system has a resolution of 0.012 mm and within 61 cm a resolution of 0.089 mm. The position measurements were differentiated in order to estimate the velocity of the IMU during the estimated stance-phase. The second dataset was collected to evaluate the navigation performance of each mounting position in a real-world scenario, summarized in Fig. 7.7 and Tables 7.1 and 7.2. The subject crawled along a straight line of length 42.6 m ten times.

7.2.3 Data Processing

Stance-phase was detected using the parameters identified in Section 7.1. The AMV detector was used for the knee mounting positions with a window size of 0.15 s. The SHOE detector [51] was used for the heel-mounted IMU with a window size of 0.07 s. When stance-phase was detected, the velocity of the IMU (residual velocity) was estimated using the electromagnetic position sensor. The residual velocity collected across 50 gait cycles was used to create a distribution, shown in Fig. 7.8, which was assumed to be Gaussian and the standard deviation was defined to be the velocity uncertainty. In order to eliminate misdetection of stance-phase, the highest 1 % of residual velocities were omitted.

ZUPT-aided INS was implemented for the navigation experiments in the EKF framework of Section 7.1. Initially, the subject was stationary for 10 s before the experiment began in order to estimate the IMU biases. Next, after the first 10 m of travel, the yaw angle of all trajectories was aligned in order to eliminate alignment error.

While the stance-phase detection threshold can be optimized for the navigation trials based on accuracy, the threshold for the 50 gait cycles cannot be optimized using the same method. In order to rigorously determine the stance-phase threshold for these trials, we first manually

set the threshold for the 50 gait cycles in order to derive an initial velocity uncertainty. This was used to tune the stance-phase threshold for the 10 navigation trials. Following, step detection based on the algorithm described in [234] was performed in order to estimate the average stance-phase duration across the 10 navigation trials. Finally, the stance-phase threshold for the 50 gait cycles was adjusted to match the stance-phase duration of the navigation trials and the process was repeated again to ensure convergence.

7.2.4 Results

A summary of the velocity uncertainty, stance-phase duration, and navigation accuracy of the IMU mounting positions is provided in Tables 7.1 and 7.2 and Fig. 7.7. The velocity uncertainty values illustrate that the knee is more stable compared to the heel for the types of crawling used and, further, the center-knee is the most stable mounting position during the stance-phase. The average stance-phase duration is the longest for the center knee mounting position as well, indicating this position will enable the most reliable stance-phase detection. The best navigation performance was achieved via the center knee with a CEP of 0.84 m and 0.79 m and vertical RMSE of 0.25 m and 0.23 m for hand/knee and duck crawling, respectively. This marks a 1.9 and 3.7 times improvement in CEP for the hand/knee and duck crawling methods, respectively, for the center-knee compared to any of the other mounting positions.

Placement of the IMU package in the center of the knee, while minimizing the applied force, increases the complexity of the setup. In this experiment it was observed that even though the velocity uncertainty was generally lower for knee mounting positions, the navigation accuracy was comparable to the heel mounting position, excluding the center knee. One reason for this is the underestimate of the trajectory length for the inside and outside knee mounting positions. An underestimate of the residual velocity has been linked to an underestimate of the trajectory length, but this may not be the case in the conducted experiment since the

uncertainty was precisely evaluated for the experimental conditions [76]. Another potential explanation is that zero-velocity pseudo-measurements in the EKF implementation assumed uncorrelated, zero-mean, Gaussian measurement noise when in reality the IMU will still be moving, causing unmodeled errors. The movement of the IMU during stance-phase is generally due to rotation of the body, for example the knee or foot, with respect to the contact point creating a moment arm. Crawling methods, in particular, involve rotation of the body around the knee, as discussed in Section 7.1, exacerbating these effects. This same effect likely contributes to the high vertical RMSE for the inside and outside knee for the duck crawling methods.

Measuring residual velocity directly confirms that the center of the knee pad is the placement least disturbed during stance, and that it holds across both crawling methods. It provides no such anchor while walking, however, and firefighters alternate between walking and crawling within a single deployment. Choosing one placement therefore concedes accuracy in whichever gait it does not serve. Chapter 8 instead instruments both the shoe and the knee, and couples the two solutions so that each is anchored by the other in the gait where its own stance-phase is least reliable.

7.3 Conclusion

In this chapter, stance-phase detection and sensor placement were studied to improve the accuracy of ZUPT-aided pedestrian inertial navigation during firefighter crawling. A parameter search over stance-phase detector type, window length, and detection threshold was conducted for heel- and knee-mounted IMUs over a 42.6 m crawling trajectory. Accelerometer-based detectors were shown to outperform the angular rate energy and SHOE detectors for a knee-mounted IMU, which was attributed to rotation of the body about the knee contact point rendering a less stationary stance-phase as observed by the gyroscopes. Under the optimal configuration for each, the knee-mounted IMU rendered a CEP of 0.22 m against

0.95 m for the heel-mounted IMU, a 77.3 % improvement, with a corresponding 59.1 % improvement in vertical RMSE. Four mounting positions were then evaluated across hand/knee and duck crawling, with the residual stance-phase velocity measured directly by an electromagnetic position sensor over 50 gait cycles. The center of the knee pad was identified as the placement with the longest stance-phase duration and the lowest velocity uncertainty, yielding a CEP of 0.84 m and 0.79 m for hand/knee and duck crawling, a $1.9\times$ and $3.7\times$ improvement over the remaining placements.

These results establish that the detector configuration conventionally used for walking is non-optimal for crawling, where the stance-phase at the shoe is degraded in both stability and duration. In answer to the applicable research question, the extension of ZUPT-aided navigation to firefighter crawling is enabled at the sensor and detector level by mounting the IMU at the center of the knee pad and configuring the stance-phase detector for the crawling gait, which together recover a reliable stance-phase where a walking-configured foot-mounted solution does not. This placement provides no such anchor while walking, however, and firefighters alternate between walking and crawling within a single deployment. The joint use of both placements across walking, crawling, and the transitions between them is developed in Chapter 8.

RQ3a: During firefighter crawling, the stance-phase is reliably detected by an accelerometer-based detector rather than one weighting angular rate, and the IMU is best mounted at the center of the knee pad, which yields the longest stance-phase duration and the lowest residual velocity among the placements evaluated. This configuration recovers a dependable stance-phase where the walking-configured foot-mounted solution does not.

Chapter 8

Dual-IMU ZUPT-Aided Inertial Navigation for Seamless Firefighter Walking and Crawling

This chapter develops a dual-IMU framework that extends ZUPT-aided pedestrian inertial navigation across the walking, crawling, and transitional gaits of firefighter smoke-diving operations, for which no single mounting position provides a reliable stationary interval¹. Chapter 7 identified the knee as the placement that anchors crawling and the shoe as the placement that anchors walking, but neither alone serves both gaits. This chapter addresses **RQ3b**, how the navigation accuracy that ZUPT provides during walking can be maintained across the crawling gaits of firefighter operation, where no single sensor placement provides a reliable stance-phase in both. To combine their strengths, Section 8.1 describes the dual-IMU architecture, and Section 8.2 develops the joint estimator, in which shoe- and knee-mounted

¹This chapter is based on the following publication in preparation:

- A. R. Parrish, E. Sangenis, D. Pachori, C. S. Jao, and A. M. Shkel, “Dual-IMU ZUPT-Aided Inertial Navigation for Seamless Firefighter Walking and Crawling,” *IEEE Sensors Journal*, Under Review, 2026.

solutions are coupled in a joint 30-state error-state EKF through the ZPD constraint on the rigid lower leg. A multi-gait dataset combining walking and multiple crawling methods was collected for evaluation, Section 8.3, and a random forest classifier was developed to identify the type of motion and adjust the EKF parameters accordingly, Section 8.4. Section 8.5 presents the multi-gait navigation results. Finally, the chapter is concluded with Section 8.6.

8.1 Dual-IMU Architecture

The shoe- and knee-mounted solutions are propagated independently and coupled in a joint 30-state error-state EKF through a kinematic constraint on the rigid lower leg, so that the IMU anchored in the current gait constrains the one that is not. The constraint requires no additional ranging hardware, and a random forest classifier configures the stance-phase detector and measurement noise for the gait being performed. The framework is evaluated on a multi-gait dataset in which subjects alternate between walking and both crawling methods within a single continuous trajectory.

8.2 EKF Formulation and Kinematic Constraint

The proposed dual-IMU framework is realized in a joint 30-state error-state EKF. The per-IMU error state $\delta\vec{x}_i$ ($i \in \{s, k\}$, shoe or knee) consists of the attitude, velocity, position, gyroscope-bias, and accelerometer-bias errors,

$$\delta\vec{x}_i = \begin{bmatrix} \delta\vec{\theta}_i^\top & \delta\vec{v}_i^\top & \delta\vec{p}_i^\top & \delta\vec{b}_{g,i}^\top & \delta\vec{b}_{a,i}^\top \end{bmatrix}^\top \in \mathbb{R}^{15}, \quad (8.1)$$

where $\delta\vec{\theta}_i$ is the body-frame attitude error, $\delta\vec{v}_i$ and $\delta\vec{p}_i$ are the velocity and position errors expressed in the NED frame, and $\delta\vec{b}_{g,i}$ and $\delta\vec{b}_{a,i}$ are the gyroscope and accelerometer bias errors along the three axes of the IMU body frame. The joint state stacks the two per-IMU




Figure 8.1: IMU mounting positions on the foot and knee for the dual-IMU setup. Lever arms from each IMU to the ankle were measured per subject with a ruler and are expressed in the respective IMU body frames shown. The average values were $\vec{L}^s = [0 \ 0 \ 7.6]^\top$ cm and $\vec{L}^k = [-38.9 \ 0 \ 2]^\top$ cm, where the 2 cm Z^k -component accounts for the knee-pad standoff. Residual lever-arm error is absorbed by the ZPD measurement noise.

error vectors,

$$\delta \vec{x} = \begin{bmatrix} \delta \vec{x}_s^\top & \delta \vec{x}_k^\top \end{bmatrix}^\top \in \mathbb{R}^{30}. \quad (8.2)$$

In the prediction step, the nominal states corresponding to each IMU are propagated with the strapdown INS algorithm [179]. The linearized continuous-time error dynamics of each IMU follow the strapdown error model of Section 2.1.7, Eq. 2.17, evaluated per IMU: the direction cosine matrix is $\mathbf{C}(q_n^{b_i})$, corresponding to the estimated attitude quaternion of IMU i , and the specific force is $\vec{f}_i^n$, the bias-compensated specific force of IMU i expressed in the navigation frame. This yields a 15×15 system matrix $\mathbf{A}_i(t)$ for each IMU [171]. As in Section 2.1.7, Earth rotation and transport rate are neglected owing to the limited duration and spatial extent of pedestrian navigation. The corresponding process-noise matrix is

$$\mathbf{Q}_i = \text{diag}(\sigma_{\text{ARW}}^2 \mathbf{I}_{3 \times 3}, \sigma_{\text{VRW}}^2 \mathbf{I}_{3 \times 3}, \mathbf{O}_{3 \times 3}, \sigma_{\text{RRW}}^2 \mathbf{I}_{3 \times 3}, \sigma_{\text{AcRW}}^2 \mathbf{I}_{3 \times 3}), \quad (8.3)$$

where the noise terms are those defined in Section 2.1.7. Values are listed in Table 8.1 and were derived from the manufacturer specification of the VectorNav VN-200.

Because the two IMUs observe independent inertial physics, the joint state-transition and process-noise matrices are block-diagonal,

$$\mathbf{A}(t) = \begin{bmatrix} \mathbf{A}_s(t) & \mathbf{O}_{15 \times 15} \\ \mathbf{O}_{15 \times 15} & \mathbf{A}_k(t) \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \mathbf{Q}_s & \mathbf{O}_{15 \times 15} \\ \mathbf{O}_{15 \times 15} & \mathbf{Q}_k \end{bmatrix}. \quad (8.4)$$

Table 8.1: Process-noise parameters for the EKF, based on an 800 Hz IMU update rate.

Hyper-parameter	Value
σ_{ARW}	2.1597×10^{-5}
σ_{VRW}	4.8557×10^{-4}
σ_{RRW}	1.7141×10^{-6}
σ_{AcRW}	1.3873×10^{-5}

8.2.1 Filter Initialization

The initial state covariance $\mathbf{P}_0$ is set based on a brief static calibration period of approximately 10 seconds, during which the IMUs are assumed stationary. This period allows estimation of the IMU biases, which are initialized from the mean of the stationary measurements. Because the two IMUs were sampled on independent clocks, their streams were first time-aligned by a constant timing-offset estimate before fusion, as detailed in Sec. 8.2.1. To align the initial heading between the IMUs, a yaw alignment procedure is performed using a fixed-lag smoothing approach, detailed in Sec. 8.2.1. This smoothing step is acknowledged as a limitation of the current method; alternative approaches such as extended calibration routines, initial calibration runs, GPS integration, visual aiding [235], ultrasonic ranging [72], or other sensor fusion techniques could be employed to synchronize and fuse the sensors without requiring this smoothing process.

Temporal Alignment of the Dual-IMU Dataset

The shoe and knee IMUs were recorded by independent acquisition units and were therefore not sampled on a common clock. A constant timing offset between the two streams was estimated and removed prior to gait classification and EKF fusion, since an uncorrected offset degrades both the classifier features and the cross-IMU ZPD constraint.

The alignment exploited the anti-phase pitch relationship of the lower leg during walking: the knee-IMU pitch angle reaches its maximum at the same point in the gait cycle at which the shoe-IMU pitch angle reaches its minimum [236]. A pitch-signal reference of this kind was also used by Shabani et al. [237] to synchronize foot- and shank-mounted IMUs. For each trial, the pitch angle of both IMUs was computed, and the timing offset was taken as the time difference between a matched knee-pitch maximum and the corresponding shoe-pitch minimum. This offset was then applied to resample the shoe stream onto the knee time base. Fig. 8.2(top) illustrates the matched landmark across the recorded gait cycles.

The two streams were assumed to differ by a constant lag with negligible relative clock drift over the trial. This assumption is reasonable for the short trajectories used here, but would not hold for extended recordings, for which a drift-compensating alignment would be needed [237].

Mounting-Yaw Alignment via Fixed-Lag Smoothing

The knee IMU was mounted with an unknown, approximately constant yaw offset $\Delta\psi$ relative to the shoe IMU frame. The cross-IMU ZPD constraint requires at least an approximate correction of this offset at initialization; otherwise the two IMUs are fused in inconsistent heading frames. Because $\Delta\psi$ is unobservable while the subject is stationary, the constraint was held off until it became observable from the first translation, and the initial segment was recovered with a fixed-lag smoother.




Figure 8.2: Temporal alignment of shoe (blue) and knee (orange) IMU data, shown as gyroscope-derived change in pitch angle $\Delta\theta$ for walking data. (Top) $\Delta\theta$ over 11 gait cycles; the alignment landmark, circled, is where the knee-pitch maximum coincides with the shoe-pitch minimum. (Bottom) The 11 cycles superimposed, each centered on the knee-pitch peak, $t = 0$.

Let t_0 be the first navigation epoch. The heading ψ_m^i of IMU $i \in \{s, k\}$ was taken as the direction of its horizontal displacement from t_0 . To avoid estimating this bearing from position noise, the estimation window $\mathcal{S}$ was the first complete stance-phase for the shoe. Following [75], the offset was the circular mean of the yaw difference $\psi_m^s - \psi_m^k$ over $\mathcal{S}$,

$$\Delta\psi = \arctan\left(\frac{\sum_{m \in \mathcal{S}} \sin(\psi_m^s - \psi_m^k)}{\sum_{m \in \mathcal{S}} \cos(\psi_m^s - \psi_m^k)}\right). \quad (8.5)$$

In the multi-gait dataset, walking preceded crawling, so the shoe served as the anchor and the knee was aligned to it. If crawling is performed first, the roles reverse.

Let t_a be the epoch at which (8.5) first yields a valid estimate. Over $[t_0, t_a]$ a provisional forward filter (per-IMU ZUPT only, cross-IMU ZPD constraint off) produced the solution used in (8.5). The knee was then re-initialized at t_0 with $\Delta\psi$ embedded in its initial attitude, its position and covariance reset accordingly, and the buffer reprocessed over $[t_0, t_a]$ with the cross-IMU ZPD constraint enabled. For $t > t_a$ the filter continued as an ordinary causal recursion. The procedure is a fixed-lag smoother of lag $L = t_a - t_0$, bounded by the time of

one stride and independent of trial length; only the initial L seconds of the knee trajectory are revised, with $L = 4.9$ s on average for the nine trials.

The accuracy demanded of Eq. (8.5) is set by the sensitivity of the navigation solution to $\Delta\psi$, which is measured directly in Fig. 8.3. The navigation solution was run for varying knee yaw offsets displaced from the calibrated value by $0, \pm 3, \pm 6, \pm 12$, and $\pm 24^\circ$, with calibration, gait classification, ZUPT, and ZPD otherwise unchanged. The shoe trajectory is reported: the perturbed quantity belongs to the knee and reaches the shoe solution only through the ZPD constraint, so the figure isolates the propagation of a mounting-calibration error from one IMU into the navigation solution of the other. The Easting error at maximum displacement, 42.6 m, grows linearly with the yaw offset at 0.215 m per degree; a $\pm 24^\circ$ offset yields ± 5.1 m, or 12 % of maximum displacement, and holding this error below 1 m requires $\Delta\psi$ initialization accurate to 4.7° .




Figure 8.3: Sensitivity of the shoe trajectory to error in the knee mounting-yaw offset $\Delta\psi$. Each curve is a complete estimator run, identical except for $\Delta\psi$, displaced from its calibrated value by $0, \pm 3, \pm 6, \pm 12$, and $\pm 24^\circ$. Trajectories are rotated to a common heading at 10 m, which removes any rigid heading offset; circles mark maximum displacement.

8.2.2 ZUPT Measurement Model

The zero-velocity update developed in Section 2.1.8 is applied to each IMU independently whenever its own stance-phase detector declares ground contact, using the detectors and window sizes of Section 7.2. Because the velocity error is a sub-block of each per-IMU state, the measurement matrix of Section 2.1.8 is embedded in the joint state as a single 3×30 row block, occupying columns 1–15 for the shoe or 16–30 for the knee and leaving the opposing IMU unconstrained by that update. The measurement noise is the isotropic covariance $\mathbf{R}_i^{\text{ZUPT}} = \sigma_{\text{ZUPT}}^2 \mathbf{I}_{3 \times 3}$ for $i \in \{s, k\}$, with σ_{ZUPT} configured per IMU as described in Section 7.2 and per gait by the classifier of Section 8.4.

8.2.3 Zero-Position-Difference Constraint

The ZPD constraint [73] enforces the rigid-link condition that the ankle position estimated from the shoe IMU must agree with the ankle position estimated from the knee IMU. With the body-to-navigation direction cosine matrix $\mathbf{C}(q_n^{b_i})$ and the body-frame lever arm $\vec{L}^i$ from each IMU to the ankle, the two ankle estimates are:

$$\begin{cases} \vec{p}_{A_s} = \vec{p}_s + \mathbf{C}(q_n^{b_s}) \vec{L}^s \\ \vec{p}_{A_k} = \vec{p}_k + \mathbf{C}(q_n^{b_k}) \vec{L}^k \end{cases} \quad (8.6)$$

and the residual is their difference:

$$\vec{z}_{\text{ZPD}} = \vec{p}_{A_s} - \vec{p}_{A_k}. \quad (8.7)$$

Linearizing through the lever arm with the attitude-error convention $\delta(\mathbf{C}\vec{L}) = -\mathbf{C}[\vec{L} \times] \delta\vec{\theta}$ shows that each IMU's contribution affects only its own attitude and position columns of the joint 30-state error vector. The constraint is applied symmetrically: both IMUs are corrected simultaneously by a single 3-dimensional residual, distributing the correction between them

according to their relative uncertainties in the joint covariance. The measurement matrix is therefore a single 3×30 row block,

$$\begin{aligned} \mathbf{H}^{\text{ZPD}} = & \left[-\mathbf{C}(q_n^{b_s}) [\vec{L}^s \times] \quad \mathbf{O}_{3 \times 3} \quad + \mathbf{I}_{3 \times 3} \quad \mathbf{O}_{3 \times 6} \right. \\ & \left. + \mathbf{C}(q_n^{b_k}) [\vec{L}^k \times] \quad \mathbf{O}_{3 \times 3} \quad - \mathbf{I}_{3 \times 3} \quad \mathbf{O}_{3 \times 6} \right] \in \mathbb{R}^{3 \times 30}, \end{aligned} \quad (8.8)$$

with the shoe-mounted IMU occupying columns 1–15 and the knee-mounted IMU occupying columns 16–30. The constraint is activated whenever a valid ankle-position estimate is available from both IMUs.

This symmetric treatment relies on the joint EKF covariance to allocate the correction: at each update the Kalman gain places more of the innovation onto whichever IMU currently carries the larger position and attitude uncertainty. During walking, the shoe-mounted IMU is repeatedly re-anchored by ZUPT and its covariance stays tight, so most of the ZPD correction is absorbed by the knee-mounted IMU; during crawling, the role is reversed and the shoe-mounted IMU takes the majority of the correction.

The measurement noise encodes the ankle-model error (lever-arm calibration error, soft-tissue artifact, and small violations of the rigid-link assumption) and is taken as an isotropic covariance scaled by a per-gait standard deviation, $\mathbf{R}^{\text{ZPD}} = \sigma_{\text{ZPD}}^2 \mathbf{I}_{3 \times 3}$, with σ_{ZPD} a per-gait residual standard deviation. A single covariance suffices regardless of which IMU is better anchored, since the joint covariance rather than the measurement noise allocates the correction. The standard deviation is set looser during the walking and crawling gaits than when stationary to account for the greater soft-tissue artifact and knee-pad slip in those modes, and looser still during transitions to absorb the instants when posture, and thus lever-arm orientation, changes fastest. The per-gait values are listed in Table 8.2.

Table 8.2: Per-gait ZPD position uncertainty (one standard deviation). The ZPD constraint is symmetric, so a single isotropic covariance $\mathbf{R} = \sigma^2 \mathbf{I}_{3 \times 3}$ is applied, with σ the per-gait value listed.

Class	Gait	σ [cm]
0	Stationary	4.2
1–3	Walk, hand/knee / duck crawl	10.4
4–7	Transitions	300

8.3 Multi-Gait Dataset

The framework was evaluated on a multi-gait dataset that combined all gait methods in a single, continuous trajectory, exercising the walk-to-crawl transitions central to firefighter movement. The dataset was intended to mimic a real-world smoke-diving scenario, in which a firefighter alternates between gaits depending on the environment being traversed. Three subjects, each of approximately 178 cm height and 77 kg mass, performed the protocol. As in Section 7.2, the subjects were not trained firefighters but replicated the walking and crawling motions as closely as possible. Each wore the firefighter PPE described in Section 7.2, approximately 20 kg, with IMUs mounted on the shoe and the center of the knee pad.

Each trial began with the subject standing still for approximately 10 s to initialize the IMU biases, Section 8.2.1. The subject then walked in a straight line for approximately 21 m, transitioned down into hand/knee crawling for a further 21 m, stood and turned around to walk a third 21 m segment, and finally transitioned into duck crawling for the remaining 21 m back to the start. Each trial ended with the subject standing and precisely aligned to the starting position, so that the start and end coincide, providing a known endpoint for evaluating final-position error. The trajectory had a total path length of 84.6 m and an average duration of 189 s per trial, and three trials were collected per subject, for nine trials in total.

8.4 Gait Classification

The ZUPT stance-phase detection threshold and measurement-noise parameter (stance-phase residual velocity RMS) are configured for each gait, so the filter requires an online estimate of the current motion class, produced for each sample from a preceding fixed-length window. A random forest classifier was therefore developed to label each window of inertial data as one of eight motion classes: the four gait classes of the multi-gait dataset (Section 8.3) and four directional transitions between them. The transitions are merged into a single transition class for reporting, giving the five classes of Table 8.3 and Fig. 8.4; a window whose true and predicted labels are both transitions is scored as correct. The predicted class selects the per-IMU ZUPT detector thresholds and the per-gait measurement-noise parameters of Section 8.2.

