



University of
Strathclyde
Glasgow

Unmanned Aerial Vehicles for
Contact-Based Inspection

Robert Watson

Department of Electronic and Electrical Engineering

Centre for Ultrasonic Engineering

University of Strathclyde

A thesis submitted for the degree of

Doctor of Philosophy

2023

Declaration & Copyright

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed:

A handwritten signature in black ink, consisting of the letters 'R', 'S', and 'W' in a cursive style.

Date: *4th June 2023*

“Remember kids: the only difference between screwing around and science
is writing it down.”

*Adam Savage,
on the set of MythBusters, 2012*

Acknowledgements

First, to my academic supervisors Charles MacLeod, S. Gareth Pierce, and Gordon Dobie, thank you for your continuous support and encouragement throughout this studentship and your patience with my drive to always do better. The unique opportunities you have opened to me made this thesis possible and I thoroughly enjoyed working through each of them.

Thank you to my industrial supervisors Gary Bolton, Tom Carey, Tariq Dawood, Antoine Joly and your colleagues in the National Nuclear Laboratory and renewables team at the EDF R&D UK Centre. I am most grateful for the warm welcomes you offered me and the experience I gained amongst such bright and enthusiastic teams. I also remain appreciative for the funding provided through the Engineering and Physical Sciences Research Council (EPSRC) and the Industrial Cooperative Awards in Science & Technology (iCASE) program.

To Dayi Zhang, Taiyi Zhao, William Jackson, Momchil Vasilev, Euan Foster and all the other friends and colleagues I have worked with in the Centre for Ultrasonic Engineering (CUE) and the Sensor Enabled Automation Robotics and Control Hub (SEARCH), thank you for assistance, insight, and making the long hours in the lab as enjoyable. Particular thanks also to Walter Galbraith and Dave Lines for your sage-like advice regarding all things electronic and ultrasonic when conventional sources ran dry.

To my parents and my friends outside academia, thank you for your love, support, the many luxurious distractions you offered, and the welcome respite they brought. I look forwards to seeing more of you.

Lastly, to my brother and long-suffering flat mate Euan, thank you most of all for your unwavering faith that I would eventually finish this thesis. You will be happy to know my days of complaining and neglecting housework are (almost) over.

Abstract

Multirotor Unmanned Aerial Vehicles (UAVs) have become effective tools for in situ inspection of energy sector infrastructure where aging-related degradation causes critical failures but asset scale and inaccessibility make frequent manual assessment cost-prohibitive. Minimising human access to wind, nuclear and oil and gas facilities through rapid visual screening, UAVs have drastically reduced the associated risks of work at height, radiation exposure, and hazardous atmospheres. To further improve structural insight and cut downtime, this thesis examines airborne deployment of other established Non-Destructive Evaluation (NDE) methods via physical surface contact.

An over-actuated multirotor deploying a dry-coupled ultrasonic wheel probe is developed as a novel thickness mapping strategy. Using bi-axial tilting propellers in a unique thrust-vectoring tricopter layout, this system may efficiently apply interaction forces from omnidirectional flight. Through laboratory testing, stable and repeatable inspection is then characterised in various representative operations. Against a reference sample mounted vertically or beneath a 45° overhang, static point-thickness measurements consistently show mean absolute error under 0.10 mm. This error also remains below 0.28 mm in dynamic rolling measurements suited to area thickness profiling, highlighting successful airborne delivery of multi-modal ultrasonic testing.

A novel multirotor-crawler hybrid vehicle is also developed for contact-based assessment of pipes, tanks, and other cylindrical assets. Utilising a multidirectional propeller array under energy-optimised interaction control, this may bypass surface obstructions and adhere itself to the target in a unique fly-crawl-fly inspection mode. Laboratory testing demonstrates enhanced stability, showing mean absolute static pose deviation below 1.03 mm and well-regulated helical translation around the full pipe circumference. Deploying immediate-proximity visual inspection, line features and defects larger than $102\ \mu\text{m}$ are then resolvable, with extensions to photogrammetric reconstruction and orthogonal imaging recording 3D geometry and full location context. Quantified operational performance bounds of each system thereby support industrial adoption and further development of airborne NDE.