This approach follows the framework of Sangenis et al. [66], who classified pedestrian activities from a single foot-mounted IMU to adapt the ZUPT detector threshold and zero-velocity measurement variance to each activity. Each classification instance is a window of $W = 100$ samples, 0.125 s at the 800 Hz IMU rate, of the three-axis accelerometer and gyroscope signals and the pitch angle of each IMU, extracted at a stride of one sample as a trailing window, giving $C = 14$ channels for the two-IMU configuration and $C = 7$ for a single IMU. The channel samples within a window are concatenated into a raw vector of dimension $D = C \times W$, 1400 and 700, respectively. The classifier was trained and evaluated on data collected on separate days. Training used four single-gait recordings per subject, 42.6 m each, which supplied the four sustained-gait classes, together with a dedicated transition recording that supplied the four directional transition classes. Evaluation used the multi-gait datasets, the same continuous multi-gait trajectories used for the navigation evaluation in this chapter (Section 8.3), collected on a different day and held out entirely from training. All balancing and dimensionality reduction were fit on the training data alone, so no information from the multi-gait evaluation data influenced the model. For class balancing on the training

split, the four gait classes were randomly undersampled toward 20 000 windows each, and the four directional transition classes, drawn from the dedicated transition recording, were oversampled to the same target with the SMOTE [238]. The PCA was then fit on the balanced training windows, reducing the representation to 15 components that retain 96.2 % of the total variance, and the identical projection was applied to the evaluation recordings. A random forest of 300 trees was selected as the final classifier.

To assess generalization, the classifier was trained on each subject individually and on all three combined, and each model was evaluated on every subject’s test set, for three sensor configurations: both IMUs, the shoe IMU alone, and the knee IMU alone. All testing used a separate multi-gait dataset recorded on a different day. In Tables 8.3a–8.3c, a mark indicates the subject(s) used for training and each accuracy column reports the test accuracy for that subject; single-mark rows therefore measure cross-subject generalization, while the all-subject (top) row measures cross-session performance on subjects also seen in training. The dual-IMU classifier trained on all three subjects reached a mean test accuracy of 88.6 % over the three subjects, 85.6 %, 87.0 %, and 93.1 %, exceeding both single-IMU configurations trained the same way. For a single IMU, models trained on one subject generalized poorly to the others, reflecting intersubject variability in walking and crawling dynamics, whereas training on all three subjects gave consistent cross-subject accuracy. The confusion matrix of Fig. 8.4 shows this accuracy concentrated in the sustained stationary, walking, and crawling classes, with the residual error dominated by transitions, which were most often absorbed into the adjacent walking or crawling class. As analyzed in Section 8.5, this transition confusion has little effect on the navigation solution, because the cross-IMU constraint is suppressed during transitions regardless of the predicted label.

Table 8.3: Random forest gait-classification accuracy (%) on the held-out multi-gait dataset by training-subject composition and sensor configuration.

(a) Classifier trained and evaluated on combined shoe- and knee-IMU features.

Subject Trained on			Accuracy		
Subj I	Subj II	Subj III	Subj I	Subj II	Subj III
x	x	x	85.6%	87.0%	93.1%
x			85.1%	76.2%	93.1%
	x		77.9%	86.2%	85.2%
		x	77.3%	70.7%	86.6%

(b) Classifier trained and evaluated on shoe-IMU features only.

Subject Trained on			Accuracy		
Subj I	Subj II	Subj III	Subj I	Subj II	Subj III
x	x	x	84.7%	81.9%	89.9%
x			84.5%	62.0%	81.2%
	x		75.2%	82.4%	83.1%
		x	72.1%	59.9%	74.5%

(c) Classifier trained and evaluated on knee-IMU features only.

Subject Trained on			Accuracy		
Subj I	Subj II	Subj III	Subj I	Subj II	Subj III
x	x	x	83.9%	86.7%	87.0%
x			82.3%	81.5%	86.0%
	x		69.9%	82.3%	74.4%
		x	79.6%	72.8%	80.9%




Figure 8.4: Confusion matrix for the random forest gait classifier, trained on all three subjects and applied to Subjects I, II, and III, showing accuracy above 86.4% except for transitions between walking/crawling.

8.5 Navigation Results and Discussion

To isolate the contributions of the classifier and the cross-IMU ZPD constraint, three filter configurations of increasing capability were compared. The baseline used the shoe IMU alone, since firefighter navigation systems are conventionally foot-mounted; during crawling the shoe still enters a stance-phase, though a degraded one relative to the knee. Two shoe-only configurations were defined: Independent (ZUPT), which applied a constant detector threshold γ , and Independent (ZUPT-ML), which adapted γ to the classified gait. The third configuration, Joint-EKF, additionally fused the shoe and knee IMUs through the ZPD constraint, with the per-gait configuration of Section 8.2. These configurations were evaluated on the nine multi-gait trials of Section 8.3. Because each trial returned to and aligned with its starting point, the final-position error equals the loop-closure error, reported as horizontal CEP and vertical RMSE in Table 8.4 and Figs. 8.6, 8.7, and 8.8.




Figure 8.5: Stance-phase detection statistic (log-likelihood) and the resulting gait-classifier output for the shoe IMU, (a), and the knee IMU, (b). Each IMU’s statistic is most stable in its anchor gait, the shoe during walking and the knee during crawling. The colored bar at right encodes the classified state, and the green labels above (a) mark the true gait, with the small gaps corresponding to the walk-to-crawl transitions.

The Joint-EKF reduced the horizontal CEP to 0.58m and the vertical RMSE to 1.26m, corresponding to 0.69 % and 1.49 % of the 84.6m path length. Relative to the shoe-only baseline, 3.86 m and 3.06 m, these are reductions of 85 % in CEP and 59 % in vertical RMSE, a factor of 6.7 and 2.4, respectively. The two aiding steps contributed differently. Adapting the threshold to the classified gait, Independent (ZUPT-ML), reduced CEP from 3.86 m to 2.30 m




Figure 8.6: Multi-gait North-East trajectories over the nine trials under the three filter configurations described in the text; black stars mark the true trajectory endpoint.




Figure 8.7: Multi-gait vertical trajectories versus time over the nine trials for each of the three filter configurations.




Figure 8.8: Final horizontal position error over the nine multi-gait trials for each of the three filter configurations; red markers are the per-trial trajectory endpoints, the black star is the true endpoint, and the red circle indicates the CEP, summarized in Table 8.4.




Figure 8.9: Easting error accumulation across the multi-gait protocol for all three navigation configurations. Each trajectory is heading-aligned so the corridor lies along the north axis; the plotted quantity is the median absolute cross-track error over the nine trials at five checkpoints: start, the three labeled gait transitions (walk to crawl, crawl to walk, walk to duck), and the endpoint. The shaded band shows the interquartile range of the corresponding value.

and vertical RMSE from 3.06 m to 1.29 m, recovering most of the vertical drift on its own. Adding the cross-IMU constraint, Joint-EKF, reduced CEP further to 0.58 m, delivering the majority of the remaining horizontal improvement, while vertical RMSE improved modestly to 1.26 m.

These results follow from the stance-phase behavior in Fig. 8.5(a),(b). The shoe detection statistic was stable during walking but degraded during crawling, and the knee statistic showed the opposite pattern, so a single constant threshold either missed crawl stances or admitted noisy ones on the shoe. Adapting γ per gait retimed the shoe ZUPT to its degraded crawl stance, which accounts for the large vertical-drift reduction, since mistimed zero-velocity updates particularly corrupt the vertical channel. The residual horizontal error persisted because the shoe stance is intrinsically weak during crawling; the Joint-EKF addressed this through the symmetric cross-IMU ZPD constraint: during crawling the knee IMU carries the tighter covariance, so the ZPD constraint directs most of the ankle-position correction onto the shoe, drawing it toward the better-anchored knee solution. The effect is visible in Fig. 8.6, where the shoe-only trajectories spread in both heading and length while the Joint-EKF trajectories cluster near the true path, and in Fig. 8.7, where the shoe-only vertical drift reached tens of meters against roughly ± 2 m for the Joint-EKF.

The checkpoint-resolved cross-track error in Fig. 8.9 confirms this across the trajectory: the Joint-EKF holds the lowest median error and the tightest interquartile band at every checkpoint, and reaches the endpoint at 0.58 m against 1.0 m for Independent (ZUPT-ML) and 2.75 m for Independent (ZUPT). The shoe-only configurations partially cancel opposing-sign drift between the crawl-to-walk and walk-to-duck segments, so their intermediate checkpoints understate the accumulated error, whereas the Joint-EKF remains the lowest at every checkpoint without relying on such cancellation.

Several factors bounded the residual accuracy. The classifier confusion matrix, Fig. 8.4, reports above 86.4% accuracy for the stationary, walking, and crawling states but only

Table 8.4: Final-position error over the nine multi-gait trials for each filter configuration.

Configuration	CEP [m]	Vertical RMSE [m]
Independent (ZUPT)	3.86	3.06
Independent (ZUPT-ML)	2.30	1.29
Joint-EKF	0.58	1.26

19.3 % for transitions, which were most often mislabeled as walking, 29.8 %, or as one of the crawling gaits, 24.5 % and 23.6 %. Transitions are brief relative to the sustained walking and crawling segments, so they comprise only a small fraction of each trajectory, and the overall per-sample accuracy therefore remained above 85 %, Section 8.4. Their misclassification also had limited effect on the navigation solution: during transitions the ZUPT thresholds for both IMUs were set low, so that only an exceptionally stable stance triggered an update, and the ZPD covariance was inflated so that the cross-IMU constraint was effectively disabled. Because a transition is inherently non-stationary, no update or constraint would be applied regardless of the label, so a mislabeled transition is largely benign. The dominant transition error, marked Error Type I in Fig. 8.5(a),(b), arose from the large pitch excursions during the toe-down phase of walking, a high-motion interval that resembles the raising and lowering of a true transition. As both are non-stationary, this confusion did not degrade the solution.

The second-largest error, marked Error Type II in Fig. 8.5(a),(b), was confusion between hand/knee and duck crawling, 12.5 % and 6.5 % in Fig. 8.4. These errors occurred mainly during the dynamic phase rather than during stance. Even within a stance-phase, the two crawling gaits share similar detector thresholds γ , so ZUPT still fired correctly, and the misclassified intervals were brief enough that at least part of each stance was still captured.

Three limitations bound the scope of these results and motivate further study. First, the reported trajectories were produced offline: the shoe and knee streams were time-aligned by a post-hoc constant-offset estimate, Section 8.2.1, and the initial heading was recovered with a fixed-lag smoother, so a causal, real-time implementation remains future work. Second, the evaluation used three subjects over nine trials, and the subjects were not trained firefighters

but replicated the walking and crawling motions in full PPE; operational fireground data would be needed to confirm performance under realistic conditions. Third, a single leg was instrumented, so the crawl-segment analysis assumes the instrumented knee provides the stance anchor throughout; multi-leg instrumentation would remove this assumption.

8.6 Conclusion

In this chapter, ZUPT-aided pedestrian inertial navigation was extended across the walking and crawling gaits used during firefighter smoke-diving operations, for which no single mounting position provides a reliable stance-phase. Because the shoe is anchored during walking and the knee during crawling, the shoe- and knee-mounted solutions were fused in a joint 30-state error-state EKF coupled through the ZPD constraint on the rigid lower leg. The constraint was applied symmetrically, so that the joint covariance directs the correction onto whichever IMU is less well anchored in the current gait. A random forest gait classifier configured the per-IMU stance-phase detector and the per-gait measurement noise online, reaching a mean classification accuracy of 88.6% across the three subjects. Evaluated over nine multi-gait trials with three subjects in full firefighter PPE, the joint solution reduced the horizontal CEP from 3.86 m to 0.58 m and the vertical RMSE from 3.06 m to 1.26 m relative to a conventional foot-mounted solution.

ZUPT bounds the growth of inertial error only if there is a reliable stance-phase to estimate IMU biases and errors; without it, the aiding degrades and error accumulates unchecked. The stationary interval is recovered here by drawing the zero-velocity reference from whichever body segment is anchored in the gait being performed, without ranging hardware or prior infrastructure, so that much of the accuracy a pedestrian INS attains while walking is carried across the gaits firefighters use during smoke-diving operations. This directly addresses the extension of ZUPT-aided pedestrian navigation from conventional walking to firefighter crawling and to the transitions between gait modes: a single wearable configuration, reconfig-

ured online by the gait classifier, sustains bounded navigation error across both gaits and the transitions between them. What remains is the residual heading drift that no zero-velocity measurement can observe, which the sensor developed in Part I is intended to constrain.

RQ3b: The navigation accuracy that ZUPT provides during walking is maintained across crawling by mounting an IMU on both the shoe and the knee and fusing them in a joint EKF coupled through the ZPD constraint on the rigid lower leg, so that the segment anchored in the current gait constrains the one that is not, with a gait classifier configuring the per-IMU detector and measurement noise.

Chapter 9

Conclusions

This dissertation reports on the development of fused silica DSG and on the navigation algorithms in which these sensors are employed, toward the goal of self-contained pedestrian localization in environments where external aiding signals are unavailable. On the sensor side, two approaches were investigated. In the first, metallization is retained and the capacitive electrodes are defined directly within the bulk of the fused silica by FLICE, forming the IPE-DSG and eliminating the separately assembled electrode substrate. In the second, metallization of the resonating element is removed entirely, and driving, tuning, and sensing are instead performed by IPA of the outer shell, laser ablation trimming, and optical detection. On the algorithm side, a meso-scale inertial navigation platform was developed, and the ZUPT-aided pedestrian inertial navigation framework was extended to the non-standard gaits encountered by firefighters during smoke-diving operations, where the walking gait assumption underlying conventional foot-mounted implementations no longer holds. Specific contributions of this work are summarized below.

9.1 Contributions

9.1.1 Fused Silica Dual-Shell Gyroscope Contributions

- A method of digitally defining self-aligned, high-aspect-ratio capacitive electrodes directly within the bulk of fused silica using FLICE was developed, eliminating both the separately fabricated electrode substrate and the precise assembly step conventionally required to form the capacitive gap. An experimental parameter study identified the laser pulse energy and KOH etching duration that minimize channel width while achieving complete release, yielding capacitive gaps as small as $2\text{ }\mu\text{m}$ with aspect ratios as high as 150:1. Applied to a DSG to realize the IPE-DSG, the process yielded a maximum as-fabricated Q-factor of 4.57 M for the $n = 2$ wineglass mode, indicating that FLICE patterning does not substantially increase surface-related energy loss. Because the electrodes are formed within the same monolith as the resonator, the capacitance non-uniformity introduced by assembly tilt and misalignment is eliminated. Open-loop rate gyroscope operation was demonstrated with a scale factor of $0.07\text{ mV}/(\text{deg/s})$, and a projection based on the optimized FLICE parameters indicates a potential $14\times$ improvement in scale factor. Because FLICE can pattern electrodes in any preserved region of an optically transparent dielectric, the approach is applicable to other fused silica resonator geometries beyond the DSG.
- The mechanism governing IPA transduction efficiency between the outer and inner shells of the DSG was identified, and a design framework for maximizing it was established. The dimensionless γ ratio, defined as the distance from the central stem to the shell junction divided by the length of the inner shell, was identified as the dominant geometric parameter, with predicted displacement maximized at an intermediate value of γ . Ten DSG designs spanning a range of γ were evaluated by adapting the RAFT to treat the DSG as a lumped mass-spring-damper system. Experimental

characterization of five PZT-coated DSGs agreed with the RAFT predictions, demonstrating a $6.5\times$ improvement in displacement amplitude and a $12.3\times$ improvement in Q-normalized transduction efficiency over non-optimized designs. Two distinct design paths to large displacement amplitude were identified: a low- γ path achieving 9.1 nm at 200 kV cm^{-1} through a Q-factor of 1.06 million, and a high- γ path achieving comparable displacement at a reduced Q-factor of 278k. The low- γ path is preferred for gyroscope applications, where high Q-factor improves rate sensitivity.

- Non-metallized tuning of the DSG was demonstrated by laser ablation trimming, since electrostatic frequency tuning is no longer possible without metallization. A semi-automated procedure was developed to identify the principal axes of elasticity by characterizing the $n = 2$ mode shape around the rim of the resonator, combining automated stage motion, LDV, and lock-in amplifier control. Material was then removed at the antinodes of the lower-frequency mode. Two trimming approaches were developed. Direct-rim trimming reduced the frequency split to as low as 0.39 Hz, but degraded the Q-factor, by 73 % on average, since material is ablated into the resonating shell itself. Defining teeth-like-tines on the rim by FLICE mitigates Q-factor reduction by relocating the ablation damage off the resonating shell, so that trimming perturbs the modal mass with little change to the stiffness of the vibrating shell. Across 17 shells the teeth-like-tines approach reached a split as low as 0.10 Hz while limiting the average Q-factor loss to 21 %, achieving an average Q-factor of 1.4 M after trimming.
- Non-metallized gyroscope operation of a microfabricated HRG was demonstrated, to the best of our knowledge for the first time, with driving, tuning, and sensing all performed without metallization of the resonating element. A custom dual LDV setup was developed that positions and tilt-corrects two vibrometers on a precision rate table to probe the drive and sense antinodes simultaneously, serving as a benchtop prototype of a MZI implemented in a PIC. IPA was used to selectively excite the

$n = 2$ wineglass mode, and open-loop rate operation was demonstrated with a scale factor of $0.02 \text{ mV}/(\text{deg/s})$. Because feedthrough capacitance is absent in the non-metallized configuration, electromechanical amplitude modulation was neither required nor possible.

9.1.2 Self-Contained Pedestrian Navigation Contributions

- A meso-scale inertial navigation platform was developed for pedestrian navigation, with a total volume of 7.1 cm^3 excluding battery and case, small enough for integration into the sole of a shoe. A barometer-aided ZUPT-aided INS was implemented in an EKF framework with sufficient computational headroom to retain accelerometer and gyroscope biases in the filter states, which prior embedded implementations had omitted in order to maintain update rate. Ten indoor walking experiments including flat planes, stairs, and an elevator ride demonstrated a CEP of 0.66 m over a 125 m trajectory, or 0.53% of the trajectory length.
- The parameters governing stance-phase detection during crawling were characterized through a search over detector type, window size, and detection threshold for heel- and knee-mounted IMUs. Under the experimental conditions, the knee-mounted IMU rendered a 77.3% improvement in CEP over the heel-mounted IMU across a 42.6 m straight-line trajectory. The detector configuration conventionally used for walking was shown to be insufficient for reliable stance-phase detection during crawling, establishing that detector parameters must be selected per gait.
- The optimal IMU mounting position for firefighter crawling was determined by evaluating four candidate placements across two crawling methods, with the residual stance-phase velocity measured directly using an electromagnetic position sensor to quantify the velocity uncertainty of each placement. The center of the knee pad was identified as the placement with the longest stance-phase duration and the lowest velocity uncer-

tainty, rendering a $1.9\times$ and $3.7\times$ improvement in CEP over the remaining mounting positions for hand/knee and duck crawling, respectively.

- A dual-IMU ZUPT-aided INS framework was developed that provides a reliable stance-phase across both walking and crawling without additional ranging hardware, addressing the fact that no single mounting position serves both gaits. Shoe- and knee-mounted solutions were fused in a joint 30-state error-state EKF that shares a block-structured covariance between the two per-IMU strapdown solutions and transfers information across them through simultaneous ZUPT and ZPD measurement updates derived from the rigid lower-leg kinematic constraint. The position-difference constraint was applied symmetrically, allowing the joint covariance to allocate the correction to whichever IMU carries the larger uncertainty in the current gait. A random forest gait classifier was developed to identify the current motion state and to configure the stance-phase detector threshold and measurement noise for each IMU and gait. Evaluated on three subjects in full firefighter equipment over an 84.6 m, 189 s multi-gait trajectory comprising nine trials, the framework reduced the horizontal CEP from 3.86 m to 0.58 m and the vertical RMSE from 3.06 m to 1.26 m relative to a conventional foot-mounted solution, corresponding to reductions of 85 % and 59 % and to 0.69 % and 1.49 % of the path length.

9.2 Answers to the Research Questions

The research questions posed in Section 1.6 are answered concisely below.

- **RQ1.** The transduction of the fused silica DSG was improved along two parallel paths, each removing one requirement of conventional electrostatic mechanization. The first eliminates the separately assembled electrode substrate while retaining electrostatic transduction; the second removes the metallization of the resonating element, replacing

the driving, tuning, and sensing that electrostatics provided, and the shell geometry governing the efficiency of that replacement drive was then optimized.

- **RQ1a.** Capacitive electrodes are defined without an assembly step by stopping back-lapping before release and patterning the gaps directly into the bulk fused silica with FLICE. The minimum gap and maximum aspect ratio are governed by the laser pulse energy and the KOH etching duration.
- **RQ1b.** Rate operation with no conductive film on the resonating element is achieved by replacing each function of metallization independently: IPA of the outer shell for directional driving, dual-LDV optical detection for sensing, and laser ablation trimming for permanent correction of the frequency split.
- **RQ1c.** The dual-shell geometry that maximizes IPA displacement amplitude is one of intermediate γ ratio, since transduction efficiency rises with γ while the Q-factor falls to anchor loss over the same range. Displacement, the product of the two, is therefore optimized by balancing these effects rather than by maximizing transduction efficiency alone.
- **RQ2.** Implementing pedestrian inertial navigation in a shoe-integrated platform imposes coupled computational, thermal, and mechanical constraints; a self-contained platform small enough for the sole of a shoe was built and evaluated to establish both the accuracy achievable within those constraints and the constraints themselves.
 - **RQ2a.** The platform attains meter-level accuracy, a CEP of 0.66 m, or 0.53 % of the 125 m trajectory.
 - **RQ2b.** Its integration is bounded by computation, which limits the filter states retainable at the update rate; by processor heat, which biases the IMU as it warms

toward steady state; and by mechanical loading of the sole during stance, which appears as residual lateral drift.

- **RQ3.** The ZUPT-aided navigation framework was extended from walking to the crawling gaits of firefighter operation, both by characterizing stance-phase detection for those gaits and by fusing two IMUs so that walking accuracy is maintained across both gaits.
 - **RQ3a.** During crawling the stance-phase is reliably detected by an accelerometer-based detector rather than one weighting angular rate, and the IMU is best mounted at the center of the knee pad, which yields the longest stance-phase duration and the lowest residual velocity among the placements evaluated.
 - **RQ3b.** Walking-grade accuracy is maintained across crawling by mounting an IMU on both the shoe and the knee and fusing them in a joint EKF coupled through the ZPD constraint on the rigid lower leg, so the segment anchored in the current gait constrains the one that is not, with a gait classifier configuring the per-IMU detector and measurement noise.

9.3 Projected Gyroscope Performance

The two transduction methods developed in Part I were directed at improving the noise performance of the wafer-level DSG. The ARW attainable by each is projected here within a common framework, so that the potential of each method for reaching navigation-grade performance can be assessed. Both devices are modeled as FRB gyroscopes, for which the Brownian noise floor is given by Eq. 2.32. This noise floor scales as $Q^{-1/2}$ and as A^{-1} , so it is improved both by increasing Q-factor and by increasing the displacement amplitude. For each method, the ARW set by the dominant electronic noise source is compared against this Brownian floor, Fig. 9.1, and the projected performance is taken as the larger of the two. These are projections rather than demonstrations since neither device was operated in

FRB mode of operation, and the assumed Q-factors and frequency splits were not achieved simultaneously.