Contents

Declaration & Copyright.....	i
Acknowledgements	iii
Abstract	iv
Contents	v
List of Figures	xi
List of Tables.....	xvii
Abbreviations & Acronyms	xviii
CHAPTER 1 Introduction	1
1.1 Industrial Context & Motivation	1
1.1.1 UK Electricity Supply	1
1.1.2 Wind Turbines.....	2
1.1.3 Nuclear Power Plants	6
1.1.4 Oil & Gas	10
1.1.5 Remote Airborne Inspection	14
1.2 Problem Statement	16
1.3 Research Objectives	18
1.4 Contributions to Knowledge	19
1.4.1 System I: An Over-Actuated UAV for Dry-coupled Airborne Ultrasonic Thickness Measurement.....	19
1.4.2 System II: A Surface Crawling UAV for Visual Inspection.....	20
1.5 Thesis Outline.....	21
1.6 Dissemination of Findings.....	22
1.6.1 Arising from this Thesis.....	22
1.6.2 Also by the Author	24

1.6.3	Further Output.....	25
CHAPTER 2	A Review of Multirotor Unmanned Aerial Vehicles and Non-Destructive Evaluation	26
2.1	Introduction	26
2.1.1	Common NDE Methods & UAV Compatibility.....	26
2.1.2	Chapter Structure	29
2.2	Remote Visual Inspection & Photogrammetry	29
2.2.1	Remote Visual Inspection	30
2.2.2	3D Reconstruction & Photogrammetry.....	32
2.2.3	Reconstruction Sensors & Software	33
2.2.4	Orthoimage Mapping	34
2.2.5	Evaluation	36
2.3	Ultrasonic NDE	36
2.3.1	Wave Modes & Propagation Speed	36
2.3.2	Reflection, Transmission & Refraction.....	40
2.3.3	Piezoelectric Transducers.....	43
2.3.4	Thickness Measurement & Flaw Detection	48
2.3.5	Evaluation	52
2.4	Multirotor UAVs	52
2.4.1	Propeller Thrust Generation.....	53
2.4.2	Collective Actuation.....	57
2.4.3	Free-Flight Model	61
2.4.4	Flight Control Structures.....	66
2.4.5	Evaluation	69
2.5	Airborne Non-Destructive Evaluation.....	71
2.5.1	Non-Contact Systems.....	71

2.5.2	Self-Supported Contact Systems.....	80
2.5.3	Externally Supported Contact Systems.....	86
2.5.4	Other Technology Trends	91
2.5.5	Evaluation	95
2.6	Summary	99
CHAPTER 3 Dry-Coupled Airborne Ultrasonic Non-Destructive Evaluation....		101
3.1	Introduction	101
3.1.1	Dry-Coupling Ultrasonic Probes.....	103
3.1.2	System Conceptualisation	104
3.1.3	Chapter Structure	106
3.2	System Overview	106
3.2.1	Voliro Tri-Copter	106
3.2.2	Contact NDE Payload	110
3.2.3	Acoustic Coupling Behaviour	114
3.3	Experimental Assessment Methodology	123
3.3.1	Flight Testing Facility	123
3.3.2	Inspection Samples.....	124
3.3.3	Ultrasonic Thickness Measurement	128
3.3.4	Sample-Relative Measurement Localisation.....	133
3.4	Results	136
3.4.1	Vertical Wall Thickness Measurement	136
3.4.2	Point Inspection Repeatability	139
3.4.3	Accessing the Underside of Overhanging Structures.....	142
3.4.4	Dynamic Scanning Measurement Acquisition.....	144
3.5	Discussion	151
3.6	Summary	154

CHAPTER 4 A Hybrid Unmanned Aerial Vehicle – Crawler for Airborne Pipe
Inspection 155

4.1	Introduction	155
4.1.1	Pipe Exterior Inspection Robots.....	156
4.1.2	System Conceptualisation	158
4.1.3	Chapter Structure	160
4.2	System Modelling & Prototype Vehicle	160
4.2.1	Prototype Vehicle.....	161
4.2.2	Rigid-Body Interaction Model	163
4.2.3	Vehicle Design Feasibility	167
4.3	Optimal Thrust Support Setpoint	173
4.3.1	Wrench Energy Optimization	173
4.3.2	Optimal Supporting Leg Torque	175
4.4	System Control.....	180
4.4.1	Flight Control.....	182
4.4.2	Interaction Control	186
4.4.3	Wrench Allocation	189
4.4.4	Surface Crawling.....	189
4.5	Experimental Assessment Methodology	193
4.5.1	Motion Tracking.....	193
4.5.2	Imaging Performance	194
4.5.3	Geometry & Surface Texture Inspection	196
4.6	Results	203
4.6.1	Interaction Stability	203
4.6.2	Imaging Performance	205
4.6.3	Landing Repeatability	211