The IPA method, Chapter 5, lowers the noise floor by increasing the displacement amplitude through the product of transduction efficiency and Q-factor. Because the drive amplitude increases linearly with Q-factor, evaluated here at Q-factors of 278k, 1 M, and 4.5 M, the Brownian-floor ARW scales as $Q^{-3/2}$ rather than the $Q^{-1/2}$ of a fixed-amplitude device. The displacement at the sense axis is optically detected, and two optical detection methods are considered against this floor, Fig. 9.1(a). Chip-scale photonic detection, which has already been demonstrated on a DSG [159], has a projected displacement noise floor of $0.83 \text{ pm}/\sqrt{\text{Hz}}$ on the bare fused silica shell, while benchtop LDV detection has a lower floor of $0.05 \text{ pm}/\sqrt{\text{Hz}}$ referred to the 46 kHz DSG 4n resonance [230]. Against both floors the projected ARW is detection-limited across the range of frequency splits considered, with the Brownian floor lying below the detection-limited ARW at each Q-factor. For chip-scale photonic detection, the projected ARW reaches the sub-degree-per-square-root-hour range once the Q-factor approaches 1 M and the frequency split is reduced below approximately 0.1 Hz. Benchtop LDV detection, having the lower displacement noise floor, approaches the Brownian floor at the highest Q-factor considered.

The IPE method, Chapter 3, instead lowers the electronic contribution to the ARW by increasing the sensing capacitance. This contribution scales inversely with the displacement-to-charge conversion coefficient of the capacitive pickoff, $K_{yq} = V_{\text{bias}} C_{sn}/g$, where V_{bias} is the bias voltage, C_{sn} the sense capacitance, and g the capacitive gap [9]. The capacitance can be raised by reducing the gap, but a smaller gap also reduces the drive amplitude sustainable before capacitive-gap nonlinearity, and therefore the signal. The high-capacitance geometry of Section 3.7 avoids this tradeoff by raising the capacitance at a fixed gap, preserving the drive amplitude. The electronic-noise contribution was estimated from the noise floor of the DSG reported in [9]. As shown in Fig. 9.1(b), the electronic-noise-limited ARW falls




Figure 9.1: Projected ARW of the two transduction methods versus frequency split. (a) IPA: detection-limited ARW for chip-scale PIC and benchtop LDV optical floors, with the mechanical-thermal (Brownian) floor in black; the drive amplitude is scaled with Q from the measured Q -normalized displacement. (b) IPE on Sample A: electronic-noise-limited ARW estimated from the noise model of [9], evaluated at the Q -factor values and capacitances shown, with the Brownian floor in black; the drive amplitude is set by the gap and the out-of-plane to in-plane displacement ratio.

well below the Brownian floor across the range considered, so the projected performance is limited by the Brownian floor. At mode-matched operation the projected ARW ranges from $4.5 \times 10^{-4} ^{\circ}/\sqrt{h}$ at a Q -factor of 1 M to $1.6 \times 10^{-4} ^{\circ}/\sqrt{h}$ at a Q -factor of 4.5 M.

Both methods are therefore projected to reach ARWs within the range required for pedestrian inertial navigation, through different dominant noise mechanisms: the IPA projection is limited by the optical readout, while the IPE projection is limited by the Brownian floor once the capacitance is sufficient to suppress the electronic noise below it.

9.4 Bridging Sensor Development and Pedestrian Navigation

The accuracy attainable by a ZUPT-aided INS is bounded by the performance of its yaw-axis gyroscope. Reaching that bound additionally requires a complete six-axis IMU and the real-time computation that fuses its measurements, both of which were developed in Part II. Every state of the navigation solution is observable for ZUPT-aided INS except the yaw angle and the bias of the gyroscope measuring it, which accumulate without correction, Section 1.5. The cross-track position uncertainty derived there, Eq. (1.2), is therefore governed by the noise of the yaw-axis gyroscope, and Fig. 1.3 shows this uncertainty falling by orders of magnitude as that gyroscope advances from consumer grade toward navigation-grade HRG performance. The sensor development of Part I was directed at supplying such a gyroscope with lower SWaP+C than the gyroscopes that meet those performance specifications today, as discussed in Section 1.5.1.

Two transduction methods for the wafer-level DSG were developed toward this requirement, and both were projected in Section 9.3 to reach ARWs within the range required for pedestrian inertial navigation. Both act on the same FRB noise floor of Eq. 2.32, and both were evaluated as candidate yaw-axis gyroscopes for the same ZUPT-aided pedestrian INS. This shared noise floor is the connection that unites the sensor development of Part I with the navigation development of Part II.

Both methods therefore represent promising approaches to achieving navigation-grade gyroscope performance in a wafer-level device, and both were shown to have the potential to reach this level. Neither, however, has demonstrated gyroscope operation at this level, and the projected performance was not verified experimentally. The navigation studies of Part II were conducted with commercial IMUs, so the benefit of the DSG to the navigation solution remains a projection.

9.5 Limitations

The contributions of this dissertation are subject to several limitations which are consolidated here.

The ARW was used throughout as the indicator of gyroscope noise performance, but it is not the only noise coefficient governing the navigation solution. As shown in Section 1.5, the cross-track position uncertainty of Eq. (1.2) depends on both the ARW and the RRW of the yaw-axis gyroscope, which enter as t^3 and t^5 respectively. The ARW therefore governs the cross-track error at the trajectory durations demonstrated in this dissertation, while the RRW governs at the longer durations of a sustained mission. Unlike the ARW, the RRW cannot be projected analytically from the design parameters of the resonator, and must instead be characterized experimentally on a fabricated and operating device. The projections of Section 9.3 therefore establish a path toward navigation-grade ARW, but do not bound the RRW that would govern performance over a full mission duration.

The gyroscope contributions are projections rather than demonstrations. The DSG was not operated as a closed-loop rate sensor in this work, and the Q-factors and frequency splits assumed in the projections were not achieved simultaneously on a single DSG.

The DSG was not incorporated into the Sugar-Cube platform or the firefighter navigation system, and the navigation studies of Part II were conducted with commercial IMUs. The benefit of the DSG to the navigation solution is therefore projected through the position-uncertainty model of Section 1.5. Realizing a self-contained ZUPT-aided INS built on the DSG requires development beyond the scope of this dissertation. The pickoff and control electronics must be miniaturized, as the characterization reported here relied on benchtop lock-in amplifiers. The device requires thermal characterization and correction, and characterization of its sensitivity to gravity, vibration, and humidity. Vacuum packaging, and integration with a platform that offers thermal isolation and shock-mitigation are also re-

quired. Each of these stands between the projected performance of the DSG and operation within a full INS.

9.6 Future Research Directions

9.6.1 Fused Silica Dual-Shell Gyroscope

Improvement of Annealing Profile After glassblowing, the DSG is removed from the furnace at 1700 °C and cooled to room temperature. The thermal shock of this transition induces significant residual stress, which must be relieved by annealing. In this work, a 920 °C, 30 hour anneal was demonstrated to substantially increase Q-factor; however, this profile may not be optimal. During annealing the glass structure relaxes toward the equilibrium configuration corresponding to the annealing temperature, so annealing at the lowest achievable temperature leaves the structure closest to equilibrium once returned to room temperature, minimizing residual stress and thereby reducing internal loss and improving Q-factor. The hydroxyl content of the glass sets a lower bound on how low this temperature can be. Hydroxyl groups are a consequence of the manufacture of synthetic fused silica and are distributed through the bulk; Corning 7980 fused silica contains an expected 800–1000 ppm [115]. Grades with lower hydroxyl content are available, but at considerably greater cost. Lunin *et al.* showed that annealing drives dehydroxylation of a near-surface zone: hydroxyl groups diffuse out of the surface region and escape into the atmosphere, which places the surface in tension. Below approximately 850 °C the structural relaxation rate is too slow to hold this stress, and fractures form on the surface [121, 206]. The 920 °C anneal used here was selected to remain above this bound. However, it is possible that this bound does not apply to glassblown DSG, since the near-surface hydroxyl distribution is established at 1700 °C during glassblowing rather than during the subsequent anneal, in which case no new dehydroxylation layer would form and a lower annealing temperature could be used. Verifying this requires a measurement of residual stress that is independent of Q-factor,

since Q-factor cannot differentiate the internal loss contribution from anchor and surface loss. Photoelastic imaging polarimetry, such as the StrainScope instruments (ilis GmbH, Erlangen, Germany), resolves stress birefringence in transparent glass and would serve this purpose.

Thickness-Profile Tailoring of the DSG. Two routes are available for tailoring the wall thickness profile of the DSG: patterning the device wafer prior to glassblowing, as demonstrated by Najafi’s group for blowtorch-molded shells [239], or post-forming removal on the inner shell using FLICE. Either provides a design variable independent of the γ ratio [100] and of the hemitoroidal topology [199] for reducing stress concentration at the junction between the inner and outer shells. A lower as-fabricated $n = 2$ frequency is desirable because the decay time scales as $\tau = Q/(\pi f)$ [89]; however, thickness reduction also reduces modal mass, so the profile must be optimized against thermomechanical noise rather than for frequency alone.

Anchor Stress Isolation. Anchor loss is the dominant remaining dissipation mechanism in the DSG, and it is governed by the geometry of the structure, in particular the γ ratio [6, 100]. Junction geometry is currently set by the layout of the etched cavities and is therefore constrained to what glassblowing can produce. However, FLICE offers a route to modify the inner shell region after glassblowing, adding slots that reflect acoustic energy back toward the outer rim of the inner shell and prevent its transmission into the anchor. Prior work cut isolation slots by FIB into polycrystalline diamond hemispherical resonators and measured no improvement in Q-factor; in those devices, however, anchor loss may not have been the limiting mechanism as Q-factor was measured at 87k [103]. Exploring this method for the DSG, which has a maximum reported Q-factor of over 4.5 M (as of Chapter 3 of this dissertation), may prove worthwhile.

In-Plane Electrode High-Capacitance Designs. The single-cut electrode demonstrated for the IPE-DSG is the simplest of many patternable geometries. Section 3.7 develops a pseudo-parallel-plate array that is expected to increase capacitance substantially over the single-cut design, though it has not been fabricated and its realization remains future work.

The primary challenge to realizing these high-capacitance designs is metallization. As established, conformal coverage of 150:1 aspect ratio structures is challenging. In this dissertation metal coating is deposited by PE-ALD; however, more conformal coverage could be achieved with thermal ALD of either Pt or aluminum-doped zinc oxide (AZO). The primary advantage of AZO over Pt is a lower CTE and less nucleation delay on fused silica [216]. A further benefit of the high-capacitance design is that concentrating additional mass at the rim increases the modal mass while preserving the low-anchor-loss geometry of the DSG, lowering the thermomechanical noise floor.

The IPE-DSG can also be designed to reduce the ratio of in-plane to out-of-plane displacement for the $n = 2$ wineglass mode, since this ratio is set by the shell geometry and decreases as the radius-to-height ratio is lowered by blowing a taller shell at a fixed radius [89]. Because the IPEs sense in-plane motion, which is also the motion limited by the capacitive-gap nonlinearity, a geometry that minimizes in-plane displacement while retaining large out-of-plane displacement is expected to preserve linear operation at smaller gaps while sustaining a strong drive signal, improving SNR. This mitigates the tradeoff between sensitivity and linearity that smaller gaps otherwise impose [211].

Simultaneous Trimming of $n = 2$ and $n = 3$ Modes. Simultaneous trimming of the $n = 2$ and $n = 3$ mode pairs would enable both pairs of a single DSG to be operated concurrently as independent CVGs. Ge *et al.* demonstrated this concept on a planar disk resonator, reducing the frequency split of both mode pairs to approximately 100 mHz via mass loading [240]. Because the two mode pairs exhibit different temperature sensitivities while their bias

drifts remain strongly correlated, the common disturbance becomes observable and can be canceled by fusing the two rate measurements, yielding a substantial reduction in long-term drift [241, 242].

Femtosecond Laser Induced Tuning of Anisodamping. The laser ablation trimming described in this dissertation has already been extended through FLICE-based trimming, which reduces material damage relative to ablation [243]. Both methods, however, address only the frequency split Δf . It is also highly desirable to correct anisodamping ($\Delta(1/\tau)$) as this causes angle drift in WA mode and increases quadrature bias drift in FRB mode [110]. Correcting anisodamping has previously been proposed through selective metal deposition, lowering the Q-factor of the higher-Q-factor mode until the two match [244].

Femtosecond laser irradiation may offer an alternative that does not require material removal/addition. The CTE of fused silica can be locally and permanently tuned by laser modification, and the sign of the change is set by the modification regime: type I modifications densify the exposed volume and lower its CTE toward $0.3 \times 10^{-6} \text{ K}^{-1}$, roughly half the unmodified value, while type II self-organized nanogratings expand and raise it by approximately the same factor [245]. The two regimes likewise leave the surrounding material in stress states of opposite sign [246]. This change in CTE will impact TED and could be used as a mechanism of anisodamping tuning. Note, this method of tuning would impact both anisodamping and anisostiffness as it was shown that the Young's modulus of fused silica will change and even become anisotropic depending on the polarity of type II laser modification [247]. If the effect can be established, a two-pass scheme could be used: a first pass trims $\Delta(1/\tau)$ using unetched nanogratings, and a second FLICE pass trims Δf . Provided the second pass can be made significantly $\Delta(1/\tau)$ -neutral [243], this would enable independent perturbation of the two asymmetries.

9.6.2 Self-Contained Pedestrian Navigation

Integration of the DSG with the Sugar Cube platform The ultimate vision is a die-level integration of a MEMS IMU with computational processing, termed the μ NavChip [44, 45]. The eventual goal is to integrate the DSG into the Sugar Cube platform, both to reduce yaw angle drift and to expose the challenges such an integration presents. The primary barrier is miniaturization of the DSG signal processing. The current benchtop setup relies on lock-in amplifiers and discrete pickoff and drive electronics, all of which must be reduced in size, and the control algorithms must be implemented on an FPGA [91].

Non-zero velocity update for the knee-mounted IMU As discussed in Section 7.1, the knee rotates during crawling as the body pivots about it. This is the primary reason the AMV detector outperformed the SHOE detector in stance-phase detection, and it implies that, although the stance-phase is more stable at the knee, the IMU is not perfectly stationary. This residual motion is not accounted for in the EKF, which assumes zero-mean, white measurement noise. Estimating the residual velocity of the knee pad during the stance-phase, following the approach of [180], could therefore further reduce navigation error for knee-mounted IMUs in firefighter smoke-diving scenarios.

Bibliography

- [1] Robert Lutwak. Micro-Technology for Positioning, Navigation, and Timing Towards PNT Everywhere and Always. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, pages 1–4, Laguna Beach, CA, USA, Feb. 2014.
- [2] Wen Zhang, Mounir Ghogho, and Baolun Yuan. Mathematical model and matlab simulation of strapdown inertial navigation system. *Modelling and Simulation in Engineering*, 2012, 2012.
- [3] Mohammad H Asadian, Danmeng Wang, and Andrei M Shkel. Fused Quartz Dual-Shell Resonator Gyroscope. *IEEE/ASME Journal of Microelectromechanical Systems*, 31(4):533–545, 2022.
- [4] Austin R. Parrish, Mohammad H. Asadian Ardakani, Lois Meira-Lopez, and Andrei M. Shkel. Laser Patterned Self-Aligned Electrodes for Hemispherical Resonator Gyroscope. U.S. Patent Application, 2025. UC Case No. 2025-819-1.
- [5] Jeffrey S. Pulskamp, Sarah S. Bedair, Ronald G. Polcawich, Gabriel L. Smith, Joel Martin, Brian Power, and Sunil A. Bhave. Electrode-shaping for the excitation and detection of permitted arbitrary modes in arbitrary geometries in piezoelectric resonators. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 59(5):1043–1060, 2012.
- [6] Lois Meira-Lopez, Austin R Parrish, Esther Wong, Danmeng Wang, and Andrei M Shkel. From Stress to Success: Demonstrating Ultrahigh Quality Factor Dual-Shell Resonator Gyroscopes with Low Stress Topology. In *Solid-State Sensors, Actuators and Microsystems Workshop (TRANSDUCERS)*, Orlando, FL, USA, Jun. 2025.
- [7] VectorNav Technologies, LLC. VN-200 GNSS/INS (Data Sheet). Datasheet, Rev. 2, 2022. Accessed: 2023.
- [8] Polhemus. Polhemus Patriot. <https://polhemus.com/>, 2008. Accessed: Jul. 2023.
- [9] Danmeng Wang and Andrei M. Shkel. Effect of EAM on Capacitive Detection of Motion in MEMS Vibratory Gyroscopes. *IEEE Sensors Journal*, 22(3):2271–2281, 2022.

- [10] Chun Yang, Thao Nguyen, and Erik Blasch. Mobile positioning via fusion of mixed signals of opportunity. In *IEEE Aerospace and Electronic Systems Magazine*, volume 29, St. Louis, MO, USA, 2014.
- [11] Fran van Diggelen and Per Enge. The World’s first GPS MOOC and Worldwide Laboratory using Smartphones. In *International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+)*, Tampa, FL, USA, Sep. 2015.
- [12] Paul Groves. *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Artech, Norwood, MA, USA, 2nd edition, 2013.
- [13] CloudNav, Inc. CloudNav. <https://cloudnav.net/>, 2014. Accessed: Jul. 2026.
- [14] SmartPoint. <https://www.smartpointco.com/>, 2017.
- [15] Dorota A. Grejner-Brzezinska, Charles K. Toth, Terry Moore, John F. Raquet, Mikel M. Miller, and Allison Kealy. Multisensor Navigation Systems: A Remedy for GNSS Vulnerabilities? *Proceedings of the IEEE*, 104(6):1339–1353, 2016.
- [16] Gregory Falco. When satellites attack: Satellite-to-satellite cyber attack, defense and resilience. In *ASCEND*, Virtual, Nov. 2020.
- [17] J. Rantakokko, P. Händel, M. Fredholm, and F. Marsten-Eklöf. User requirements for localization and tracking technology: A survey of mission-specific needs and constraints. *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2010.
- [18] Richard Campbell. United States Firefighter Injuries. National Fire Protection Association (NFPA), <https://www.nfpa.org/en/education-and-research/research/nfpa-research/fire-statistical-reports/firefighter-injuries-in-the-united-states>, 2025. Accessed: Jul. 2026.
- [19] Richard Campbell and Ben Evarts. United States Firefighter Injuries in 2020. Technical report, National Fire Protection Association, Dec. 2021.
- [20] Joseph W. Starr and B. Y. Lattimer. Evaluation of Navigation Sensors in Fire Smoke Environments. *Fire Technology*, 50:1459–1481, 2014.
- [21] Carl Fischer and Hans Gellersen. Location and navigation support for emergency responders: A survey. *IEEE Pervasive Computing*, 9(1):38–47, 2010.
- [22] Jerry Davis and Sean Gallagher. Physiological demand on firefighters crawling during a search exercise. *International Journal of Industrial Ergonomics*, 44(6):821–826, 2014.
- [23] Eudald Sangenis and Andrei M. Shkel. BUG-FIRE: A 3D Path Planning Algorithm with Augmented Reality Integration for Firefighter Navigation in Unknown Multi-floor Buildings. *Fire Technology*, 62(2), 2026.

- [24] Eudald Sangenis Rafart. *Algorithms and Implementation of Self-Contained Indoor Localization and Navigation for Dismounted Agents*. PhD thesis, University of California, Irvine, Irvine, CA, USA, 2026.
- [25] J. Rantakokko, P. Strömbäck, and P. Andersson. Foot- and knee-mounted INS for firefighter localization. In *International Technical Meeting of the Institute of Navigation (ION ITM)*, pages 145–153, San Diego, CA, USA, Jan. 2014.
- [26] National Institute of Standards and Technology. PSCR Awards \$8M for First Responder 3D Indoor Tracking Prize. <https://www.nist.gov/news-events/news/2021/04/pscr-awards-8m-first-responder-3d-indoor-tracking-prize>, April 2021. Accessed: Jul. 27, 2026.
- [27] National Institute of Standards and Technology. Saving Seconds, Saving Lives: NIST-Funded Challenge Crowns Winners in 3D Tracking Technology. <https://www.nist.gov/news-events/news/2024/01/saving-seconds-saving-lives-nist-funded-challenge-crowns-winners-3d>, January 2024. Accessed: Jul. 27, 2026.
- [28] Rainer Mautz. *Indoor Positioning Technologies*. PhD thesis, ETH Zurich, 2012.
- [29] Christoph Fuchs, Nils Aschenbruck, Peter Martini, and Monika Wieneke. Indoor tracking for mission critical scenarios: A survey. *Pervasive and Mobile Computing*, 7(1):1–15, 2011.
- [30] Estefanía Munoz-Diaz, Dina Bousdar Ahmed, and Susanna Kaiser. A review of indoor localization methods based on inertial sensors. In Jordi Conesa, Antoni Pérez-Navarro, Joaquín Torres-Sospedra, and Raul Montoliu, editors, *Geographical and Fingerprinting Data to Create Systems for Indoor Positioning and Indoor/Outdoor Navigation*, pages 311–333. Academic Press, 2019.
- [31] Naser El-Sheimy and You Li. Indoor navigation: state of the art and future trends. *Satellite Navigation*, 2(1), 2021.
- [32] Luca Mainetti, Luigi Patrono, and Ilaria Sergi. A survey on indoor positioning systems. In *International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 111–120, Split, Croatia, Sep. 2014.
- [33] Chi-Shih Jao, Danmeng Wang, and Andrei M. Shkel. UWB-Foot-SLAM: Bounding Position Error of Foot-mounted Pedestrian INS with Simultaneously Localized UWB Beacons. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Monterey, CA, USA, Apr. 2023.
- [34] Laura Ruotsalainen, Heidi Kuusniemi, and Ruizhi Chen. Overview of methods for visual-aided pedestrian navigation. In *Ubiquitous Positioning Indoor Navigation and Location Based Service*, Kirkkonummi, Finland, Oct. 2010.

- [35] Changwei Chen and Solmaz S. Kia. Cooperative Localization Using Learning-Based Constrained Optimization. *IEEE Robotics and Automation Letters*, 7(3):7052–7058, 2022.
- [36] Chi-Shih Jao, Ali A. Abdallah, Changwei Chen, Min-Won Seo, Solmaz S. Kia, Zagher M. Kassas, and Andrei M. Shkel. PINDOC: Pedestrian Indoor Navigation System Integrating Deterministic, Opportunistic, and Cooperative Functionalities. *IEEE Sensors Journal*, 22(14), 2022.
- [37] Chi-Shih Jao, Yusheng Wang, Sina Askari, and Andrei M Shkel. A Closed-Form Analytical Estimation of Vertical Displacement Error in Pedestrian Navigation. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Portland, OR, USA, Apr. 2020.
- [38] Chi-Shih Jao, Austin R. Parrish, and Andrei M. Shkel. ”Sugar-Cube” PLT: A Real-time Pedestrian Localization Testbed Utilizing Foot-mounted IMU/Barometer/Ultrasonic Sensors. In *IEEE Sensors Conference (SENSORS)*, Sydney, Australia, Oct. 2021.
- [39] Chi-Shih Jao, Yusheng Wang, and Andrei M. Shkel. A Zero Velocity Detector for Foot-mounted Inertial Navigation Systems Aided by Downward-facing Range Sensor. In *IEEE Sensors Conference (SENSORS)*, Rotterdam, Netherlands, Oct. 2020.
- [40] Andre Filipe Goncalves Ferreira, Duarte Manuel Azevedo Fernandes, Andre Paulo Catarino, and Joao L. Monteiro. Localization and Positioning Systems for Emergency Responders: A Survey. *IEEE Communications Surveys and Tutorials*, 19(4):2836–2870, 2017.
- [41] Chi Shih Jao, Yusheng Wang, Yu Wei Lin, and Andrei M. Shkel. A hybrid barometric/ultrasonic altimeter for aiding ZUPT-based inertial pedestrian navigation systems. In *International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+)*, pages 2500–2517, St. Louis, MO, USA, Sep. 2020.
- [42] John Olof Nilsson, Amit K. Gupta, and Peter Händel. Foot-mounted inertial navigation made easy. In *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Busan, Korea, Oct. 2014.
- [43] Michael Perlmutter and Laurent Robin. High-performance, low cost inertial MEMS: A market in motion! In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Myrtle Beach, SC, USA, Apr. 2012.
- [44] Alexandra Efimovskaya. *Multi-Axis Solutions for MEMS Inertial Sensors*. PhD thesis, University of California, Irvine, 2017.
- [45] Yu Wei Lin, Alexandra Efimovskaya, and Andrei M. Shkel. Folded MEMS Platform Based on Polymeric Flexible Hinges for 3D Integration of Spatially-Distributed Sensors. *IEEE/ASME Journal of Microelectromechanical Systems*, 30(6):907–914, 2021.

- [46] Oliver J. Woodman. An Introduction to Inertial Navigation. Technical Report UCAM-CL-TR-696, University of Cambridge, Computer Laboratory, Aug. 2007.
- [47] Eric Foxlin. Pedestrian Tracking with Shoe-Mounted Inertial Sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.
- [48] Johan Wahlstrom and Isaac Skog. Fifteen Years of Progress at Zero Velocity: A Review. *IEEE Sensors Journal*, 21(2):1139–1151, 2021.
- [49] Inertial Elements. Inertial Elements. <https://inertiaelements.com/index.html>, 2014. Accessed: Jul. 2026.
- [50] Robotic Research, LLC. WarLoc. <https://www.roboticresearch.com/warloc/>, 2019. Accessed: 2023.
- [51] Isaac Skog, Peter Händel, John-Olof Nilsson, and Jouni Rantakokko. Zero-velocity detection in pedestrian navigation systems — an algorithm evaluation. *IEEE Transactions on Biomedical Engineering*, 57(11):2657–2666, 2010.
- [52] Yusheng Wang and Andrei M. Shkel. Adaptive Threshold for Zero-Velocity Detector in ZUPT-Aided Pedestrian Inertial Navigation. *IEEE Sensors Letters*, 3(11):1–4, 2019.
- [53] Rui Zhang, Hai Yang, Fabian Höflinger, and Leonhard M. Reindl. Adaptive Zero Velocity Update Based on Velocity Classification for Pedestrian Tracking. *IEEE Sensors Journal*, 17(7):2137–2145, 2017.
- [54] Eudald Sangenis, Chi-Shih Jao, and Andrei M Shkel. SVM-based Motion Classification Using Foot-mounted IMU for ZUPT-aided INS. In *IEEE Sensors Conference (SENSORS)*, Dallas, TX, USA, Oct. 2022.
- [55] Jörg F Wagner, Michael Kohl, and Benedikt Györfi. Reevaluation of Algorithmic Basics for ZUPT-based Pedestrian Navigation. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Avignon, France, May 2022.
- [56] Chi-Shih Jao and Andrei M. Shkel. ZUPT-aided INS Bypassing Stance Phase Detection by Using Foot-Instability-Based Adaptive Covariance. *IEEE Sensors Journal*, 21(21):24338–24348, 2021.
- [57] Chi-Shih Jao, Danmeng Wang, and Andrei M. Shkel. Prio-IMU: Prioritizable IMU Array for Enhancing Foot-Mounted Inertial Navigation Accuracy. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Kauai, HI, USA, Mar. 2023.
- [58] Yusheng Wang, Sina Askari, and Andrei M. Shkel. Study on Mounting Position of IMU for Better Accuracy of ZUPT-Aided Pedestrian Inertial Navigation. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, pages 116–119, Naples, FL, USA, Apr. 2019.

- [59] A. Peruzzi, U. Della Croce, and A. Cereatti. Estimation of stride length in level walking using an inertial measurement unit attached to the foot: A validation of the zero velocity assumption during stance. *Journal of Biomechanics*, 44(10):1991–1994, 2011.
- [60] Isaac Skog, John Olof Nilsson, and Peter Händel. An open-source multi inertial measurement unit (MIMU) platform. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Laguna Beach, CA, USA, Feb. 2014.
- [61] Sina Askari, Chi Shih Jao, Yusheng Wang, and Andrei M. Shkel. A Laboratory Testbed for Self-Contained Navigation. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Naples, FL, USA, Apr. 2019.
- [62] Chi Shih Jao, Danmeng Wang, Austin R. Parrish, and Andrei M. Shkel. A Neural Network Approach to Mitigate Thermal-Induced Errors in ZUPT-aided INS. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Avignon, France, May 2022.
- [63] Gavin P. Horn, Richard M. Kesler, Steve Kerber, Kenneth W. Fent, Trevor J. Schroeder, William S. Scott, Patricia C. Fehling, Bo Fernhall, and Denise L. Smith. Thermal response to firefighting activities in residential structure fires: impact of job assignment and suppression tactic. *Ergonomics*, 61(3):404–419, 2018.
- [64] Xin Zhang, Seungbae Park, and Michael W. Judy. Accurate assessment of packaging stress effects on MEMS sensors by measurement and sensor-package interaction simulations. *IEEE/ASME Journal of Microelectromechanical Systems*, 16(3):639–649, 2007.
- [65] Zhelong Wang, Hongyu Zhao, Sen Qiu, and Qin Gao. Stance-Phase Detection for ZUPT-Aided Foot-Mounted Pedestrian Navigation System. *IEEE/ASME Transactions on Mechatronics*, 20(6):3170–3181, 2015.
- [66] Eudald Sangenis, Chi-Shih Jao, and Andrei M. Shkel. Effect of Adding Time Correlation to SVM-Based Motion Classification in Pedestrian Navigation. *IEEE Journal of Indoor and Seamless Positioning and Navigation*, 3:32–42, 2025.
- [67] Johan Wahlström, Isaac Skog, Fredrik Gustafsson, Andrew Markham, and Niki Trigoni. Zero-Velocity Detection: A Bayesian Approach to Adaptive Thresholding. *IEEE Sensors Letters*, 3(6), Jun. 2019.
- [68] Brandon Wagstaff, Valentin Peretroukhin, and Jonathan Kelly. Improving Foot-Mounted Inertial Navigation Through Real-Time Motion Classification. In *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Sapporo, Japan, Sep. 2017.
- [69] Sean Gallagher, Jonisha Pollard, and William L. Porter. Locomotion in restricted space: Kinematic and electromyographic analysis of stoopwalking and crawling. *Gait and Posture*, 33(1):71–76, 2011.

- [70] Li Cao, Gerard A. Davis, Sean Gallagher, Mark C. Schall, and Richard F. Sesek. Characterizing posture and associated physiological demand during evacuation. *Safety Science*, 104:1–9, 2018.
- [71] Huiju Park, Seonyoung Kim, Kristen Morris, Melissa Moukperian, Youngjin Moon, and Jeffrey Stull. Effect of firefighters’ personal protective equipment on gait. *Applied Ergonomics*, 48:42–48, 2015.
- [72] Michel Laverne, Michael George, Dale Lord, Alonzo Kelly, and Tamal Mukherjee. Experimental Validation of Foot to Foot Range Measurements in Pedestrian Tracking. In *International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+)*, Portland, OR, USA, Sep. 2011.
- [73] Miaoxin Ji, Xiangbo Xu, Zhe Li, Jiasheng Wang, and Jinhao Liu. A Zero-Position-Difference ZUPT Method for Foot-Shank-Mounted Pedestrian Inertial Navigation Systems. *IEEE Sensors Journal*, 21(22):25649–25658, 2021.
- [74] Hojin Ju and Chan Gook Park. A pedestrian dead reckoning system using a foot kinematic constraint and shoe modeling for various motions. *Sensors and Actuators A: Physical*, 284:135–144, 2018.
- [75] Dina Bousdar Ahmed, Estefanía Munoz-Diaz, and Juan Jesús García Domínguez. Novel multi-IMU tight coupling pedestrian localization exploiting biomechanical motion constraints. *Sensors*, 20, 2020.
- [76] Yusheng Wang, Yu Wei Lin, Sina Askari, Chi-Shih Jao, and Andrei M. Shkel. Compensation of Systematic Errors in ZUPT-Aided Pedestrian Inertial Navigation. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Portland, OR, USA, Apr. 2020.
- [77] C. V. Heer. History of the Laser Gyro. In *Physics of Optical Ring Gyros*, Snowbird, UT, USA, 1984.
- [78] Vittorio M.N. Passaro, Antonello Cuccovillo, Lorenzo Vaiani, Martino De Carlo, and Carlo Edoardo Campanella. Gyroscope technology and applications: A review in the industrial perspective. *Sensors (Switzerland)*, 17(10), 10 2017.
- [79] Erik. J. Eklund. *Microgyroscope Based on Spin-Polarized Nuclei*. PhD thesis, University of California, Irvine, 2008.
- [80] Lisa M. Lust and Dan W. Youngner. Chip Scale Atomic Gyroscope. US 7,359,059 B2, 2008.
- [81] Elizabeth. A. Donley and John. Kitching. Nuclear magnetic resonance gyroscopes. In *Optical Magnetometry*, chapter 19, pages 369–386. Cambridge University Press, 2011.
- [82] John Kitching, Svenja Knappe, and Elizabeth A. Donley. Atomic sensors - A review. *IEEE Sensors Journal*, 11(9):1749–1758, 2011.

- [83] Radwan M. Noor. *MEMS components for NMR atomic sensors*. PhD thesis, University of California, Irvine, 2019.
- [84] Doruk. Senkal, S. Askari, M. J. Ahamed, and E. J. Ng. 100K Q-Factor Toroidal Ring Gyroscope implemented in Wafer-Level Epitaxial Silicon Encapsulation Process. In *IEEE International Conference on Micro Electro Mechanical Systems (MEMS)*, San Francisco, CA, USA, Jan. 2014.
- [85] Parsa Taheri-Tehrani, Oleg Izyumin, Igor Izyumin, Chae H. Ahn, Eldwin J. Ng, Vu A. Hong, Yushi Yang, Thomas W. Kenny, Bernhard E. Boser, and David A. Horsley. Disk resonator gyroscope with whole-angle mode operation. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Hapuna Beach, HI, USA, Mar. 2015.
- [86] Ahmad Rahbar Ranji, Vijayakanthan Damodaran, Kevin Li, Zilang Chen, Shahpour Alirezaee, and Mohammed Jalal Ahamed. Recent Advances in MEMS-Based 3D Hemispherical Resonator Gyroscope (HRG)—A Sensor of Choice. *Micromachines*, 13(10):1–30, 2022.
- [87] David M. Rozelle. The Hemispherical Resonator Gyro: From Wineglass to the Planets. In *AAS/AIAA Space Flight Mech. Meeting*, Savannah, GA, USA, Feb. 2009.
- [88] Jill Kipnis. Northrop Grumman’s Hemispherical Resonator Gyro Achieves 50 Million Operating Hours in Space. <https://news.northropgrumman.com/news/releases>, 2019. Accessed: 2023.
- [89] Mohammad H. Asadian. *Development of 3D High-Q Fused Quartz Micro Structures for Precision Coriolis Vibratory Gyroscopes*. PhD thesis, University of California, Irvine, 2019.
- [90] Doruk. Senkal and A. M. Shkel. *Whole-Angle MEMS Gyroscopes*. Wiley-IEEE Press, 2020.
- [91] Doruk Senkal. *Micro-glassblowing Paradigm for Realization of Rate Integrating Gyroscopes*. PhD thesis, University of California, Irvine, 2015.
- [92] Andrei M. Shkel. Type I and Type II micromachined vibratory gyroscopes. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 586–593, San Diego, CA, USA, Apr. 2006.
- [93] Daryosh Vatanparvar. *Study of Methods for Microfabrication of Fused Silica Planar CVGs and Direct Angle Measurements*. PhD thesis, University of California, Irvine, 2022.
- [94] Matthews Anthony. The operation and mechanization of the hemispherical resonator gyroscope. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Monterey, CA, USA, Apr. 2018. IEEE.

- [95] Erik Jesper Eklund and Andrei M Shkel. Glass Blowing on a Wafer Level. *IEEE/ASME Journal of Microelectromechanical Systems*, 16(2):232–239, 2007.
- [96] D. Senkal, C. R. Raum, A. A. Trusov, and A. M. Shkel. Titania Silicate / Fused Quartz Glassblowing for 3-D Fabrication of Low Internal Loss Wineglass Micro-structures. In *Solid-State Sensors, Actuators and Microsystems Workshop (TRANSDUCERS)*, pages 267–270, Anchorage, AK, USA, Jun. 2012.
- [97] Doruk Senkal, Mohammed J. Ahamed, Alexander A. Trusov, and Andrei M. Shkel. Achieving sub-hz frequency symmetry in micro-glassblown wineglass resonators. *IEEE/ASME Journal of Microelectromechanical Systems*, 23(1):30–38, 2014.
- [98] David D. Lynch. Bell Gyro and Improved Means for Operating Same. U.S. Patent US 3,656,354, Apr. 1972.
- [99] Lois Meira-Lopez, Danmeng Wang, Austin R Parrish, and Andrei M Shkel. A Method to Extract and Model Structural Asymmetries in Dual-Shell Gyroscopes to Study Frequency Degeneracy. In *Solid-State Sensors, Actuators and Microsystems Workshop (Hilton Head)*, Jun. 2024.
- [100] Wei Guan, Danmeng Wang, Mohammad H. Asadian, Yusheng Wang, and Andrei M. Shkel. Effect of Geometry on Energy Losses In Fused Silica Dual-Shell Gyroscopes. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Avignon, France, May 2022.
- [101] Daryosh. Vatanparvar and A. M. Shkel. Rate-Integrating Gyroscope Operation in The Non-Linear Regime. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Avignon, France, May 2022.
- [102] Kun Lu, Xiang Xi, Wei Li, Yan Shi, Zhanqiang Hou, Ming Zhuo, Xuezhong Wu, Yulie Wu, and Dingbang Xiao. Research on precise mechanical trimming of a micro shell resonator with T-shape masses using femtosecond laser ablation. *Sensors and Actuators A: Physical*, 290:228–238, 2019.
- [103] Jonathan J. Bernstein, Mirela G. Bancu, Joseph M. Bauer, Eugene H. Cook, Parshant Kumar, Eric Newton, Tenzin Nyinjee, Gayatri E. Perlin, Joseph A. Ricker, William A. Teynor, and Marc S. Weinberg. High Q diamond hemispherical resonators: Fabrication and energy loss mechanisms. *Journal of Micromechanics and Microengineering*, 25(8), 2015.
- [104] Youwang Hu, Haoning Zheng, Yao Wang, Yalong Wang, Xiaoyan Sun, and Ji'an Duan. Frequency Tuning Method on Teeth-Like Tines of the Fused Silica Micro-Hemispherical Resonator Using Femtosecond Laser. *IEEE/ASME Journal of Microelectromechanical Systems*, 33(1):78–87, 2024.
- [105] Bin Li, Xiang Xi, Kun Lu, Yan Shi, Dingbang Xiao, and Xuezhong Wu. Study of Electrode Ablation for Femtosecond Laser Trimming of Micro Hemispherical Resonator Gyroscope. *IEEE Sensors Journal*, 23(3):1933–1942, 2023.

- [106] Wenming Zhang, Yinan Zhang, Haoyu Gu, Zhihui Lin, Qi Wei, Bin Zhou, and Rong Zhang. The High-Precision Balance Method For Hemispherical Resonator Based on Synchronous Trimming System. *IEEE Sensors Journal*, 24(20):31780–31788, 2024.
- [107] Randall Kubena, Richard Joyce, Robert Thomas M’Closkey, and A. Dorian Challner. Microgyro Tuning using Focused Ion Beams, 2004.
- [108] Yuting Wang, Yao Pan, Tianliang Qu, Yonglei Jia, Kaiyong Yang, and Hui Luo. Decreasing frequency splits of hemispherical resonators by chemical etching. *Sensors (Switzerland)*, 18(11), 2018.
- [109] Yusheng Wang, Mohammad H. Asadian, and Andrei M. Shkel. Compensation of frequency split by directional lapping in fused quartz micro wineglass resonators. *Journal of Micromechanics and Microengineering*, 28(9), 2018.
- [110] David. D. Lynch. Vibratory Gyro Analysis by the Method of Averaging. In *Saint Petersburg International Conference on Gyroscopic Technology and Navigation*, Saint Petersburg, Russia, May 1995.
- [111] Sina Askari, Mohammad H. Asadian, Kasra Kakavand, and Andrei M. Shkel. Vacuum sealed and getter activated MEMS Quad Mass Gyroscope demonstrating better than 1.2 million quality factor. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Laguna Beach, CA, USA, Mar. 2016.
- [112] N. Gireesh Sharma, T. Sundararajan, and Sachin Singh Gautam. Identification of limiting damping mechanisms in a high quality factor hybrid resonator of space application gyroscope. *Advances in Space Research*, 69(3):1662–1679, 2022.
- [113] Tal Nagourney, Jae Yoong Cho, Ali Darvishian, Behrouz Shiari, and Khalil Najafi. Effect of Metal Annealing on the Q-Factor of Metal-Coated Fused Silica Micro Shell Resonators. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Hapuna Beach, HI, USA, Mar. 2015.
- [114] Ron Lifshitz and M L Roukes. Thermoelastic damping in micro- and nanomechanical systems. *Physical Review B*, 61(8):5600–5609, 2000.
- [115] Corning Inc. *Corning® HPFS® 7979, 7980, 8655 Fused Silica*, Jul. 2015.
- [116] Philip. E. Miller, T. I. Suratwala, L. L. Wong, M. D. Feit, J. A. Menapace, P. J. Davis, and R. A. Steele. The distribution of subsurface damage in fused silica. In *Laser-Induced Damage in Optical Materials*, Boulder, CO, USA, Sep. 2005.
- [117] John C. Lambropoulos. From Abrasive Size To Subsurface Damage in Grinding. In *Optical Fabrication and Testing*, Québec City, Canada, Jun. 2000.
- [118] Jianning Chu, Xin Liu, Changlin Liu, Jianguo Zhang, Junfeng Xiao, Xuefang Wang, Xiao Chen, and Jianfeng Xu. Fundamental investigation of subsurface damage on the quality factor of hemispherical fused silica shell resonator. *Sensors and Actuators A: Physical*, 335, 2022.

- [119] Y. J. Zhai, Y. Pan, Y. L. Jia, J. P. Liu, Z. Q. Tan, K. Y. Yang, and H. Luo. Surface evolution of cylindrical fused silica resonator and its implication on Q factor. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Naples, FL, USA, Apr. 2019.
- [120] A. Esmaeilzare, A. Rahimi, and S. M. Rezaei. Investigation of subsurface damages and surface roughness in grinding process of Zerodur® glass-ceramic. *Applied Surface Science*, 313(15):67–75, 2014.
- [121] B. S. Lunin, A. N. Kharlanov, and S. E. Kozlov. Dehydroxylation and formation of KU-1 silica glass surface defects during annealing. *Moscow University Chemistry Bulletin*, 65(1):34–37, 2010.
- [122] B. S. Lunin and V. M. Lopatin. Surface Internal Friction in High-Q Fused Quartz Resonators. *Inorganic Materials*, 58(6):636–642, 2022.
- [123] Sajal Singh, Jae Yoong Cho, Tal Nagourney, and Khalil Najafi. Effect of adsorbed water on quality-factor of fused silica uncoated and Platinum-coated shell resonators. *IEEE Sensors Letters*, Feb. 2026.
- [124] Mohammed J. Ahamed, Doruk Senkal, and Andrei M. Shkel. Effect of Annealing on Mechanical Quality Factor of Fused Quartz Hemispherical Resonator. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Laguna Beach, CA, USA, Feb. 2014.
- [125] Yusheng Wang and Andrei M. Shkel. Study on surface roughness improvement of Fused Quartz after thermal and chemical post-processing. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Laguna Beach, CA, USA, Feb. 2016.
- [126] Danmeng Wang, Mohammad H. Asadian, Wei Guan, Doreen Hii, Austin R. Parrish, and Andrei M. Shkel. Effect of Metallization on Quality Factor and Noise Characteristics in Fused Silica Dual-Shell Gyroscopes. *IEEE/ASME Journal of Microelectromechanical Systems*, 31(6):877–887, 2022.
- [127] Jae Yoong Cho, Jialiang Yan, Jeffrey A. Gregory, Harald W. Eberhart, Rebecca L. Peterson, and Khalil Najafi. 3-Dimensional Blow Torch-Molding of Fused Silica Microstructures. *IEEE/ASME Journal of Microelectromechanical Systems*, 22(6):1276–1284, 2013.
- [128] Sajal Singh, Jae Yoong Cho, Tal Nagourney, and Khalil Najafi. Impact of Metallization on Quality Factor of Fused-Silica Shell Resonators. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Destin, FL, USA, May 2026.
- [129] V. Stankevič and G. Račiukaitis. Formation of rectangular channels in fused silica by laser-induced chemical etching. *Lithuanian Journal of Physics*, 54(3):136–141, 2014.
- [130] C. Hnatovsky, R. S. Taylor, E. Simova, P. P. Rajeev, D. M. Rayner, V. R. Bhardwaj, and P. B. Corkum. Fabrication of microchannels in glass using focused femtosecond

- laser radiation and selective chemical etching. *Applied Physics A: Materials Science and Processing*, 84:47–61, 2006.
- [131] Calum A. Ross, David G. MacLachlan, Debaditya Choudhury, and Robert R. Thomson. Optimisation of ultrafast laser assisted etching in fused silica. *Optics Express*, 26(19):24343–24356, 2018.
 - [132] Daryosh Vatanparvar and Andrei M. Shkel. Digital Manufacturing of Resonance MEMS From A Single-Layer Fused Silica Material. In *IEEE International Conference on Micro Electro Mechanical Systems (MEMS)*, Tokyo, Japan, Jan. 2022.
 - [133] Martin Hermans, Jens Gottmann, and Frank Riedel. Selective, laser-induced etching of fused silica at high scan-speeds using KOH. *Journal of Laser Micro Nanoengineering*, 9(2):126–131, 2014.
 - [134] Agnė Butkutė and Linas Jonušauskas. 3D Manufacturing of Glass Microstructures Using Femtosecond Laser. *Micromachines*, 12(5), 2021.
 - [135] V. S. Kuznetsov, D. V. Maiorov, and P. A. Andreev. Fabrication of Tuning-Fork Gyroscope Resonator using Directed Chemical Etching of Fused Silica. *Herald of the Bauman Moscow State Technical University. Series Instrument Engineering*, (4):32–44, 2019.
 - [136] Tao Zhao, Ming Zhuo, Xin Zhou, Xiang Xi, Xuezhong Wu, and Dingbang Xiao. Fused Silica Gyroscope Resonator Manufactured With Femtosecond Laser Assisted Wet Etching. *IEEE/ASME Journal of Microelectromechanical Systems*, 31(3):315–317, 2022.
 - [137] Davide Pavesi, Lois Meira-Lopez, Eudald Sangenis, Valentina Zega, and Andrei M. Shkel. Design, Fabrication and Characterization of a High-Quality Factor Fused Silica Toroidal Ring Gyroscope. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Lindau, Germany, Mar. 2025.
 - [138] Austin R Parrish, Lois Meira-Lopez, Mohammad H Asadian, Esther Wong, and Andrei M Shkel. Laser Patterned Self-Aligned Electrodes for Dual-Shell Hemispherical Resonator Gyroscope. In *Solid-State Sensors, Actuators and Microsystems Workshop (TRANSDUCERS)*, Orlando, FL, USA, Jun. 2025.
 - [139] David M. Rozelle, Youngmin Albert Choi, and James M. Pavell. Wide Rim Vibratory Resonant Sensors. U.S. Patent US 10,119,820 B2, Nov. 2018.
 - [140] Youngmin A. Choi and David M. Rozelle. Inner-Forcer Milli-Hemispherical Resonator Gyroscope. U.S. Patent US 7,839,059 B2, Nov. 2010.
 - [141] Jae Yoong Cho and Khalil Najafi. Vibrating Gyroscopes with Resonator Attachments. U.S. Patent US 11,874,112 B1, Jan. 2024.
 - [142] Khalil Najafi, Jae Yoong Cho, Ali Darvishian, Guohong He, Behrouz Shiari, and Tal Nagourney. Gyroscope and Fabrication Process. U.S. Patent Application US 2017/0248422 A1, Aug. 2017.