4.6.4	Fly-Crawl-Fly Inspection	213
4.6.5	Surface Area Inspection	215
4.7	Discussion	225
4.8	Summary	229
CHAPTER 5	Conclusion.....	231
5.1	Summary	231
5.2	Key Findings & Implications	237
5.3	Limitations & Future Work.....	242
5.3.1	Sensor-Based Interaction Control	242
5.3.2	Environmental Influence & Field Trials	243
5.3.3	Applicability to Generic Asset Geometries.....	244
5.3.4	Additional NDE Sensory Modes.....	245
5.3.5	Enhanced Automation.....	246
5.3.6	BVLOS Operation & Site Residency.....	247
5.3.7	NDE Data Fusion	248
APPENDIX A	Photogrammetric Theory	250
A.1	The Pinhole Camera Model.....	250
A.1.1	Simple Linear Representation	250
A.1.2	Principal Point Offset & Pixel Space	253
A.1.3	Arbitrary Camera Pose	253
A.1.4	Revised Camera Model	254
A.1.5	Lens Distortion.....	255
A.1.6	Multiple Viewpoints & Epipolar Geometry.....	257
A.2	3D Reconstruction	259
A.2.1	Image Acquisition	262
A.2.2	Feature Detection	262

A.2.3	Tie-Point Matching	263
A.2.4	Camera Alignment	264
A.2.5	Point Cloud Densification	266
A.2.6	Mesh Generation	268
A.2.7	Mesh Texturing	269
APPENDIX B	Extended Aspects of Ultrasonic NDE	272
B.1	Beam Spread.....	272
B.2	V-Path Thickness Measurement.....	274
B.3	Cross-Sectional & Volumetric Imaging	281
B.4	Full Matrix Capture Methods	283
B.5	Alternate Methods for Ultrasound Transduction.....	284
APPENDIX C	Comparative Review of UAV Control Theory	286
C.1	Common UAV Control Algorithms	286
C.1.1	Proportional-Integral-Derivative Control.....	286
C.1.2	Linear Quadratic Regulators	288
C.1.3	Sliding Mode Control.....	290
C.1.4	Backstepping Control.....	291
C.1.5	Model Predictive Control.....	293
C.2	Comparative Evaluation	294
References	297

List of Figures

<i>Figure 1-1:</i> Examples of wind turbine blade damage.	3
<i>Figure 1-2:</i> Age and geographic distribution of operational UK wind turbines.	5
<i>Figure 1-3:</i> Examples of structural monitoring in nuclear applications.....	6
<i>Figure 1-4:</i> Example of reconnaissance to determine the state of an enclosed structure.	8
<i>Figure 1-5:</i> Examples of flow accelerated corrosion.	9
<i>Figure 1-6:</i> Examples of corroded and degraded pipework within the oil and gas sector.	11
<i>Figure 1-7:</i> Examples of steel containment structures within the oil and gas sector.	12
<i>Figure 1-8:</i> Example of work at height conducted from an offshore platform.	13
<i>Figure 2-1:</i> A diagram of the pinhole camera model.	31
<i>Figure 2-2:</i> An example of an aerially captured orthoimage of a light industrial facility [131].	35
<i>Figure 2-3:</i> Common wave modes used in ultrasonic NDE shown by their distortion of a 2D grid.	37
<i>Figure 2-4:</i> Ultrasonic wave reflection, refraction, and mode conversion at the boundary between a fluid and a solid material.....	42
<i>Figure 2-5:</i> Cut-away view of a single element ultrasonic probe with an added delay line.....	45
<i>Figure 2-6:</i> Single-element ultrasound probe configurations for inspection.	48
<i>Figure 2-7:</i> A bounce diagram showing ultrasonic wave propagation through a sample containing a defect.	50
<i>Figure 2-8:</i> A diagram of the primary forces and torques exerted on the UAV body by a single propeller in static conditions, arising from the aerodynamic effects of its rotation.	54
<i>Figure 2-9:</i> Secondary propeller aerodynamic effects arising from rotor flapping.	55
<i>Figure 2-10:</i> Quadcopter actuation and free-body diagram.	57