- [143] Yu Wang, Chaoliang Guan, Yifan Dai, and Shuai Xue. On-Machine Measurement of Profile and Concentricity for Ultra-Precision Grinding of Hemispherical Shells. *Micro-machines*, 13(10), 2022.
- [144] Danmeng Wang. *Micromachined Vibratory Gyroscopes Designed for Precision of Measurements and Environmental Robustness*. PhD thesis, University of California, Irvine, 2022.
- [145] Zhihu Ruan, Xukai Ding, Zhengcheng Qin, Jia Jia, and Hongsheng Li. Modeling and Compensation of Assembly Inclination Error of Micro Hemispherical Resonator Gyroscope under Force-to-Rebalance Mode. *IEEE Sensors Journal*, 21(13):14726–14738, 2021.
- [146] Andrei M. Shkel, Alexander A. Trusov, Igor P. Prikhodko, and Sergei A. Zotov. Three-Dimensional Wafer-Scale Batch-Micromachined Sensor and Method of Fabrication for the Same. U.S. Patent Application US 2011/0239763 A1, Oct. 2011.
- [147] Sajal Singh, Jae Yoong Cho, Jong Kwan Woo, Ester Bentley, and Khalil Najafi. Shell-in-Shell (SiS): 3D Shell Resonator with 3D Conformal Shell Electrodes. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Hiroshima, Japan, Mar. 2020.
- [148] Yinan Zhang, Mingwei Ma, Bin Zhou, Haoyu Gu, Qi Wei, and Rong Zhang. Assembly-Free Electrode Manufacturing Process for Micro-Hemispherical Resonant Gyroscopes with Enhanced Capacitance. In *IEEE Sensors Conference (SENSORS)*, Vancouver, Canada, Oct. 2025.
- [149] Farrokh Ayazi, Peng Shao, Xin Gao, and Vahid Tavassoli. Micro-Hemispherical Resonators and Methods of Making the Same. U.S. Patent Application US 2016/0344368 A1, Nov. 2016.
- [150] Tiwei Wei, Jian Cai, Qian Wang, Yang Hu, Lu Wang, Ziyu Liu, and Zijian Wu. Optimization and Evaluation of Sputtering Barrier/Seed Layer in Through Silicon Via for 3-D Integration. *Tsinghua Science and Technology*, 19(2), 2014.
- [151] I. J.M. Erkens, M. A. Verheijen, H. C.M. Knoop, W. Keuning, F. Roozeboom, and W. M.M. Kessels. Plasma-assisted atomic layer deposition of conformal Pt films in high aspect ratio trenches. *Journal of Chemical Physics*, 146(5), 2 2017.
- [152] Feng Zhu, Jia Liu, Chen Liao, Xuezhong Wu, Yan Shi, and Xiang Xi. Fabrication of highly homogeneous metal films for hemispherical resonators by a 3D mask based on a sputter-coating process. *Thin Solid Films*, 821, 7 2025.
- [153] Chenguang Xin, Yingkun Xu, Zhongyao Zhang, and Mengwei Li. Micro-Opto-Electro-Mechanical Systems for High-Precision Displacement Sensing: A Review. *Micromachines*, 15(8), 2024.

- [154] Wenyi Huang, Senyu Zhang, Jamal N. A. Hassan, Xing Yan, Dingwei Chen, Guangjun Wen, Kai Chen, Guangwei Deng, and Yongjun Huang. High-precision angular rate detection based on an optomechanical micro hemispherical shell resonator gyroscope. *Optics Express*, 31(8):12433–12448, 2023.
- [155] Jamal N.A. Hassan, Wenyi Huang, Maoyuan Wang, Senyu Zhang, Guangjun Wen, and Yongjun Huang. Optomechanical Gyroscope Based on Micro-Hemispherical Shell and Optical Ring Resonators. *IEEE Photonics Journal*, 16(4), 2024.
- [156] Marcel W. Pruessner, Doewon Park, Todd H. Stievater, Dmitry A. Kozak, and William S. Rabinovich. An optomechanical transducer platform for evanescent field displacement sensing. *IEEE Sensors Journal*, 14(10):3473–3481, 2014.
- [157] Onur Kilic, Hyejun Ra, Onur Can Akkaya, Michel J.F. Digonnet, and Olav Solgaard. Haltere-Like Optoelectromechanical Gyroscope. *IEEE Sensors Journal*, 16(11):4274–4280, 2016.
- [158] Ozan Berk Boyraz, Ceren Babayigit, Lois Meira-Lopez, Austin R Parrish, Shafiqul Islam, Andrei M Shkel, and Ozdal Boyraz. Fused-Silica Dual-Shell Resonators with Photonic Integrated Readout. In *Conference on Lasers and Electro-Optics (CLEO)*, Long Beach, CA, USA, May 2025.
- [159] Ozan Berk Boyraz, Lois Meira-Lopez, Austin R. Parrish, Md Shafiqul Islam, Ceren Babayigit, Ataberk Atalar, Xun Li, Andrei M. Shkel, and Ozdal Boyraz. Chip-Scale Precision Optical Displacement Readout of Fused Silica Dual-Shell Resonators. *Optics Express*, 34(3):4212–4227, 2026.
- [160] Eugene Freeman and Chad Fertig. System and Methods for Optical Excitation of Acoustic Resonance in Resonator Gyroscopes. U.S. Patent Application US 2020/0173779 A1, Jun. 2020.
- [161] Lois Meira-Lopez, Melika Momenzadeh, Maxim R. Shcherbakov, and Andrei M. Shkel. Experimental Demonstration of Optical Excitation of a Dual-Shell Resonator. *IEEE Sensors Letters*, 9(8), 2025.
- [162] Melika Momenzadeh, Lois Meira-Lopez, Austin R. Parrish, Andrei M. Shkel, and Maxim R. Shcherbakov. Spatially Programmable Optical Excitation of Orthogonal Hemispherical Micro-Mechanical Resonator Modes. In *Solid-State Sensors, Actuators and Microsystems Workshop (Hilton Head)*, Hilton Head Island, SC, USA, May 2026.
- [163] Danmeng Wang, Nicholas A. Strnad, Yusheng Wang, Austin R. Parrish, Robert R. Benoit, Ryan R. Knight, and Andrei M. Shkel. Indirect Excitation of micro-HRG Using Segmented Piezoelectric ALD PHT Actuator. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Kauai, HI, USA, Mar. 2023.
- [164] Danmeng Wang, Nicholas A. Strnad, and Andrei M. Shkel. Initial Demonstration of Fused Silica Dual-Shell Gyroscope Using Indirect Method of Piezoelectric Excitation. *IEEE Sensors Letters*, 7(9), 2023.

- [165] Andrei M. Shkel, Mohammad H. Asadian Ardakani, and Yusheng Wang. Fused Quartz Dual Shell Resonator and Method of Fabrication. U.S. Patent US 11,703,330 B2, Jul. 2023.
- [166] Nicholas A. Strnad, Ryan R. Knight, Ryan Q. Rudy, Robert R. Benoit, Wendy L. Sarney, and Jeffrey S. Pulskamp. Atomic Layer Deposition and Sputtering of Piezoelectric Thin Films for Improved IMU Performance. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Monterey, CA, USA, Apr. 2023.
- [167] TDK InvenSense. ICM-42605: High Performance 6-Axis MEMS MotionTracking Device. Datasheet, Document DS-000292, 2020. Accessed: Jul. 2026.
- [168] STMicroelectronics. LSM6DSV: iNEMO Inertial Module with Embedded Sensor Fusion, AI, Qvar, High Accuracy, High Stability 6-Axis IMU. Datasheet, Document DS13971, 2023. Accessed: Jul. 2026.
- [169] Bosch Sensortec GmbH. BMI323: Compact, Low Power Inertial Measurement Unit. Datasheet, Document BST-BMI323-DS000, 2023. Accessed: Jul. 2026.
- [170] Analog Devices, Inc. ADIS16500: Precision, Miniature MEMS IMU (Data Sheet). Datasheet, 2019. Accessed: Jul. 2026.
- [171] Yusheng Wang, Andrei Chernyshoff, and Andrei M. Shkel. Study on estimation errors in ZUPT-Aided pedestrian inertial navigation due to IMU noises. *IEEE Transactions on Aerospace and Electronic Systems*, 56(3):2280–2291, 2020.
- [172] Xsens Technologies B.V. MTi-3 AHRS: MTi 1-Series Module. Product documentation, 2024. Accessed: Jul. 2026.
- [173] Gladiator Technologies. G200D Dual-Axis MEMS Gyroscope. Datasheet, 2025. Accessed: Jul. 2026.
- [174] Safran Electronics & Defense. Iconyx Series: High-Performance HRG Inertial Measurement Units. Product brochure, 2026. Accessed: Jul. 2026.
- [175] Honeywell Aerospace. HG9900 Inertial Measurement Unit. Datasheet, 2004. Accessed: Jul. 2026.
- [176] Northrop Grumman Corporation. Scalable SIRU Family of Products. Product brochure, Document DS-509-JDK-0421, 2021. Accessed: Jul. 2026.
- [177] Wei Li, Xiang Xi, Kun Lu, Yan Shi, Zhanqiang Hou, Yulie Wu, Xuezhong Wu, and Dingbang Xiao. A Novel Micro Shell Resonator Gyroscope with Sixteen T-Shape Masses. In *Solid-State Sensors, Actuators and Microsystems Workshop (TRANSDUCERS)*, pages 434–437, Berlin, Germany, Jun. 2019.
- [178] Chi-Shih Jao, Eudald Sangenis, Paula Simo, Alexandra Voloshina, and Andrei M. Shkel. An Inverted Pendulum Model of Walking for Predicting Navigation Uncertainty of Pedestrian in Case of Foot-Mounted Inertial Sensors. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Kauai, HI, USA, Mar. 2023.

- [179] David H. Titterton and John L. Weston. *Strapdown Inertial Navigation Technology*, volume 17 of *IET Radar, Sonar, Navigation and Avionics Series*. The Institution of Engineering and Technology (IET), Stevenage, UK, 2nd edition, 2004.
- [180] Eudald Sangenis, Chi-Shih Jao, Danmeng Wang, and Andrei M. Shkel. Foot Velocity in Stance Phase Is Not Zero: Revisiting Velocity Updates for Foot-Mounted Inertial Navigation Systems. *IEEE Sensors Journal*, 26(3):5077–5085, Feb. 2026.
- [181] Jacquelin Perry and Judith M. Burnfield. *Gait Analysis: Normal and Pathological Function*. SLACK Incorporated, Thorofare, NJ, USA, 2nd edition, 2010.
- [182] Austin R. Parrish, Chi-Shih Jao, and Andrei M. Shkel. Stance Phase Detection for ZUPT-Aided INS Using Knee-Mounted IMU in Crawling Scenarios. *IEEE Sensors Letters*, 6(5), 2022.
- [183] Austin R. Parrish, Chi-shih Jao, Danmeng Wang, and Andrei M. Shkel. Study of IMU Mounting Position for ZUPT-Aided INS in the Case of Firefighter Crawling. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Kauai, HI, USA, Mar. 2023.
- [184] Analog Devices, Inc. ADIS16490: Tactical Grade, Six Degrees of Freedom Inertial Sensor (Data Sheet). <https://www.analog.com/media/en/technical-documentation/data-sheets/adis16490.pdf>, 2020. Rev. D. Accessed: Jul. 2026.
- [185] Valerii V. Chikovani et al. Performance Parameters Comparison of Ring Laser, Coriolis Vibratory, and Fiber-Optic Gyros Based on Allan Variance Analysis. In *IEEE International Conference on Actual Problems of Unmanned Aerial Vehicles Development (APUAVD)*, Kyiv, Ukraine, Oct. 2013.
- [186] Jae Hong Lee and Chan Gook Park. Mitigation of a Heading Drift in Pedestrian Dead-reckoning Caused by the Sensor Bandwidth. *International Journal of Control, Automation and Systems*, 19(8):2882–2890, 2021.
- [187] Michael S. Braasch. Inertial Navigation Systems. In Alexander V. Nebylov and Joseph Watson, editors, *Aerospace Navigation Systems*, chapter 1. John Wiley & Sons, Ltd., Chichester, UK, 2016.
- [188] David Y. Hsu. An Accurate and Efficient Approximation to the Normal Gravity. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Palm Springs, CA, USA, Apr. 1998.
- [189] Alberto Olivares, Javier Ramírez, Juan M. Górriz, Gonzalo Olivares, and Miguel Damas. Detection of (In)activity periods in human body motion using inertial sensors: A comparative study. *Sensors (Switzerland)*, 12(5):5791–5814, 2012.
- [190] Rudolf E. Kalman and R. S. Bucy. New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83(1):95–108, 1961.

- [191] Alexander A. Trusov. Allan Variance Analysis of Random Noise Modes in Gyroscopes. Technical report, MicroSystems Laboratory, University of California, Irvine, Irvine, CA, USA, Mar. 2011.
- [192] G H Bryan. On the beats in the vibrations of a revolving cylinder or bell. *Proceedings of the Cambridge Philosophical Society*, 7(3):101–111, 1890.
- [193] Robert P. Leland. Mechanical-thermal noise in MEMS gyroscopes. *IEEE Sensors Journal*, 5(3):493–500, 2005.
- [194] Mohammad H. Asadian, Sina Askari, and Andrei M. Shkel. An Ultrahigh Vacuum Packaging Process Demonstrating over 2 Million Q-Factor in MEMS Vibratory Gyroscopes. *IEEE Sensors Letters*, 1(6), 12 2017.
- [195] Clarence Zener. Internal Friction in Solids I. Theory of Internal Friction in Reeds. *Physical Review*, 52, 1937.
- [196] Clarence Zener. Internal friction in solids II. General theory of thermoelastic internal friction. *Physical Review*, 53:90–99, 1938.
- [197] David S. Bindel, Emmanuel Quévy, Tsuyoshi Koyama, Sanjay Govindjee, James W. Demmel, and Roger T. Howe. Anchor loss simulation in resonators. In *IEEE International Conference on Micro Electro Mechanical Systems (MEMS)*, Miami, FL, USA, Jan. 2005.
- [198] Daniel Sherman. *An Investigation of MEMS Anchor Design for Optimal Stiffness and Damping*. PhD thesis, University of California, Berkeley, 1996.
- [199] Mohammad H. Asadian, Yusheng Wang, and Andrei M. Shkel. Development of 3D Fused Quartz Hemi-Toroidal Shells for High-Q Resonators and Gyroscopes. *IEEE/ASME Journal of Microelectromechanical Systems*, 28(6):954–964, 2019.
- [200] Janna Rodriguez, Saurabh A. Chandorkar, Christopher A. Watson, Grant M. Glaze, C. H. Ahn, Eldwin J. Ng, Yushi Yang, and Thomas W. Kenny. Direct Detection of Akhiezer Damping in a Silicon MEMS Resonator. *Scientific Reports*, 9(1):1–10, 12 2019.
- [201] Yusheng Wang, Mohammad H. Asadian, and Andrei M. Shkel. Modeling the effect of imperfections in glassblown micro-wineglass fused quartz resonators. *ASME Journal of Vibration and Acoustics*, 139(4), 2017.
- [202] Mohammad H Asadian, Danmeng Wang, Yusheng Wang, and Andrei M Shkel. 3D Dual-Shell Micro-Resonators for Harsh Environments. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 1467–1471, Portland, OR, USA, Apr. 2020.
- [203] Austin R. Parrish, Esther Wong, Mohammad H. Asadian, Lois Meira-Lopez, and Andrei M. Shkel. Self-Aligned, High-Aspect Ratio Electrodes for Fused Silica Resonators Digitally Defined by Laser Modification of Material. *IEEE/ASME Journal of Microelectromechanical Systems*, 2026. Early Access.

- [204] Danmeng Wang, Mohammad H. Asadian, Doreen Hii, and Andrei M. Shkel. Fused Silica Dual-Shell Gyroscope with In-Plane Actuation by Out-Of-Plane Electrodes Realized Using Glassblowing and Thru-Glass-Vias Fabrication. In *IEEE International Conference on Micro Electro Mechanical Systems (MEMS)*, Tokyo, Japan, Jan. 2022.
- [205] Daryosh Vatanparvar, Andrei M. Shkel, and Danmeng Wang. Laser Micromachining of MEMS Resonators from Bulk Optically Transparent Material. U.S. Patent Application US 2025/0091861 A1, Mar. 2025.
- [206] Boris S. Lunin and Kirill V. Tokmakov. Reduction in internal friction in silica glass with high OH content. *Journal of the American Ceramic Society*, 102(6):3329–3340, 2018.
- [207] Jens Gottmann, Martin Hermans, Nikolai Repiev, and Jürgen Ortmann. Selective laser-induced etching of 3D precision quartz glass components for microfluidic applications up scaling of complexity and speed. *Micromachines*, 8(4), 2017.
- [208] Clayton R. Paul. *Analysis of Multiconductor Transmission Lines*. Wiley-Interscience : IEEE Press, Hoboken, New Jersey, 2 edition, 2008.
- [209] Clark Nguyen. *Micromechanical Signal Processors*. PhD thesis, University of California, Berkeley, 1994.
- [210] Alexander A. Trusov and Andrei M. Shkel. A novel capacitive detection scheme with inherent self-calibration. *IEEE/ASME Journal of Microelectromechanical Systems*, 16(6):1324–1333, 2007.
- [211] Doruk Senkal, Mohammed J. Ahamed, Mohammad H. Asadian, Sina Askari, and Andrei M. Shkel. Demonstration of 1 million Q -factor on microglassblown wineglass resonators with out-of-plane electrostatic transduction. *IEEE/ASME Journal of Microelectromechanical Systems*, 24(1):29–37, 2015.
- [212] Ali Darvishian, Christopher Boyd, Sajal Singh, Jae-Yoong Cho, Jong-Kwan Woo, Guohong He, and Khalil Najafi. Effect of Electrode Design on Frequency Tuning in Shell Resonators. In *IEEE International Symposium on Inertial Sensors and Systems (IN-ERTIAL)*, Naples, FL, USA, Apr. 2019.
- [213] Changhong Wang, Youhuan Ning, Yan Huo, Lishan Yuan, Wei Cheng, and Zhen Tian. Frequency splitting of hemispherical resonators trimmed with focused ion beams. *International Journal of Mechanical Sciences*, 261, 2024.
- [214] Kun Lu, Dingbang Xiao, Yan Shi, Xuezhong Wu, Xiang Xi, and Yongmeng Zhang. On-line Trimming Device and Method for Micro-Shell Resonator Gyroscope. U.S. Patent US 11,703,356 B1, Jul. 2023.
- [215] Yan Shi, Dingbang Xiao, Kun Lu, Xiang Xi, Xuezhong Wu, Yulie Wu, and Yongmeng Zhang. Elastically Supported Electrode Substrate for Detecting Unbalanced Mass of Resonant Gyroscope. U.S. Patent US 11,885,641 B2, Jan. 2024.

- [216] Bart Macco, Mike L. van de Poll, Bas W.H. van de Loo, Tim M.P. Broekema, Saravana B. Basuvalingam, Cristian A.A. van Helvoirt, Wilhelmus J.H. Berghuis, Roel J. Theeuwes, Nga Phung, and Wilhelmus M.M. Kessels. Temporal and spatial atomic layer deposition of Al-doped zinc oxide as a passivating conductive contact for silicon solar cells. *Solar Energy Materials and Solar Cells*, 245:111869, 2022.
- [217] Maobo Wang, Qingsong Li, Kai Wu, Xinyu Wang, Zhanqiang Hou, Yan Shi, Xuezhong Wu, and Dingbang Xiao. A Novel All-Fused Silica MEMS Gyroscope With an Aspect Ratio Exceeding 50:1. *IEEE/ASME Journal of Microelectromechanical Systems*, 34(5):516–518, 2025.
- [218] Mario Ochoa, Pablo Roldán-Varona, José Francisco Algorri, José Miguel López-Higuera, and Luis Rodríguez-Cobo. Polarisation-independent ultrafast laser selective etching processing in fused silica. *Lab on a Chip*, 23(7):1752–1757, 2023.
- [219] M K Bhuyan, F Courvoisier, P-a Lacourt, M Jacquot, L Furfaro, and M J Withford. High aspect ratio taper-free microchannel fabrication using femtosecond Bessel beams. *Optics Express*, 18(2):566–574, 2010.
- [220] Alexander A. Trusov, Doruk Senkal, and Andrei M. Shkel. Microfabrication of High Quality Three Dimensional Structures Using Wafer-Level Glassblowing of Fused Quartz and Ultra Low Expansion Glasses. U.S. Patent US 9,139,417 B2, Sep. 2015.
- [221] Austin R Parrish, Nicholas A Strnad, Zachariah J Klutts, Lois Meira-Lopez, Danmeng Wang, Jeffrey S Pulskamp, and Andrei M Shkel. Non-Metallized Dual-Shell Hemispherical Resonator Gyroscope. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Salt Lake City, UT, USA, Apr. 2025.
- [222] Yan Shi, Xiang Xi, Bin Li, Yimo Chen, Yulie Wu, Dingbang Xiao, Xuezhong Wu, and Kun Lu. Micro Hemispherical Resonator Gyroscope with Teeth-Like Tines. *IEEE Sensors Journal*, 21(12):13098–13106, 2021.
- [223] Tao Zhao, Yan Shi, Kun Lu, Xiang Xi, Xuezhong Wu, and Dingbang Xiao. Six million Q factor micro fused silica shell resonator with teeth-like tines released by femtosecond laser-assisted chemical etching. *Science China Technological Sciences*, 67, 2024.
- [224] Ozan Berk Boyraz, Lois Meira-Lopez, Austin R. Parrish, Md Shafiqul Islam, Ataberk Atalar, Andrei M. Shkel, and Ozdal Boyraz. Dual-Channel PIC Optical Readout for Modal Characterization of Fused-Silica Hemispherical Resonators. In *Proceedings of SPIE MOEMS and Miniaturized Systems XXV*, San Francisco, CA, USA, Jan. 2026.
- [225] Jonathan M. Puder, Jeffrey S. Pulskamp, Ryan Q. Rudy, Cristian Cassella, Matteo Rinaldi, Guofeng Chen, Sunil A. Bhave, and Ronald G. Polcawich. Rapid Harmonic Analysis of Piezoelectric MEMS Resonators. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 65(6):979–990, 2018.
- [226] Allan F. Bower. *Applied Mechanics of Solids*. Cambridge University Press, Cambridge, UK, 2010.

- [227] Jiashi Yang. *An Introduction to the Theory of Piezoelectricity*. Springer Science, 2005.
- [228] Nicolas Ledermann, Paul Muralt, Jacek Baborowski, Sandrine Gentil, Kapil Mukati, Marco Cantoni, Andreas Seifert, and Nava Setter. {100}-textured, piezoelectric $\text{Pb}(\text{Zr}_x\text{Ti}_{1-x})\text{O}_3$ thin films for MEMS: integration, deposition and properties. *Sensors and Actuators A: Physical*, 105(2):162–170, 2003.
- [229] S. Sivaramakrishnan, P. Mardilovich, A. Mason, A. Roelofs, T. Schmitz-Kempen, and S. Tiedke. Electrode size dependence of piezoelectric response of lead zirconate titanate thin films measured by double beam laser interferometry. *Applied Physics Letters*, 103(13), 2013.
- [230] Polytec GmbH. Metrological Properties VibroFlex Connect VFX-F-110.
- [231] Sang Jin Park, Huinam Rhee, Young Chul Choi, and Jin Ho Park. The Effect of Stem on the Bryan’s Factor of a Hemispherical Resonator. *Journal of Sound and Vibration*, 538, 11 2022.
- [232] Kris Winer. ICM-42605 Library. <https://github.com/kriswiner/ICM42605>, 2019. GitHub repository. Accessed: 2022.
- [233] Jerry Yip. Seed ICM20600 + AK09918 Library. https://github.com/Seed-Studio/Seed_ICM20600_AK09918, 2018. GitHub repository. Accessed: 2022.
- [234] Xin Liu, Ning Li, Geng Xu, and Yonggang Zhang. A Novel Robust Step Detection Algorithm for Foot-Mounted IMU. *IEEE Sensors Journal*, 21(4):5331–5339, 2021.
- [235] Chi-Shih Jao, Yusheng Wang, and Andrei M. Shkel. Pedestrian Inertial Navigation System Augmented by Vision-Based Foot-to-foot Relative Position Measurements. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 900–907, Portland, OR, USA, Apr. 2020.
- [236] Nimsiri Abhayasinghe, Iain Murray, and Shiva Sharif Bidabadi. Validation of thigh angle estimation using inertial measurement unit data against optical motion capture systems. *Sensors (Switzerland)*, 19(3), 2019.
- [237] Shaban Shabani, Alan K. Bourke, Amir Muaremi, Jens Praestgaard, Kate O’Keeffe, Rob Argent, Martin Brom, Celeste Scotti, Brian Caulfield, and Lorcan C. Walsh. An Automatic Foot and Shank IMU Synchronization Algorithm: Proof-of-concept. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4210–4213, Glasgow, UK, Jul. 2022.
- [238] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.
- [239] Sajal Sagar Singh. *Fused Silica Precision Shell Integrating (PSI) Navigation-Grade Micro Gyroscopes*. PhD thesis, University of Michigan, 2020.

- [240] Howard H. Ge, Dennis Kim, and Robert T. M'Closkey. Simultaneous Exploitation of the Fundamental and Higher Order Wineglass Modes in a Vibratory Gyro. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Hapuna Beach, HI, USA, Mar. 2015.
- [241] Howard H. Ge and Robert T. M'Closkey. Simultaneous Angular Rate Estimates Extracted From a Single Axisymmetric Resonator. *IEEE Sensors Journal*, 17(22):7460–7469, Nov. 2017.
- [242] Howard H. Ge, Amir H. Behbahani, and Robert T. M'Closkey. MEMS Gyro Drift Compensation Using Multiple Rate Measurements Derived from a Single Resonator. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 288–293, Monterey, CA, USA, Apr. 2018.
- [243] Lois Meira-Lopez, Austin R. Parrish, and Andrei M. Shkel. Chemical Etching Based Trimming Process for Sub-Hz Permanent Frequency Tuning of Fused Quartz Vibratory Gyroscopes. In *Solid-State Sensors, Actuators and Microsystems Workshop (Hilton Head)*, Hilton Head Island, SC, USA, Jun. 2026. Presented; to appear.
- [244] Jonathan Bernstein, Marc Steven Weinberg, Murali Chaparala, and Peter G. Sherman. MEMS Hemispherical Resonator Gyroscope. U.S. Patent US 9,423,253 B2, Aug. 2016. Assignee: The Charles Stark Draper Laboratory, Inc. Filed Oct. 31, 2012.
- [245] Pieter Vlugter, Erica Block, and Yves Bellouard. Local Tuning of Fused Silica Thermal Expansion Coefficient Using Femtosecond Laser. *Physical Review Materials*, 3(5), May 2019.
- [246] Yves Bellouard, Audrey Champion, Benjamin McMillen, Sebabrata Mukherjee, Robert R. Thomson, Charles Pépin, Philippe Gillet, and Ya Cheng. Stress-State Manipulation in Fused Silica via Femtosecond Laser Irradiation. *Optica*, 3(12):1285–1293, Dec. 2016.
- [247] Pieter Vlugter and Yves Bellouard. Elastic Properties of Self-Organized Nanogratings Produced by Femtosecond Laser Exposure of Fused Silica. *Physical Review Materials*, 4(2), Feb. 2020.
- [248] Carl Tyren. Magnetic Anomalies as a Reference for Ground-speed and Map-matching Navigation. *The Journal of Navigation*, 35(2):242–254, 1982.
- [249] Aaron Canciani and John Raquet. Airborne Magnetic Anomaly Navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 53(1):67–80, 2017.
- [250] Viki Bankey, Alejandro Cuevas, David Daniels, Carol A. Finn, and Israel Hernandez. Magnetic Anomaly Map of North America. United States Geological Survey (USGS), <https://mrdata.usgs.gov/magnetic/>, 2002. Accessed: Jul. 2026.
- [251] Austin R. Parrish, Radwan Noor, and Andrei M. Shkel. Microfabricated Optically Pumped Gradiometer with Uniform Buffer Gases. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Virtual, Mar. 2021. IEEE.

- [252] Lauren Mount. *Navigation using Vector and Tensor Measurements of the Earth's Magnetic Anomaly Field*. PhD thesis, Air Force Institute of Technology, 2018.
- [253] Zhao Ta, Chen Yu Wei, Zhou Zhi Jian, and Cheng de Fu. Research on three-component geomagnetic field differential measurement method for underwater vehicle. *International Journal of Multimedia and Ubiquitous Engineering*, 11(2):89–98, 2016.
- [254] Gautam Sharma, Isabelle A. G. Laureyns, and Kartik B. Ariyur. The Use of Natural Signals for Localization and Navigation with Application to Centimeter sized UAVs. In *American Control Conference (ACC)*, Baltimore, MD, USA, Jun. 2010.
- [255] Francisco Zampella, Mohammed Khider, Patrick Robertson, and Antonio Jiménez. Unscented Kalman filter and Magnetic Angular Rate Update (MARU) for an improved Pedestrian Dead-Reckoning. *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 129–139, 2012.
- [256] William Storms, Jeremiah Shockley, and John Raquet. Magnetic Field Navigation in an Indoor Environment. In *Ubiquitous Positioning Indoor Navigation and Location Based Service*, Kirkkonummi, Finland, 2010. IEEE.
- [257] Isaac Skog, Gustaf Hendeby, and Felix Trulsson. Magnetic-field Based Odometry - An Optical Flow Inspired Approach. In *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Lloret de Mar, Spain, Nov. 2021.
- [258] Peter D.D. Schwindt, Svenja Knappe, Vishal Shah, Leo Hollberg, John Kitching, Li Anne Liew, and John Moreland. Chip-scale atomic magnetometer. *Applied Physics Letters*, 85(26):6409–6411, 2004.
- [259] Svenja Knappe. Microfabricated Optically-Pumped Magnetometers. In *Conference on Lasers and Electro-Optics (CLEO)*, San Jose, CA, USA, Jun. 2014.
- [260] Vishal Shah. QZFM Gen-2 – QuSpin. <https://quspin.com/products-qzfm/>.
- [261] Ingolf V. Hertel and Claus-Peter Schulz. *Atoms, Molecules and Optical Physics 1*. Springer, Berlin, Germany, 2015.
- [262] Daniel Adam Steck. Rubidium 85 D Line Data. <http://steck.us/alkalidata>, 2013.
- [263] Daniel Adam Steck. Rubidium 87 D Line Data. <http://steck.us/alkalidata>, 2001.
- [264] Radwan M. Noor, Nikita Kulachenkov, Mohammad H. Asadian, and Andrei M. Shkel. Study on MEMS Glassblown Cells for NMR Sensors. In *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Naples, FL, USA, Apr. 2019.
- [265] Thad. G. Walker and Michael. S. Larsen. Spin-Exchange-Pumped NMR Gyros. In *Advances in Atomic, Molecular and Optical Physics*, volume 65, pages 373–401. Elsevier, 2016.

- [266] W. Happer, E. Miron, S. Schaefer, D. Schreiber, W. A. van Wijngaarden, and X. Zeng. Polarization of the nuclear spins of noble-gas atoms by spin exchange with optically pumped alkali-metal atoms. *Physical Review A*, 29(6), 1984.
- [267] Thad G. Walker and William Happer. Spin-exchange optical pumping of noble-gas nuclei. *Reviews of Modern Physics*, 69(2):629–642, 1997.
- [268] B. C. Grover. Noble-gas NMR detection through noble-gas-rubidium hyperfine contact interaction. *Physical Review Letters*, 40(6):391–392, 1978.
- [269] S. R. Schaefer, G. D. Cates, Ting-Ray Chien, D. Gonatus, W. Happer, and T. G. Walker. Frequency shifts of the magnetic-resonance spectrum of mixtures of nuclear spin-polarized noble gases and vapors of spin-polarized alkali-metal atoms. *Physical Review A*, 39(11):5613–5623, 1989.
- [270] Claude. Cohen-Tannoudji, J. Dupont Roc, S. Haroche, and F. Laloe. Diverses Resonances. *Revue de Physique Appliquee*, 5(95), 1970.
- [271] Claude. Cohen-Tannoudji. Transverse Optical Pumping and Level Crossing in Free and "Dressed" Atoms. In *Fundamental and Applied Laser Physics: Proceeding of the Esfahan Symposium*, pages 791–815. Wiley, 1971.
- [272] Michael Larsen, Robert Griffith, Michael Bulatowicz, and Philip Clark. Combined Electron Paramagnetic Resonance (EPR) and Nuclear Magnetic Resonance (NMR) Magnetometer System, 2017.
- [273] Felix Bloch. Nuclear Induction. *Physical Review*, 70(7-8):460–474, 1946.
- [274] H. G. Dehmelt. Modulation of a Light Beam by Precessing Absorbing Atoms. *Physical Review*, 6(105):1924–1925, 1957.
- [275] David. A. Van Baak. Resonant Faraday rotation as a probe of atomic dispersion. *American Journal of Physics*, 64(6):724–735, 1996.
- [276] Kishore Padmaraju. Faraday Rotation. *American Journal of Physics*, 67(8), 1999.
- [277] Dmitry. Budker, W. Gawlik, D. F. Kimball, S. M. Rochester, V. V. Yashchuk, and A. Weis. Resonant nonlinear magneto-optical effects in atoms. *Reviews of Modern Physics*, 74(4):1153–1201, 2002.
- [278] D. Sheng, A. R. Perry, S. P. Krzyzewski, S. Geller, J. Kitching, and Svenja Knappe. A microfabricated optically-pumped magnetic gradiometer. *Applied Physics Letters*, 110(3):031106, 2017.
- [279] Anthony P. Colombo, Tony R. Carter, Amir Born, Yuan-Yu Jau, Cort N. Johnson, Amber L. Dagel, and Peter D. D. Schwindt. Four-channel optically pumped atomic magnetometer for magnetoencephalography. *Optics Express*, 24(14):15403, 2016.

- [280] Yu Ji, Jintang Shang, Jin Zhang, and Guoliang Li. Magnetic Gradient Measurement Using Micro-Fabricated Shaped Rubidium Vapor Cell. *IEEE Transactions on Magnetics*, 55(9), 2019.
- [281] Mark. A. Rosenberry, J. P. Reyes, D. Tupa, and T. J. Gay. Radiation trapping in rubidium optical pumping at low buffer-gas pressures. *Physical Review A*, 75, 2007.
- [282] Matthew D. Rotondaro and Glen P. Perram. Collisional broadening and shift of the rubidium D1 and D2 lines ($52S_{1/2} \rightarrow 52P_{1/2}$, $52P_{3/2}$) by rare gases, H₂, D₂, N₂, CH₄ and CF₄. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 57(4):497–507, 1997.
- [283] Nicole Allard and John Kielkopf. The effect of neutral nonresonant collisions on atomic spectral lines. *Reviews of Modern Physics*, 54(4):1103–1182, 1982.
- [284] Greg A. Pitz, Andrew J. Sandoval, Tiffany B. Tafoya, Wade L. Klennert, and David A. Hostutler. Pressure broadening and shift of the rubidium D1 transition and potassium D2 transitions by various gases with comparison to other alkali rates. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 140:18–29, 2014.
- [285] Radwan M. Noor, Mohammad H. Asadian, and Andrei M. Shkel. Design Considerations for Micro-Glassblown Atomic Vapor Cells. *IEEE/ASME Journal of Microelectromechanical Systems*, 29(1), 2020.
- [286] I. A. Sulai, Z. J. Deland, M. D. Bulatowicz, C. P. Wahl, R. T. Wakai, and T. G. Walker. Characterizing atomic magnetic gradiometers for fetal magnetocardiography. *Review of Scientific Instruments*, 90(8), 2019.
- [287] B.S Lunin. Physical and Chemical Bases for Development of Hemispherical Resonators for Solid State Gyroscopes. Technical report, Moscow Aviation Institute, Moscow, 2005.

Appendix A

Microfabricated Optically Pumped Gradiometer

Magnetometers support inertial navigation in several complementary roles. Outdoors, they provide a heading reference by orienting the IMU against the Earth’s field [179, 12] and a position reference through magnetic map-matching, which compares measured spatial variations in the crustal field against known maps such as those of the US Geological Survey [248, 249, 250]¹. Both benefit from a gradiometer, since sensing the field gradient rejects common-mode environmental disturbances and removes temporal variations of the Earth’s field, whose spatial gradient is nearly zero [252, 253]. Navigation-grade heading has been shown to require sensitivity on the order of 20 pT [254]. Indoors, where ambient noise limits absolute heading, magnetometers still aid navigation through magnetic angular rate updates, fingerprinting, and odometry [255, 256, 257]. Since the first microfabricated atomic magnetometer was reported in 2004 [258], such sensors have reached sub-femtotesla sensitivity

¹This appendix is based on the following peer-reviewed publication: A. R. Parrish, R. Noor, and A. M. Shkel, “Microfabricated Optically Pumped Gradiometer with Uniform Buffer Gases,” Virtual Conference. *IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, Virtual Conference, Mar. 2021 [251].

with a scale factor independent of device geometry [259]. Most, however, are built by manual assembly and remain expensive, with early commercial units priced near 10 000 \$ each [260]. The batch-microfabricated gradiometer presented here targets this aiding role at reduced cost and size. To introduce its operation, the following sections review the atomic physics underlying optically pumped sensors, from energy-level structure and optical pumping through the magnetometer dynamics, before presenting the EPR gradiometer.

A.1 Background

First, in order to prepare the reader for the discussion on optically pumped sensors, the underlying physics of a simplified model of an atom is reviewed. Electrons are arranged around a nucleus with discrete shells. Since the nucleus contains charged particles and electrons are also charged particles, their angular momentum creates a magnetic moment. However, when each shell is full, the magnetic moments of the electrons in that shell effectively cancel out. Therefore, for alkali metals, which contain one valence electron, a simplified model of the hydrogen atom can be used to introduce the ideas. The hydrogen atom model is introduced here.

Electrons contain both spin angular momentum ($\vec{S}$) and orbital angular momentum ($\vec{L}$). The spin angular momentum can be thought of as the spin of the electron around its own axis, while the orbital angular momentum is due to its movement around the nucleus. The vector sum of these two is the total angular momentum of the electron $\vec{J} = \vec{L} + \vec{S}$. The nucleus also has inherent angular momentum ($\vec{I}$). Therefore, the total angular momentum of an atom is expressed as

$$\vec{F} = \vec{J} + \vec{I} \tag{A.1}$$

For the rest of this derivation, the focus is on the magnetic moment of the nucleus, which is described as:

$$\vec{\mu}_I = \gamma_I \vec{I} \quad (\text{A.2})$$

where γ_I is the gyromagnetic ratio, which is dependent on the specific atom (number of unpaired protons, nuclear g-factor). The magnetic properties of a magnetic moment are similar to a magnetic dipole. In a static external magnetic field ($\vec{B}_0$) a torque will align the magnetic moment of the nucleus along the field lines. However, as will be discussed in Section A.2, the angular momentum of the nuclei can take different values, causing the magnetic moment of the nuclei to orient either parallel or anti-parallel to the applied field. In addition, due to the angular momentum of the nuclei, they cannot align completely with the applied magnetic field and will precess. This precession occurs according to:

$$\frac{d\vec{\mu}_I}{dt} = \vec{\mu}_I \times \gamma \vec{B}_0 \quad (\text{A.3})$$

Further, this precession occurs at a specific frequency, termed the Larmor frequency:

$$\vec{\omega}_L = \gamma \vec{B}_0 \quad (\text{A.4})$$

When an additional magnetic field is applied ($\delta \vec{B}$) the Larmor frequency will change according to:

$$\vec{\omega}_L = \gamma(\vec{B}_0 + \delta \vec{B}) \quad (\text{A.5})$$

This is a conceptual explanation of the NMR Magnetometer.

A.2 Energy Level Splitting of Nuclei

While only a single atom has been considered so far, practical devices involve many nuclei whose electron configurations are described by the principal (n), secondary (l), magnetic (m_l), and spin (m_s) quantum numbers [261]. For the Rb valence electron $n = 5$, and the relevant sub-shells are $l = 0$ (“S”) and $l = 1$ (“P”). As illustrated in Figure A.1, coupling between the orbital ($\vec{L}$) and spin ($\vec{S}$) angular momenta of the outer electron produces the fine structure, whose total quantum number lies in the range $|l - s| \leq j \leq l + s$ [262].

The nuclear angular momentum also takes discrete values; for the most stable isotopes ^{87}Rb and ^{85}Rb the nuclear quantum number is $i = 3/2$ and $i = 5/2$, respectively. Coupling the electron and nuclear angular momenta (Eq. A.1) produces hyperfine splitting, governed by $|j - i| \leq f \leq j + i$. Finally, an applied magnetic field breaks the spherical symmetry of the atom [261], adding Zeeman sublevels $-f \leq m_F \leq f$ that, for weak fields, are more closely spaced than the hyperfine splitting. The complete energy level splitting for ^{85}Rb is shown in Figure A.1.




Figure A.1: Energy Level Splitting of ^{87}Rb . Adopted from [79].

A.3 Optical Pumping

The goal of optical pumping is to drive the Rb valence electrons into a common m_F state, orienting all nuclei parallel or anti-parallel to the applied field. This is achieved via the D_1 and D_2 transitions shown in Figure A.1: D_1 ($5^2S_{1/2} \rightarrow 5^2P_{1/2}$) corresponds to 794.77 nm in air for ^{87}Rb , and D_2 ($5^2S_{1/2} \rightarrow 5^2P_{3/2}$) to 780.241 nm [263]. On decay the electron emits a photon that could pump neighboring nuclei into extraneous orientations, so N_2 gas is added to absorb it, a process called “quenching” (Sec. A.6.1).

Because circularly polarized light carries angular momentum, absorption transfers $\Delta m_F = \pm 1$ to the electron for σ^+ and σ^- light. As shown in Figure A.2, σ^+ light drives most nuclei to the $m_F = 2$ state, aligning them parallel to the applied field.

Wall and interatomic collisions depolarize the pumped atoms, reducing the polarization “life-time.” Since wall collisions dominate, buffer gases, wall coatings, and spherical glassblown cells are used to extend it [264].




Figure A.2: Optical Pumping of ^{87}Rb using σ^+ polarized beam. Adopted from [79].

A.4 Spin Exchange

The vapor cells also contain ^{129}Xe , whose lower gyromagnetic ratio ($2\pi \times 11.86 \text{ Hz } \mu\text{T}^{-1}$ [265]) and longer relaxation time enable future NMR operation; the Xe polarization is established by spin-exchange transfer of angular momentum from optically pumped Rb [266, 267], and when polarized Xe collides with Rb the sensed field is enhanced by a factor of ≈ 500 due to hyperfine contact enhancement [268, 269].

A.5 Electron Paramagnetic Resonator Magnetometer Dynamics

This model is the basis for the EPR Magnetometer based off of derivations in [79, 270, 271, 272]. The setup for the EPR gradiometer is shown in Figure A.6; however, a model for only one magnetometer is described here. The pump beam is oriented along the z-axis, circularly polarized, and locked to the Rb D1 line (795 nm). The probe beam is oriented along the y-axis, linearly polarized, and slightly detuned from the Rb D2 line (780 nm) (as explained in Sec A.5.2). Since the probe beam is detuned from D_2 resonance, only the effect of optical pumping from the pump beam is considered in this model. In this setup, the net optical pumping is assumed to be strong and in the $\hat{z}$ direction.

The magnetic field for this setup is assumed to be:

$$\vec{B} = \begin{bmatrix} B_x \\ B_y \\ B_0 + B_c \cos(\omega_c t) \end{bmatrix} \quad (\text{A.6})$$

where B_x and B_y are the external magnetic fields the magnetometer is designed to measure, and the z-axis field is to be generated with the three-axis coil in Figure A.6. Since this

application involves many nuclei, their collective behavior is described as the sum of the magnetic moments of the individual nuclei (denoted as $\vec{M}$). The behavior of the nuclei in the magnetometer setup is shown in Figure A.3. The Bloch-like equations describing the dynamics of the magnetic moment are now developed [273]. Three primary phenomena are considered in this model: optical pumping of the nuclei along the $\hat{z}$ direction by the pump beam, thermal relaxation, and the moment exerted by the applied magnetic fields ($\vec{B}$). These phenomena are described by the three terms respectively in Eq A.7 below.




Figure A.3: Principle of operation of EPR magnetometer. Adapted from [79]

$$\frac{d\vec{M}}{dt} = \frac{\vec{M}_0 - \vec{M}}{T_1} - \frac{\vec{M}}{T_2} + \gamma \vec{M} \times \vec{B} \quad (\text{A.7})$$

where $\vec{M}_0$ is the magnetic moment at equilibrium with thermal relaxation when including the impact of only the pump beam, T_1 and T_2 are time constants describing the impact of the pump beam and thermal relaxation, and γ is the gyromagnetic ratio of Rb. Eq A.7 can be simplified by defining the following terms:

$$\frac{1}{\tau} = \frac{1}{T_1} + \frac{1}{T_2}$$

$$\vec{M}_0' = \vec{M}_0 \frac{\frac{1}{T_1}}{\frac{1}{T_1} + \frac{1}{T_2}}$$

where τ is a total time constant describing the combined impact of optical pumping and thermal relaxation. Substituting these terms into Eq A.7 becomes:

$$\frac{d\vec{M}}{dt} = \frac{1}{\tau}(\vec{M}_0' - \vec{M}) + \gamma \vec{M} \times \vec{B} \quad (\text{A.8})$$

This equation can be expanded into the following component form:

$$\frac{d\vec{M}}{dt} = \gamma \begin{vmatrix} \hat{x} & \hat{y} & \hat{z} \\ M_x & M_y & M_z \\ B_x & B_y & B_z \end{vmatrix} - \frac{1}{\tau}(M_x \hat{x} + M_y \hat{y} + (M_z - M_0') \hat{z}) \quad (\text{A.9})$$

Further, separating the equation into individual axes becomes:

$$\frac{dM_x}{dt} = \gamma(B_z M_y - B_y M_z) - \frac{1}{\tau} M_x \quad (\text{A.10a})$$

$$\frac{dM_y}{dt} = \gamma(B_x M_z - B_z M_x) - \frac{1}{\tau} M_y \quad (\text{A.10b})$$

$$\frac{dM_z}{dt} = \gamma(B_y M_x - B_x M_y) - \frac{1}{\tau}(M_0' - M_z) \quad (\text{A.10c})$$

Using the definition $M_+ = M_x + iM_y$, the equations for the x and y axes can be combined as:

$$\frac{dM_+}{dt} = iM_z \gamma(B_x + iB_y) - \frac{M_+}{\tau} - i\gamma B_z M_+ \quad (\text{A.11})$$

Rearranging this equation to match the form of Eq (16) of [270] and defining $M_0'' = i\tau M_z \gamma(B_x + iB_y)$:

$$(\frac{d}{dt} + \frac{1}{\tau} + i\gamma B_z)M_+ = \frac{M_0''}{\tau} \quad (\text{A.12})$$

The steady-state solution to this equation is a Fourier series in carrier harmonics with Bessel-function coefficients $J_n(\omega_1/\omega_c)$, where $\omega_1 = B_c\gamma$ is the Larmor frequency of the carrier field [270]. Retaining only the change in intensity sensed by Dehmelt detection (Sec. A.5.1), taking the real part ($M_x = \Re(M_+)$), and selecting the $n = 1$ resonance ($\omega_0 = -\omega_c$, the negative sign denoting counterclockwise precession about the z-axis) yields [270]:

$$\begin{aligned} \delta M_x^{(n=1)} = \tau M_z \gamma J_1\left(\frac{\omega_1}{\omega_c}\right) & \left(\sum_{p=1}^{+\infty} -B_y \left(J_{1-p}\left(\frac{\omega_1}{\omega_c}\right) + J_{1+p}\left(\frac{\omega_1}{\omega_c}\right) \right) \cos(p\omega_c t) - \right. \\ & \left. -B_x \left(J_{1-p}\left(\frac{\omega_1}{\omega_c}\right) - J_{1+p}\left(\frac{\omega_1}{\omega_c}\right) \right) \sin(p\omega_c t) \right) \end{aligned} \quad (\text{A.13})$$

Observing this signal at the carrier frequency ($p=1$) or its harmonics, B_x and B_y can be extracted from this equation. Observing the $p=1$ harmonic at the carrier frequency and considering only the $\cos(\omega_c t)$ term (this component is selected by demodulating and adjusting the phase) yields:

$$\delta M_x^{(n=1)}_{(p=1)} = -B_y \tau M_z \gamma J_1\left(\frac{\omega_1}{\omega_c}\right) \left(J_0\left(\frac{\omega_1}{\omega_c}\right) + J_2\left(\frac{\omega_1}{\omega_c}\right) \right) \cos(\omega_c t) \quad (\text{A.14})$$

Here B_y is the magnetic field to be sensed, and $-\tau M_z \gamma J_1(\frac{\omega_1}{\omega_c})(J_0(\frac{\omega_1}{\omega_c}) + J_2(\frac{\omega_1}{\omega_c}))$ is a constant in this setup. Therefore, demodulating at the carrier frequency ω_c yields a signal directly proportional to B_y :

$$\delta M_x \propto B_y \quad (\text{A.15})$$

A.5.1 Dehmelt Detection

Dehmelt detection is a method used to sense the magnetization of the Rb nuclei [274]. In this case the probe beam (Figure A.3) is circularly polarized and tuned to the D_1 resonance of Rb. As the nuclei are precessing in plane based on the previous discussion, they will appear to oscillate between pumped and un-pumped atoms when the magnetization vector is facing either towards or away from the probe beam. Therefore, the light intensity of the probe beam after passing through the cell will be proportional to the Rb magnetization in the direction of the probe beam. This setup only requires one photodetector, but pumping the Rb atoms using the probe beam will slightly skew the magnetization vector of the nuclei.