Figure 2-11: Axes of motion for a conventional quadrotor UAV shown relative to a static hover pose in grey.....	61
Figure 2-12: Generic multirotor nested flight control architecture.....	67
Figure 2-13: A screenshot from the Cyberhawk iHawk inspection reporting and analysis software [227].	72
Figure 2-14: Results of automated defect detection in UAV inspection images.	73
Figure 2-15: UAV-based photogrammetric inspection of a timber frame bridge by Khaloo <i>et al</i> [232].	74
Figure 2-16: Investigation into factors affecting UAV photogrammetric reconstruction quality of a wind turbine blade by Zhang <i>et al</i> [234]......	76
Figure 2-17: A pilot assistance algorithm for photogrammetric UAV inspection by Omari <i>et al</i> [235]......	77
Figure 2-18: Automated photogrammetric inspection of an industrial chimney by Quenzel <i>et al</i> [236]......	78
Figure 2-19: Fusion of radiation dosimetry survey data with photogrammetric mesh of a contaminated topsoil store near the site of the Fukushima power plant by Connor <i>et al</i> [238]......	79
Figure 2-20: Commercial UAV ultrasonic inspection platforms.....	81
Figure 2-21: Under-actuated multirotors for contact inspection.....	82
Figure 2-22: A comparison of multirotor configurations for environmental interaction.....	84
Figure 2-23: Aerial Manipulator UAV platforms.....	85
Figure 2-24: Crawler robots used for commercial NDE.....	87
Figure 2-25: Pipe exterior crawlers from recent literature.....	88
Figure 2-26: Wall-crawling UAVs using active thrust-based support.....	89
Figure 2-27: UAVs for passive perching and inspection on pipe-like structures. ...	90
Figure 2-28: Evaluation platforms for full-circumferential pipe access by a UAV. 91	
Figure 2-29: Collision tolerant UAV platforms for confined space inspection.....	92
Figure 2-30: Examples of commercial site-resident UAV systems.....	95
Figure 3-1: Couplant material fills microstructure at the probe-object interface, minimizing impedance changes.	103
Figure 3-2: The Voliro tricopter aerial manipulation platform.....	107

Figure 3-3: A block diagram of the Voliro control structure used to maintain omnidirectional stability and perform environmental interaction.....	108
Figure 3-4: A dual element wheel probe cross-sectional view and diagram of ultrasound wave propagation during measurement.....	111
Figure 3-5: Functional block diagram of the embedded dual channel ultrasonic transceiver.	113
Figure 3-6: Dry-coupling wheel probe coupling criteria include both compression and orientation.....	114
Figure 3-7: Experimental setup for identification of wheel probe dry-coupling requirements.....	115
Figure 3-8: Wheel probe force-coupling behaviour assessment.....	116
Figure 3-9: Ultrasound signal response to changing coupling forces.....	117
Figure 3-10: Effects of orientation on wheel probe acoustic coupling.....	120
Figure 3-11: Photographs and as manufactured dimensions of the aluminium plate sample, (a) & (b), and stepped bar sample, (c) & (d).....	125
Figure 3-12: Diagram showing extraction of reference object thickness at measurement locations using line-plane intersection.....	127
Figure 3-13: A synthetic ultrasonic A-scan signal waveform showing the inter-echo time-of-flight, T_k	130
Figure 3-14: The auto-correlation of the synthetic A-scan signals depicted previously.	131
Figure 3-15: The trend in measured thickness using the inter-echo time-of-flight, T_k	132
Figure 3-16: Projection of ultrasonic measurement locations onto the sample face.	135
Figure 3-17: The UAV positions and applies the probe to the plate sample face..	137
Figure 3-18: A map of the points on the sample outer face inspected in order of top to bottom during the test flight.....	138
Figure 3-19: Error in the remote measurement versus reference geometry plotted across the time of the flight trial.	138

Figure 3-20: Map of inspection points acquired through repeated interaction with upper right quadrant of target plate sample showing position uncertainty boundary.	140
Figure 3-21: The Voliro manipulator platform is able to enter and maintain stable contact with the underside of an overhanging surface.	142
Figure 3-22: The UAV flightpath when inspecting the underside of an overhanging structure.	143
Figure 3-23: Error plot of thickness measurements from the underside of an overhanging surface versus reference geometry.	143
Figure 3-24: A top-down view of the UAV scanning across the stepped aluminium bar.	145
Figure 3-25: Ultrasonic data captured by rolling the wheel-probe along the bar x-axis and back again.	146
Figure 3-26: Aerial thickness measurement scans of the bar sample along its x-axis.	147
Figure 3-27: The UAV flightpath while conducting a rolling scan measurement.	149
Figure 3-28: Freebody diagram of the wheel probe rolling scan interaction.	150
Figure 4-1: A diagram of the prototype hybrid vehicle.	161
Figure 4-2: Free-body diagram showing forces and torques acting on the hybrid UAV-crawler vehicle.	164
Figure 4-3: Maximum propeller thrust of the candidate vehicle in B relative to minimum required supporting force.	171
Figure 4-4: Design feasibility margin plot for the candidate vehicle design.	173
Figure 4-5: Energy optimal interaction force-torque setpoint and contact forces computed around the target surface.	175
Figure 4-6: Comparison of energy optimal leg torque setpoint (solid line) and the piecewise function setpoint (dashed line).	180
Figure 4-7: UAV-crawler hybrid controller block diagram.	181
Figure 4-8: A plot of error between the optimal interaction propeller force-torque setpoint and the approximation generated using Algorithm 1.	188
Figure 4-9: Diagram of differential steering crawler kinematic used by the hybrid vehicle.	190