A.5.2 Faraday Detection

If linearly polarized light enters a dielectric medium in the presence of a magnetic field, the polarization axis of the light changes; this is the Faraday effect. The stronger the magnetic field, the greater the change of the polarization axis; this effect is also greatly enhanced near atomic resonance [275] which is why the probe beam is slightly detuned from the Rb D_2 resonance. In this setup, Faraday detection is used to detect the magnetic moment of the Rb nuclei via the probe beam [276, 277].

In the case of the magnetometer, the precessing magnetic moment of the Rb nuclei serves as a magnetized medium, causing Faraday rotation of the probe beam as it passes through the vapor cell. After passing through the vapor cell, the probe beam will pass through a polarizing beam splitter which will separate the light into vertically and horizontally polarized components. The intensity of each component will be measured by two photodetectors. Demodulating the photodetector signals at the carrier frequency (Section A.5) yields a signal proportional to the magnetic moment of the Rb nuclei. This setup requires two photodetectors, increasing the system complexity, but it does not perturb the magnetization of the Rb atoms like Dehmelt detection does.




Figure A.4: Demonstration of the Faraday effect. The axis of polarization of the light is shown to rotate as it passes through a dielectric medium, this change is proportional to the magnetic field. Source: [83, 276].

A.6 Microfabricated Optically Pumped Gradiometer

Magnetometers are widely used in navigation to find the direction of the Earth's magnetic field, thus used as a reference for orienting the inertial measurement unit. It is also possible to use spatial variations in the Earth's magnetic field (owing to the magnetization of the Earth's crust) as a reference for magnetic navigation [249]. However, the magnetic field of the Earth is subject to temporal effects primarily due to current in the ionosphere, which is greater in the event of solar wind. The use of a gradiometer eliminates temporal effects as the spatial gradient for temporal variations is nearly zero [252].

Gradiometers permit rejection of common mode noise in magnetometers, particularly improving operation in environments with large ambient magnetic background. However, gradiometers also permit rejection of laser noise if both cells are filled with the same buffer gas content. If the magnetometer cells were filled with different buffer gas pressures, the Rb atoms would experience different shifts and broadening of the optical absorption line. This reduces the common mode noise rejection in the gradiometer [278]. The presented design

permits uniform buffer gas pressure between the two atomic magnetometers. While multi-channel gradiometers using one cell have been implemented [279] and multi-cell gradiometers have been implemented [278], to the best of the author’s knowledge, only one microfabricated gradiometer capable of multi-axis pumping and probing has been reported [280].

A.6.1 Buffer Gas Content

The use of buffer gases performs several essential functions in optically pumped magnetometers. When Rb nuclei are optically pumped they transition to excited states and will decay (releasing light) to a different ground state - over time the population of certain ground states will be greater due to optical pumping. However, the light released when Rb nuclei decay from the excited state to the ground state can potentially pump other proximal nuclei - this is not desired since they will be pumped to different orientations and it will effectively reduce the pumping rate [281] - this phenomenon is known as radiation trapping. The presence of N_2 gas can mitigate the effects of radiation trapping by absorbing or “quenching” the light released. Other purposes of buffer gases are to make the diffusion length of the gaseous nuclei smaller than the size of the pumping/probing beams for the relaxation time of the nuclei [279]. In addition, the formation of temporary van der Waals particles will increase the chance of spin-exchange between alkali and noble gas nuclei leading to greater polarization of the nuclei - the formation of these particles is only possible in the presence of buffer gases [266]. Finally, buffer gases can improve the relaxation time by minimizing collisions with the cell walls, which are much more likely to depolarize the Rb nuclei than collisions with other Rb or buffer gas nuclei [264].

A.6.2 Line Shape and Pressure Broadening

The presence of buffer gases broadens and shifts the absorption spectra of the Rb nuclei [282, 283]. The addition of buffer gases increases collision rates which increases the distribution of Rb velocities in the cell. Further, the Doppler effect increases/decreases the observed

frequency of the laser light causing frequencies slightly lower to be absorbed by nuclei moving away from the light and frequencies slightly higher to be absorbed by nuclei moving toward the light. For example, the Rb D_1 transition is known to shift by $16.3 \text{ MHz Torr}^{-1}$ in the presence of nitrogen buffer gas content [284].

It is beneficial to have uniform buffer gas content between two magnetometers in a gradiometer configuration. The primary reason for this is the rejection of laser noise in the system. In order to illustrate this, the absorption spectra of two cells, one with pure Rb and another with Rb and 300 Torr of N_2 , are shown in Figure A.5. From this, the broadening and shift of the absorption spectra with the addition of N_2 buffer gas is clear. It can also be concluded that a laser with the same wavelength will be absorbed at different levels between the two cells. Therefore, if the buffer gas content in the two cells is not uniform then frequency and intensity noise from the pump beam will scale differently impeding rejection of common mode noise.




Figure A.5: Experiment to illustrate broadening and shift of Rb absorption spectra with the inclusion of buffer gases. Figure (a) shows the setup for the experiment as the laser is detuned from the D_1 resonance of Rb nuclei, performed on a reference cell containing pure Rb and on the micro cell containing Rb and nitrogen, while (b) shows the results of the absorption spectra [285].

A.6.3 Gradiometer Setup

The setup of the EPR magnetometer is shown in Figure A.6. This setup includes the vapor cell pair, pump and probe beams, multi-layer magnetic shield, and three-axis magnetic field

generation coils. The pump beam is oriented along the z-axis, circularly polarized, and tuned to the Rb D_1 line (795 nm). The probe beam is oriented along the x-axis, linearly polarized, and slightly detuned from the Rb D_2 line (780 nm). This magnetometer is capable of sensing external magnetic fields along the x and y axes (note the NMR magnetometer is only capable of detection along the z-axis).




Figure A.6: Gradiometer implementation. (LP: linear polarizer, QWP: quarter wave plate, Faraday detector: balanced polarimeter to detect Faraday rotation). Each of the magnetometers are represented as A and B.

A.6.4 Operating Principle

The external magnetic field in the X and Y directions is captured because it manipulates the magnetic moment of the Rb nuclei ($\vec{M}$) in a predictable way. Since the magnetic moment of the Rb nuclei can be detected (Section A.5.2), it remains only to describe how the external magnetic field manipulates the magnetic moments. The derivation is described in detail in Section A.5 and the final result is that the magnetic moment of the nuclei along the x-direction will change proportionally to the applied magnetic field. As shown in Equations A.13-A.15 the magnetometer output is directly proportional to B_x or B_y (depending on the phase of demodulation).

A.6.5 Design and Fabrication

The cell network fabrication process, as introduced in [264], is shown in Fig. A.7 and begins by etching 900 μm and through-wafer cavities in a 1 mm thick Si wafer, Fig. A.7(a). Next, anodic bonding seals the etched cavities under atmospheric pressure, Fig. A.7(b). The wafer-stack is then placed in a furnace at 850 $^{\circ}\text{C}$ for 5 min to 7 min, forming into spherically shaped shells due to trapped air inside the cavities, Fig. A.7(c). Following, 100 μm deep channels are etched into backside of the Si wafer using DRIE etching, Fig. A.7(d). A second anodic bonding is performed to install a Rb dispenser pill in an environment of 85 Torr Xe, 45 Torr Ne, and 305 Torr N_2 , Fig. A.7(e). Finally, the Rb is dispensed by heating the dispenser pill using laser heating, Fig. A.7(f). The two glassblown cells are separated by a distance of 5.5 mm and each are probed as magnetometers in this experiment.

A.6.6 Gradiometer Implementation

The table-top gradiometer implementation is shown in Fig. A.6. The gradiometer is housed in a μ -metal shield with integrated three-axis magnetic field coils ($D = 6$ in) and a miniaturized oven made from thermal insulating material. A laser with optical power 4 W is used to




Figure A.7: Description of process flow: (a) DRIE etching of $900\text{ }\mu\text{m}$ and through wafer cavities, (b) first anodic bonding of glass to the etched Si wafer, (c) glassblowing of cells, (d) cell back-side opening and channel definition using DRIE etching, (e) loading Rb dispenser and second anodic bonding in environment of Xe and buffer gases, and (f) dispensing alkali metal through microchannel by laser heating of Rb source.




Figure A.8: (a) Top view of two-cell network with Rb pill shown in middle. (b) Bottom view of mid-fabrication cell network detailing channels connecting both cells enabling buffer gas content to be the same.

heat the cells to a temperature of 140 °C. The pump beam is oriented along the z-axis, circularly polarized, and locked to the Rb D1 line (795 nm). The probe beam is oriented along the y-axis, linearly polarized, and slightly detuned from the Rb D2 line (780 nm). A DC field of 2.5 μ T and an RF field of 13 kHz with an amplitude of 10.5 μ T_{pp} were applied along the z-axis. The magnetometer operates by detecting the projection of the polarization along the y-axis via Faraday rotation of the probe beam. The RF field was used as a reference for the lock-in amplifier to demodulate the output of the Faraday detector. Though the cells contain both Rb and Xe and are therefore capable of NMR operation, the implementation discussed in this paper is purely a Rb magnetometer.

A.6.7 Magnetometer Characterization

The frequency response of each magnetometer was characterized by applying a calibration field of amplitude 35 nT_{pp} with varying frequency swept from 3 Hz to 200 Hz along the y-axis similar to the process described in [286]. This calibration accounts for the mismatch in phase and relative gain between the two magnetometers. While the buffer gases in both magnetometers are identical, the alignment of the pump and probe beams accounts for the disparity in the frequency response of the two magnetometers. The frequency response of each magnetometer is reported in Fig. A.9(a,b). The peaks at 8.8 Hz and 29.5 Hz are due to the presence of ¹³¹Xe and ¹²⁹Xe and correspond to their Larmor frequency in the presence of the 2.5 μ T DC field. The measured sensitivity of the magnetometers is roughly 10 pT/ $\sqrt{\text{Hz}}$ from 3 Hz to 100 Hz.

A.6.8 Gradiometer Characterization

The CMRR of the gradiometer was measured by applying digitally synthesized magnetic noise with RMS amplitude of 70 nT along the y-axis of the setup. Two definitions of the CMRR are given in Eq's A.16,A.17. B is the magnetic field reading of the magnetometers and is defined as $B = S/M$, where S is the raw output of the magnetometer and M is the

calibrated frequency response of each magnetometer. The uncalibrated CMRR in Eq. A.16 [286] quantifies the matching of the relative gain and phase of the two magnetometers:

$$CMRR_{Uncalibrated} = \frac{M_1 B_1 + M_2 B_2}{2(M_1 B_1 - M_2 B_2)} = \frac{S_1 + S_2}{2(S_1 - S_2)} \quad (\text{A.16})$$

The gradiometer output (G), however, corresponds to the calibrated response of each magnetometer: $G = (B_1 - B_2)/L$ where L is the baseline or distance between each magnetometer. An alternative definition of CMRR, given by Eq. A.17, is used to describe the common mode noise rejection of the calibrated gradiometer output.

$$CMRR_{Calibrated} = \frac{B_1 + B_2}{2(B_1 - B_2)} \quad (\text{A.17})$$

The average uncalibrated CMRR was recorded as 72.7 and the average calibrated CMRR was recorded as 85.0 over the 3 Hz to 200 Hz bandwidth.




Figure A.9: Magnetometer and gradiometer characterization: (a,b) describe the relative gain (a) and phase (b) of the frequency response of magnetometers derived from calibration. (c) shows the measured CMRR (data shown in (c) is a moving average of 5 samples).

Appendix B

Lab Procedures

B.1 1700 °C Uniform-Temperature Glassblowing Furnace

The glassblowing process is used in the fabrication of fused silica micro-wineglass resonators. It takes advantage of the drop of viscosity of fused silica material and the increase in pressure of air trapped in the pre-etched cavity at high temperature. If the pressure in the cavity is large enough to overcome the viscosity of material, the device layer will deform and form the shape of a shell.

The equipment that is used for this process is the custom-designed KSL-1700X-Series furnace.

Safety

There are many risks associated with operating this furnace. The user should minimize these risks and be prepared for equipment failure.

The danger is the SiC rods and alumina blocks breaking apart when pulling samples out of the furnace at 1700 °C and falling to the floor. This is why we have cement pavers; if this falls on cement simply let it cool and there will be no issues. This is also why you need to

wear the aluminized lab coat and pants when pulling samples out of the furnace, to prevent you from catching on fire. Over several years of operating this furnace it is very possible it could happen. It could happen to you. Be prepared.

Things you should do before operating the furnace every time:

1. Ensure the cement pavers underneath the furnace are intact and if any high-temperature material were to fall to the floor it would protect the building. Verify that no flammable material is present near the furnace and move it away if necessary.
2. Check the fire extinguisher near the door and ensure it is charged and within the inspection date. Make sure you know how to use a fire extinguisher.

Another critical safety item is not directly associated with the furnace itself, but with when you remove the samples. In the past, samples were taken out of the furnace and immediately placed on a graphite holder because it has high thermal conductivity and will help the sample cool quickly. Freshly blown samples must not be placed directly on graphite. Graphite ignites in air at these temperatures. Instead, as described below, place the samples on old alumina blocks.

Operating Procedure

1. If not already done, install the silicon carbide/alumina sliding rack. Inspect the alumina blocks to see if they will crumble and break soon and replace if necessary.
2. If installing the sliding rack, be sure to install the locking screws for the SiC rods.
3. Flip the thermocouple in the back. This thermocouple will creep over time. You need to flip it periodically or it will be unusable. You can remove it via the two screws with the Phillips screwdriver.

4. Turn on furnace. Flip the slider on the left side, then the black switch in front.
5. Check or edit the profile by pressing the $\leftarrow$ button (A/M), then inspecting the C/T setpoints. After pressing the “A/M” button, it will show “C 01” on top and “0” at bottom. Press the button on the left side, the circle arrow. It should show “t 01” on top and “45” at bottom. This means the initial temperature is 0 °C and it takes 45 min for the temperature in the furnace to get to the next temperature point. Press the button on the left side again. It shows “C 02” on top and the second temperature point at bottom. Press the “A/M” button to select the digit that you want to edit and use up (RUN/HOLD) and down (STOP) buttons to make changes up and down. After you are finished, or if you want to edit a previous setpoint, do not press a button and it will save the profile and return to the initial state of “STOP”. Then, run the program or continue editing the profile from the start by pressing the “A/M” button again.
6. If you are happy with the profile, start heating. First, enable the heating element by pressing the green button “Heating Ready”. Then, start the profile by pressing the down button (RUN) for over 2 s.
7. Furnace will start running current through the heating elements. Loud electrical whining is normal and expected, because the electrical system uses a very large current.
8. **IMPORTANT:** The yellow bar in the temperature controller represents the amount of current. Do not start the profile without pressing the green button “Heating Ready”, as this may damage the furnace.

If the yellow bar is large but no temperature is changing, heating elements might be broken. Stop the profile and press the “Heating off” red button.

9. Do not leave the SiC assembly inside the furnace above 800 °C; it will creep over time.
You can keep the last alumina block just inside to block most of the heat and light from escaping.
10. As the furnace is getting to temperature it is best practice to set up the DSG groups you want to glassblow. Multiple samples may be processed in a single run. Preparation of 30 samples is time-consuming and should be started before the furnace reaches temperature.
11. When the furnace is at temperature, put the controller on hold by pressing the down button (RUN) for over 2s. Controller will say hold.
12. Put on the aluminized lab coat and pants and wear the glasses/goggles.
13. When mounting the samples on the alumina holder, thinner thermal gloves may be worn. Thermal gloves permit approximately 15 s of handling before becoming uncomfortably hot.




Figure B.1: Custom-designed KSL-1700X-Series furnace illustrating sliding rails and samples mounted on SiC rods with alumina blocks to improve temperature uniformity. The top-right image shows samples mounted on a 5 in AlN wafer, prepared for glassblowing.

14. Push samples in, being careful of the bump at the end of the rail.
15. Immediately use tongs to place the alumina holder on an alumina block. The HRGs will stick to the alumina block so you can tilt it slightly and they will not fall.
16. Immediately afterwards use tweezers to move the HRGs off of the alumina holder. This is critical to prevent cooling with huge CTE mismatch which will damage both the DSG and the alumina substrate.
17. Repeat steps 13–17 as needed until all samples are blown.
18. Move the furnace to the next step so it begins cooling. Press the down button (RUN) for over 2 s to resume the program. Controller will say run. Then, press the leftmost button (circle). It will show the current step. Press the up arrow to go to the next step. Leave it like that and it will go back to temperature monitoring, but the step will already be the next one. It should begin lowering temperature. If doing a glassblowing run with multiple temperatures, the process is the same to change in between temperatures.
19. For safety reasons, you can leave after the furnace temperature drops below 1000 °C. You can push the SiC assembly into the furnace at temperatures lower than 800 °C. This will insulate the furnace well, and at 600 °C you can turn it off completely as it will cool slower than the cooling temperature profile.

Temperature Program

The controller is programmed as a sequence of paired segments, where **T** entries specify segment duration in minutes and **C** entries specify the target temperature reached at the end of that segment. Ramp rate for a segment is the temperature change divided by the segment time.

Table B.1: Temperature program for the 1700 °C glassblowing profile. The manufacturer recommends ramping rates below 5 °C min⁻¹; in practice the furnace sustains substantially higher rates during glassblowing without issue.

Segment	Input Data	Description
C 01	0	Initial temperature. You can leave it as “0”.
T 01	45	Running time in the first segment, e.g., ramping time to reach temperature C02 (450 °C).
C 02	450	Temperature to be reached within T01. C02/T01 is the ramping rate in the first segment, 10 °C min ⁻¹ .
T 02	30	Running time in the second segment.
C 03	450	Temperature to be reached in the second segment. If C02 is the same value as C03, it means dwell time in the second segment.
T 03	75	Running time in the third segment.
C 04	1000	Temperature to be reached in the third segment. Ramping rate is 7.3 °C min ⁻¹ .
T 04	160	Running time in the fourth segment.
C 05	1700	Temperature to be reached in the fourth segment. Ramping rate is 4.4 °C min ⁻¹ .
T 05	120	Running time in the fifth segment.
C 06	1700	Temperature to be reached in the fifth segment. Because C06 is the same value as C05, it means dwell time; glassblowing is performed during this segment.
T 06	160	Running time in the sixth segment.
C 07	1000	Temperature to be reached in the sixth segment. Because C07 < C06, it means cooling in this segment. Cooling rate is 4.4 °C min ⁻¹ .
T 07	80	Running time in the seventh segment.
C 08	400	Temperature to be reached in the seventh segment. Cooling rate is 7.5 °C min ⁻¹ .
T 08	-121	Program finish and stop. The furnace will cool naturally.
C 09	0	Terminal state after natural cooling.




Figure B.2: Programmed temperature profile for the 1700 °C glassblowing furnace, annotated with heating and cooling rates for each segment. Glassblowing is performed during the 120 min dwell at 1700 °C, highlighted in orange. Total programmed time is 670 min, followed by natural cooling after program stop ($T_{08} = -121$).

Profiles

Substrates

Experimentally the best substrate is alumina, though AlN and SiC were also studied. AlN withstood the thermal profile, but bonded to the fused silica immediately after cooling; attempting to separate the two shattered the DSG. SiC performed well, but is expensive and available only in thin substrates. Alumina is relatively inexpensive and was observed to stick to the DSG, but because alumina is soft it peels away and leaves the DSG undamaged. Residual alumina is easily removed during lapping release.

B.2 Lapping/Release Traveler

This traveler describes the lapping and release of DSG samples following glassblowing. Material is removed sequentially with progressively finer abrasive films to open the first cavity, then the second cavity, releasing the inner shell. The final polishing steps are followed by a cleaning sequence to remove slurry residue and lapping wax from the released shell.

Before starting, the user should know the shell geometry and its dimensions (see Fig. B.6).

Lapping Procedure

1. Attach the non-magnetic platen and align it.
 - (a) Ensure the load setting is set to full prior to alignment.
2. Place mounting wax at the four corners of the mounting fixture.
 - (a) Use as little wax as possible while still securing the sample.
 - (b) Push down for at least 15 s to 20 s after the wax has melted to ensure it is cooled before releasing.
 - (c) Align the center of the sample with the center of the fixture and melt the wax using the heat gun at 400 °F. Press down tightly when the wax is melted. Let cool for ≈ 5 min before lapping.
3. Use Dia-Grid diamond discs (grit size 80 μm) to remove approximately 400 μm of material. This is intended to open the first cavity.
 - (a) Speed: 100 rpm. Load: full.
 - (b) If the stem is very thin, switch to 30 μm diamond grit film. Speed: 75 rpm. Load: 400 g. If the stem is thick, continue lapping with Dia-Grid diamond discs, but

reduce the depth pushed down. The first cavity will become very thin and must not be broken, which would damage the inner shell.

- (c) Continue lapping until the cavity is opened enough to place wax particles inside (see Fig. B.3).

4. Fill the first cavity.

- (a) Dry the cavity using the heat gun.
- (b) Put large blocks of solid wax onto the side of the cavity, then melt and blow the wax into the cavity using the heat gun at 400 °F.
- (c) Ensure the cavity is filled completely and there are no bubbles inside. Use a blade to flatten the wax film before cooling.

5. Continue to use Dia-Grid diamond discs until 500 μm before the second cavity is opened (design dependent).

- (a) Initially use small force until the wax has been lapped off and is coincident with the glass; the surface should be smooth. Wax laps away rapidly.
- (b) Following this, Speed: 100 rpm. Load: full.

6. Switch to the non-magnetic platen, calibrate the machine, and change the lapping film to 30 μm grit size. Continue lapping until 200 μm before the notch (see Fig. B.7).

- (a) Speed: 100 rpm. Load: full.
- (b) Lap until 200 μm visually from the notch.
- (c) Use the 30 μm film only twice; after this the grit is used up.

7. Lap with 15 μm film until the second cavity is barely opened (see Fig. B.5).

- (a) Calibrate before starting (at least for the final step).
 - (b) Use new 15 μm film for this step every time. Use this same film in the following step as well.
 - (c) Speed: 75 rpm. Load: 400 g. Push down a maximum of 50 μm .
 - (d) Frequently check (every 5 min) whether the cavity is opened; this is difficult to do by eye, so use a microscope.
8. Fill the second cavity with lapping wax using the same procedure as above (see Fig. B.4).
- (a) First, dry the cavity using the heat gun.
 - (b) Blow wax over the cavity. Continue reflowing around the entire edge, ensuring no bubbles are formed near the surface; bubbles deeper in the cavity are acceptable.
9. Next, re-apply wax to the first cavity to ensure it is flat, then smooth the entire wax surface with a box cutter.
10. Remove the sample from the holder and remove the wax from the holder using a box cutter. Reapply wax and lightly push down the corners (at least 15 s) to ensure it is flush with the holder surface.
- (a) Examine the first and second cavities to ensure they are completely filled and flat.
11. Switch the lapping film to 15 μm grit size. Lap $\approx 50 \mu\text{m}$ depending on the design.
- (a) Initially use old 15 μm film and lap only the wax until the sample is flush.
 - (b) Switch back to new film. Speed: 75 rpm. Load: 400 g.

- (c) Frequently inspect using the Leica to observe the second cavity width to avoid overlapping. Compare with the width of the model in COMSOL after lapping the desired distance.
 - (d) $50\text{ }\mu\text{m}$ is near the maximum lapping depth for this step; the achievable depth depends on the stopping point of the previous step, for example lap $30\text{ }\mu\text{m}$ instead.
12. Switch the lapping film to $6\text{ }\mu\text{m}$ grit size and lap $30\text{ }\mu\text{m}$. Compare the resulting surface with Fig. B.8.
- (a) Speed: 20 rpm. Load: 400 g.
 - (b) Directly push down $30\text{ }\mu\text{m}$, then lap for approximately 1 h.
13. Remove the lapping film and place Final B rayon flock on the platen. Pour $3\text{ }\mu\text{m}$ OpTech polishing slurry on the polishing cloth and turn on the lubricant (GreenLube).
- (a) The only purpose of using the lapper at this point is to polish.
 - (b) First, bring the sample into contact with the cloth. Then, pull up $150\text{ }\mu\text{m}$. Gradually lower from $150\text{ }\mu\text{m}$ and set this manually. Typically, $\approx 25\text{ }\mu\text{m}$ above the point where spindle rotation is impeded.
 - (c) Prime the rayon flock cloth with $3\text{ }\mu\text{m}$ slurry manually.
 - (d) Use the automatic dispenser with the following settings:
 - i. GreenLube: 4 dispense cycles per minute, 0.5 s pulse length.
 - ii. $3\text{ }\mu\text{m}$ diamond slurry: 1 dispense cycle per minute, 0.5 s pulse length.
 - (e) Manually apply slurry or brush it back into the path every $\approx 10\text{ min}$ and check the sample every $\approx 30\text{ min}$.

- (f) Polish until the remaining pockets are shallow and polishing further has no effect (approximately 2 h).
14. Replace the Final B polishing cloth. Pour 0.5 μm cerium oxide suspension on the cloth and again turn on the lubricant (GreenLube).
 - (a) Use the same method to determine the proper height as in the previous step (start 150 μm up). Also prime the cloth as above.
 - (b) Use the automatic dispenser with the following settings:
 - i. GreenLube: 4 dispense cycles per minute, 0.5 s pulse length.
 - ii. 0.5 μm cerium oxide slurry: 1 dispense cycle per minute, 0.5 s pulse length.
 - (c) Frequently check the surface quality with the Leica and compare with Fig. B.9. Expect 30 min to 60 min on this step.
 15. Once finished, remove the sample from the mount using the heat gun. Spray gently with acetone, avoiding direct spray onto the inner shell. Next, soak in a clean acetone bath for a minimum of 1 h, or overnight.
 16. Put the shell in two isopropyl alcohol baths for several minutes. Put the shell through >3 changes of DI water. Soak the shell in a solution of DI water with micro-soap (10:1) for 10 min.
 17. Prepare several cleanroom Q-tips. While in the soapy solution, gently grip the shell and scrub off residual slurry. The slurry residue will likely be on the outside of the outer shell; do not touch the inner shell at all. Also, be sure to scrub the sides of the wafer stack in case slurry is trapped inside the stack.

18. Place the shell in a solution of DI water and soap in the ultrasonic machine and turn it on, with heating enabled, for 20 min to 30 min. Put the shell through >3 changes of DI water.
19. Place the shell in a bath of DI water. Prepare a compressed airline so it is very gently blowing air; too much will damage the shell. Then, dip the shell in and out of the water while blowing air through the second cavity. The goal is to remove as much of the liquid trapped in the cavity as possible. Repeat this several times with fresh DI water baths.
20. This step is optional. Place the shell in acetone solution and put it in the ultrasonic shaker for 1 h, or overnight. Then, place it in at least two isopropyl alcohol baths and then >3 DI water baths.
21. Prepare piranha solution (either in the INRF or in the laboratory) and soak the shell for ≈ 30 min. Put the shell through >3 changes of DI water.
22. Repeat step 19 to remove piranha solution from the second cavity.
23. Dry the shell on a hotplate at $>150^\circ\text{C}$ inside a clean glass container for at least 1 h.

Subsurface Damage Calculations

The depth of SSD is critical to track to ensure the inner shell is damaged as little as possible. The SSD depends on the grit size of the abrasive layer and is characterized by [117]:

$$\text{SSD depth} < 2L^{0.85} \tag{B.1}$$

where L is the grit size in μm , and SSD depth is the depth of subsurface damage in μm . This equation is valid for any material and feed rate.

Table B.2: Subsurface damage upper bound as a function of material grit size.

Material grit size [μm]	SSD upper bound [μm]
Dia-Grid diamond disc	108
35	41
30	36
15	20
6	9.2
3	5.1
0.5	1.1

Reference Images




Figure B.3: Initial opening of the first cavity. Image courtesy of W. Guan.




Figure B.4: First and second cavity filled with wax. Image courtesy of W. Guan.




Figure B.5: Side-view of samples in different stages of lapping under microscope.




Figure B.6: (a) View of dual-shell after blowing and no lapping is performed, for reference. (b) On right are the cavities from Fig. B.5 labeled and left is (a) with labels explaining the viewpoint.




Figure B.7: Illustration of sample and corresponding lapping depths for each film. View is the same as Fig. B.5. The notch which is referred to in the instructions is also labeled.




Figure B.8: Leica microscope images of surface quality of dual-shell. (a) has several deep scratches. (b) is a sample after $6\mu\text{m}$ lapping is completed; all deep scratches have been eliminated.




Figure B.9: Leica microscope images of surface quality of dual-shell. (a) shows surface after $3\mu\text{m}$ polishing slurry; it was found further polishing with this slurry cannot improve the surface quality. (b) shows excellent surface quality, the desired outcome after $0.5\mu\text{m}$ polishing.

B.3 Annealing Furnace

This process is used in the annealing of fused silica micro-wineglass resonators, but is generally applicable to all fused silica glass. Internal stress in fused silica micro-wineglass resonators is primarily introduced by the rapid temperature changes the shell experiences during the glassblowing process. Annealing relaxes the microstructure of the glass, allowing these internal stresses to dissipate. The process is bounded by three characteristic temperatures of Corning 7980 fused silica [115]: the strain point (893°C), the annealing point (1042°C), and the softening point (1585°C). The annealing temperature must exceed the strain point, and higher temperatures require exponentially less time to relax internal stresses; the cooling rate must then be slow enough that the thermal gradient through the shell does not reintroduce stress. A temperature just above the strain point is generally selected so that the microstructure equilibrates at a low temperature and is locked in without reintroducing thermal stress [287].

B.3.1 Equipment

More information can be found in the Thermo Scientific Lindberg/Blue M Moldatherm 1100°C Box Furnace Manual, Thermo-Fisher Document Number 330985H01.

The operator interface is shown in Fig. B.10. It is used extensively in the procedure.



Figure B.10: Operator interface for the annealing furnace controller.

B.3.2 Operating Procedure

1. Before turning the furnace on, prepare the sample in the furnace chamber. In order to prevent contamination, it is beneficial to place the sample on a fused silica wafer. Once the sample is prepared, close the furnace door before proceeding.
2. Plug in the furnace. Flip the black switch on the front-right of the furnace.
3. An alarm will likely occur immediately. This alarm is triggered when the furnace temperature is far below the setpoint, which is common at startup. Clear the alarm.
4. Enter Level 2 and select the desired profile, following the procedure in Section B.3.3. This allows the temperature profile to be viewed and changed rather than only executed.
5. Review the temperature setpoints “TSP.1” – “TSP.8”, ramp rates “RMP.1” – “RMP.8”, and dwell times “DWEL.1” – “DWEL.8” for each segment to confirm the profile matches the intended values.
6. Start the program.
7. While the temperature profile is running, the elapsed time (“T.ELAP”) and time remaining (“T.REMN”) for a specific segment can be accessed.

B.3.3 Profiles

This furnace is equipped with 16-segment controller operation, which allows arbitrary temperature profiles to be defined. The manual counts each ramp and each dwell as a separate segment. In practice the profile is more accurately described as 8 segments, each with an associated ramp rate and dwell time.

Editing Temperature Profiles

1. Enter Level 2 by holding **PAGE** until “Lev 1 GOTO” is displayed, selecting “Lev2,” and entering the password 25.
2. Select the desired profile number under the “PROG” parameter, then return to the home page. The temperature parameters will populate for this profile.
3. Set the following parameters:
 - “TSP.1” – “TSP.8” : target setpoint for segments 1–8.
 - “RMP.1” – “RMP.8” : ramp rate for each segment.
 - “DWEL.1” – “DWEL.8” : dwell time in min for each segment.

Example Temperature Profile

An example annealing profile is shown in Fig. B.11 and described below:

- **Segment 1** – Step to 700 °C followed by a 6 h hold to ensure uniform temperature before proceeding.

Setting the ramp rate to OFF makes this a step function and the target temperature will immediately change to 700 °C.

- **Segment 2** – Long soak at 920 °C for 30 h.
- **Segment 3** – Very slow cooling at -0.2 °C min^{-1} to 650 °C followed by a 1 h soak to ensure the temperature of all samples is uniform.
- **Segment 4** – Ramp down to 15 °C, a setpoint below room temperature, ensuring the heating elements turn off and the furnace cools naturally.

Turning the ramp rate OFF makes this a step function.

- **Segments 5–8** – The remaining segments are unused; the heating elements remain off.




Figure B.11: Example annealing temperature profile.

End Type Parameter

1. Set the “END.T” parameter to “SP2”. After the temperature profile is complete, the temperature setpoint will change to the value currently under “SP2”.
2. Set the “SP2” parameter to 15 °C. This setpoint is below room temperature, so the heating elements remain off and the furnace cools naturally after operation.

B.3.4 Precautions

The temperature of the annealing furnace can get very hot. At 1000 °C or higher, wear safety glasses if it is essential to open the furnace door (this should not happen during normal operation).

Also note that unlike other furnaces, the heating element can be turned off at any time, even at high temperatures. This furnace is well insulated, and the temperature will decrease slowly regardless.

Separately, opening the door at high temperatures for prolonged periods can damage the heating elements.

B.4 Automated PAE Characterization of Dual-Shell Resonators

Scope and Objective

This traveler describes the procedure for measuring the PAE of shell resonators inside a vacuum chamber¹. It covers hardware setup, measurement execution using a LDV, and automated control of the LDV through MATLAB for mode shape acquisition.

The objective is to determine the PAE and capture mode shapes, which are required for the subsequent trimming process. The resonator is locked to the lower frequency of the $n = 2$ wineglass mode, and the Q-factor and phase of that mode are measured. These parameters are supplied to the MATLAB control script, which moves the LDV across points on the shell rim to acquire the mode shape.

The user should be familiar with the LDV and the Zurich lock-in amplifier before attempting this procedure, and should be able to take a frequency sweep, lock a PLL, and measure the Q-factor of a DSG.

Software Initialization

The three stage axes are homed and zeroed in the Micronix Motion software each time the system is power cycled; this establishes the (0,0) reference and is mandatory before acquisition. If a z -axis travel-limit error appears, the negative travel limit is set to a value less than -6 . The Micronix software port must then be closed before MATLAB is opened, because the stage accepts only one active connection and the MATLAB stage functions will otherwise fail to connect.

¹This procedure was developed by the author. S. Salmani contributed to the detailed documentation of the procedure. The mount for the stage and first iteration of the MATLAB code were developed with the help of an undergraduate senior design team comprised of Matthew Delmont, Zak Klutts, and Spencer Herbert.

Instrumentation Check

Confirm that the LDV and the vacuum chamber are connected and powered before beginning the measurement loop. The LDV bandwidth must be set above 100 kHz; a bandwidth set too low for the resonator under test is a common source of error.

Automated Acquisition

The acquisition is handled by the `Master.m` script, which controls the full experimental workflow. The initialization section of the script identifies the instruments and connects to three components: the Micronix positioning stage, the Zurich lock-in amplifier, and an oscilloscope that reads the LDV power bar voltage to confirm that the laser is focused on the surface. A reading of 2.3 V corresponds to full focus.

The stage connection routine closes any prior connection before initializing a new one. As noted in step B.4, the Micronix software port must be closed for this to succeed.

Variable definition and z -axis calibration The shells were originally intended to sit on a printed mount, but this configuration proved unstable, and the shell is now mounted directly on the stage with carbon tape. To account for the resulting height difference, `z_base` is set to -4 , placing the base coordinate at $(0, 0, -4)$.

The test radius `r_test` depends on the size of the resonator; a value of 3.2 is used here as a reference point.

Rim mapping and focus optimization The rim is located by confirming four points on it. The LDV is first centered manually on the stem with the stage at $(0, 0, z_base)$ and the laser focused there, because the mapping algorithm moves the stage radially by `r_test` relative to this initial position and an off-center start offsets all subsequent movements. At each of the four estimated rim points the laser is centered on the rim width and the




Figure B.12: Shell resonator mounted on the positioning stage using carbon tape.

point approved, with physical centering taking priority over absolute signal magnitude. The flat edge is then mapped from two points at either end to determine the orientation. The resulting coordinates are recorded so that subsequent runs skip the manual mapping, and the circle fitting, plane fitting, and edge-angle estimation that follow are automated.

Final setup and PLL locking The stage is moved to the first recorded circle point so that the laser sits on the rim, then proceed as follows.

1. **Find resonance.** Perform a frequency sweep on the lock-in amplifier to identify the peaks. If the laser sits at a node of the mode of interest, only one peak may appear; if two peaks are not observed, move the LDV to a different point on the rim and sweep again.
2. **Lock the PLL.** Lock to the lower frequency of the $n = 2$ mode. Locking to the lower frequency is critical: locking to the higher frequency causes the subsequent trimming process to increase the frequency split Δf rather than reduce it.

3. **Maximize the signal.** Adjust the PLL phase setpoint until the amplitude is maximized.
4. **Record parameters.** Measure the resonance frequency, the phase, and the decay time from the lock-in amplifier.

The measured values are entered in the mode shape characterization section of the script. No variable is defined for the decay time; instead, the measured decay time τ is inserted directly into the Q-factor calculation, $Q = \pi f \tau$, replacing the placeholder value.

Execution and monitoring Once all variables are updated, run the script. MATLAB plots the mode shape data in the figure window as the measurement proceeds. Confirm that the LDV moves to every point on the rim and that the lock-in amplifier checks the PLL lock at each step.

Because the locations of the nodes and antinodes are initially unknown, the script scans all points on the rim. A 180° phase flip and a corresponding amplitude drop are observed as the measurement passes through a node. For the $n = 2$ mode, the antinodes are spaced 90° apart.

Appendix C

MATLAB Code for Dual-IMU Constraint

The code implementing the dual-IMU framework of Chapter 8 is publicly available through the MicroSystems Laboratory GitHub: [\[GitHub\]](#). The repository contains the joint error-state EKF with the zero-position-difference constraint, the per-IMU stance-phase detectors, and the random forest gait classifier.

All code can be found under the MicroSystems Laboratory organization if the link above does not resolve: [\[GitHub\]](#).

Appendix D

List of Vendors

Airex Corporation (P15-1 Series Linear Motor Coil Assembly)

Phone: +1 (918) 627-1845

Website: <https://airex.com/>

The P15-1 Series Linear Motor Coil Assembly from Airex Corporation was used as the linear motor actuator in the FLICE system. The assembly provides precision linear electromagnetic actuation suitable for controlled motion applications.

Allied High Tech Products, Inc.

Address: Rancho Dominguez, CA

Phone: +1 (800) 675-1118

Website: <https://www.alliedhightech.com/>

Email: info@alliedhightech.com

Allied High Tech Products, Inc. provides die-level and wafer-level lapping and polishing instruments for the release of DSG, along with lapping-related products such as lapping films, slurry, lubricants, and lapping fixtures. Maintenance and repair services for the Multiprep Polishing System are also available through the vendor.

Corning Incorporated

Address: 1 Riverfront Plaza, Corning, NY 14831

Phone: +1 (607) 974-9000

Website: <https://www.corning.com/>

Corning Incorporated produces the 7980 high-purity fused silica glass used as the substrate, device, and cap wafer material for the dual-shell resonators described in this dissertation.

Dell Inc.

Address: 1 Dell Way, Round Rock, TX 78682

Phone: +1 (800) 289-3355

Website: <https://www.dell.com/en-us>

Computers, laptops, and monitors were purchased from Dell.

Digi-Key Corp

Address: One Technology Way, PO BOX 9106, Norwood, MA 02062

Phone: +1 (800) 344-4539

Website: <http://www.digikey.com/>

Retails electronic components and supplies various development boards.

Front Range Photomask Co. LLC

Address: Lake Havasu City, AZ

Phone: +1 (928) 733-6217

Website: <https://www.frontrange-photomask.com/>

Email: masks@frontrange-photomask.com

Front Range Photomask Co. LLC provides glass photomask fabrication services with fabrication accuracy ranging from 0.2 μm to 1 μm .

Ideal Aerosmith, Inc.

Address: 2205 W Lone Cactus Dr Suite 7, Phoenix, AZ 85027

Phone: +1 (701) 757-5601

Website: <http://www.ideal-aerosmith.com/>

Precision rate table systems were used to characterize inertial sensors described in this dissertation.

JLCPCB

Address: 2788 San Tomas Expy, Santa Clara, CA 95051

Phone: +1 (800) 262-5643

Website: <https://jlcpcb.com/>

Customized PCBs used on the Sugar-Cube platform were fabricated by JLCPCB.

Kurt J. Lesker Company

Address: 1925 Route 51, Jefferson Hills, PA 15025

Phone: +1 (800) 245-1656

Website: <https://www.lesker.com/>

Kurt J. Lesker Company provides supplies for Physical Vapor Deposition (PVD), including deposition materials such as Au, Cr, and Ti, as well as components and supplies for vacuum chambers and pumps.

Laird Connectivity (BL653u Bluetooth Module)

Address: W66N220 Commerce Court, Cedarburg, WI 53012

Phone: +1 (262) 375-4400

Website: <https://www.lairdconnect.com/>

The BL653u Bluetooth Low Energy (BLE) module from Laird Connectivity was used for wireless communication in the Sugar-Cube platform. It provides integrated Bluetooth 5.1 connectivity for embedded systems and supports low-power data transmission between sensor modules and external devices.

Mark Optics, Inc.

Address: Santa Ana, CA

Phone: +1 (714) 545-6684

Website: <https://www.markoptics.com/>

Email: quote@markoptics.com

Mark Optics, Inc. is a custom optics provider specializing in glass wafers. Available materials include Pyrex, Fused Quartz, Borosilicate Glass (BSG), and Ultra Low Expansion Titania Silicate Glass (ULE TSG). Wafer polishing services for etched wafers are also available through Mark Optics, Inc.

UC Berkeley Marvell Nanofabrication Laboratory

Address: 520 Sutardja Dai Hall, Berkeley, CA 94720

Phone: +1 (510) 809-8600

Website: <https://nanolab.berkeley.edu/>

Email: michael1@eecs.berkeley.edu

Marvell Nanofabrication Laboratory at UC Berkeley is a cleanroom facility providing microfabrication services. Specifically, double-side 2 μm doped Poly-Si LPCVD is used as the hard mask during HF wet etching of fused silica wafers.

MathWorks, Inc.

Address: 1 Apple Hill Drive, Natick, MA 01760

Phone: +1 (508) 647-7000

Website: <https://www.mathworks.com/>

MATLAB was used for data processing, algorithm development, simulation, visualization, and analysis of inertial navigation, sensor fusion, and experimental results presented in this dissertation.

Micronix USA

Address: 11555 Coley River Circle, Suite A, Fountain Valley, CA 92708

Phone: +1 (949) 480-0538

Website: <https://www.micronixusa.com/>

Micronix USA provides precision motion control systems including the PPS-20-11206-AtomRKLF piezoelectric motor and the MMC-100-03100 motor controller. These components were used in the vacuum chamber for automated characterization of the PAE system.

MTI Corporation (KSL-1700X-Series Furnace)

Address: 860 South 19th Street, Richmond, CA 94804

Phone: +1 (510) 525-3070

Website: <https://www.mtixtl.com/>

Email: info@mtixtl.com

The KSL-1700X-Series furnace from MTI Corporation was used as the basis for the custom glassblowing furnace developed as part of this dissertation. The alumina blocks for insulation and to mount the samples for glassblowing are also purchased from MTI.

PI Ceramic GmbH

Address: Lindenstrasse, 07589 Lederhose, Germany

Phone: +49 36604 882-0

Website: <https://www.piceramic.com/>

Email: info@piceramic.de

Ceramic piezoelectric actuators from PI Ceramic were used to drive the DSG into resonance for the principal axes of elasticity characterization setup.

PJRC (Teensy 4.0)

Address: 2230 N Pacific Ave, Unit 101, Santa Cruz, CA 95060

Phone: +1 (831) 420-1435

Website: <https://www.pjrc.com/>

The Teensy 4.0 microcontroller from PJRC was used in the Sugar-Cube platform for control and data acquisition.

Plastronics

Address: 2010 Westridge Dr, Irving, TX 75038

Phone: +1 (972) 258-2580

Website: <https://www.plastronics.com/>

Email: JDaniel@plastronics.com

Plastronics offers complete and reliable burn-in test socket solutions for the latest packaged semiconductor devices, including LCC and PLCC packages. For LCC44 packages, the socket part number used is P2044S-B-AU.

Polytec, Inc.

Address: 16400 Bake Parkway, Suites 150 & 200, Irvine, CA 92618

Phone: +1 (949) 943 3033

Website: <https://www.polytec.com/us/>

Polytec, Inc. provides Laser Doppler Vibrometer (LDV) systems used for characterization of resonators, including measurements of frequency split, Q-factor, displacement, and piezoelectric properties such as d_{33} characterization.

SparkFun Electronics

Address: 6333 Dry Creek Parkway, Niwot, CO 80503

Phone: +1 (303) 284-0979

Website: <https://www.sparkfun.com/>

Electronic breakout boards, jumper wires, lithium-ion batteries, voltage regulators, barometers, IMU, UWB, and microcontrollers, used on the Sugar-Cube were purchased from SparkFun Electronics.

TDK InvenSense, Inc.

Address: 1745 Technology Dr 200, San Jose, CA 95110

Phone: +1 (408) 988-7339

Website: <https://invensense.tdk.com/>

IMU ICM-42605, Barometer ICP-20100, and Magnetometer AK-09915 were integrated into the Sugar-Cube platform described in this dissertation.

Thermo Fisher Scientific

Address: 168 Third Avenue, Waltham, MA 02451

Phone: +1 (781) 622-1000

Website: <https://www.thermofisher.com/>

Thermo Fisher Scientific provides laboratory equipment and instrumentation used in this work, including annealing furnaces, hot plates, and general-purpose laboratory tools.

Thorlabs, Inc.

Address: 56 Sparta Avenue, Newton, NJ 07860

Phone: +1 (973) 300-3000

Website: <https://www.thorlabs.com/>

Thorlabs, Inc. provides precision optical and optomechanical components, including rotation and translation stages used for optical alignment in the Multi-LDV setup described in this dissertation.

VectorNav Technologies

Address: 10501 Markison Rd, Dallas, TX 75238

Phone: +1 (512) 772-3615

Website: <https://www.vectornav.com/>

VectorNav VN-200 IMUs used in pedestrian inertial navigation in this dissertation are available from this vendor.

Zurich Instruments USA, Inc.

Address: 400 5th Avenue, Suite 115, Waltham, MA 02451

Phone: +1 (855) 500 0056

Website: <https://www.zhinst.com/>

Email: info@zhinst.com

Zurich Instruments is a test and measurement company providing advanced instruments such as lock-in amplifiers (including the HF2LI and UHF2LI), used in the characterization of inertial sensors described in this dissertation.