Figure 4-10: Image resolution artefact used under USAF MIL-STD-150A spatial resolution assessment procedure.	195
Figure 4-11: Fitting of the reconstructed pipe mesh to a cylinder-aligned coordinate frame, C	198
Figure 4-12: Unrolling of the reconstructed pipe mesh texture. (a) Spatial attributes of the pipe in the cylinder-aligned frame C	201
Figure 4-13: Sequential image series showing the vehicle holding position at each 45° station about the pipe.....	205
Figure 4-14: Unprocessed full-frame images of the USAF artefact captured by the vehicle during static pose holding at 45° intervals around the pipe circumference.	207
Figure 4-15: Images of the central region of the USAF artefact cropped and enlarged from those of Figure 4-14 without re-sampling or any further processing steps to enhance resolution.....	208
Figure 4-16: Resolution of images captured while in static position at 45° increments around the pipe.....	209
Figure 4-17: Landed vehicle positions versus target location in $\{P\}$, which forms origin of this plot.....	212
Figure 4-18: 3D position and attitude trajectory of the vehicle in W as it conducts a flight-access circumferential scan.	214
Figure 4-19: Synchronised photographs of the fly-crawl-fly inspection.	215
Figure 4-20: Speed of the vehicle while scanning the pipe circumference.	215
Figure 4-21: 3D position and pose trajectory of the vehicle in P as it performs a helical translation along the pipe surface.	217
Figure 4-22: Vehicle speed while helically translating over the pipe surface.	218
Figure 4-23: Images of the reconstructed exterior pipe surface area.	222
Figure A-1: The simplest central projection pinhole camera model, linearly mapping a 3D point to a position on the image using similar triangles.....	251
Figure A-2: A more practical pinhole camera model.....	253
Figure A-3: Image distortion components.	255
Figure A-4: Epipolar geometry for multi-view point correspondence.....	257
Figure A-5: High-level comparison of the types of machine vision used by UAVs.	260

Figure A-6: A generic photogrammetric reconstruction workflow.	261
Figure A-7: A dense point cloud of a countryside house produced using commercial reconstruction software, Pix4D [127].	267
Figure A-8: Illustration of reconstruction texture mapping.	270
Figure B-1: Ultrasonic beam spread from a single unfocussed circular element piezoelectric transducer [140].	273
Figure B-2: Cut-away diagram of Single and Dual-Element probes showing the path of an ultrasonic wave-packet during thickness measurement.	275
Figure B-3: A simple model for V-path thickness measurement.	276
Figure B-4: Thickness measurement error due to V-path effects with varying element separation for a given number of available echoes.	280
Figure B-5: Ultrasonically imaging a structure cross-section.	281
Figure B-6: Example NDE results from ultrasonic array inspection of a steel pipe section with 9 mm wall thickness and 508 mm diameter.	282

List of Tables

<i>Table 2-1:</i> Comparison of Common UAV Control Methods	68
<i>Table 3-1:</i> Comparison of Airborne Ultrasonic NDE Systems	104
<i>Table 3-2:</i> Summary of Wheel Probe Acoustic Coupling Criteria.....	123
<i>Table 3-3:</i> Plate Sample Thickness as Manufactured.....	125
<i>Table 3-4:</i> Bar Sample Thickness as Manufactured.....	126
<i>Table 3-5:</i> Sample Ultrasonic Wave Velocity Calibration Data	129
<i>Table 3-6:</i> Wall Thickness Measurement Error Statistics	139
<i>Table 3-7:</i> Measurement Repeatability Data	141
<i>Table 3-8:</i> Tilted Measurement Accuracy Statistics.....	144
<i>Table 3-9:</i> Scanning Measurement Accuracy Statistics	148
<i>Table 4-1:</i> Comparative Overview of Pipe-Exterior Inspection Robots	159
<i>Table 4-2:</i> UAV-Crawler Hybrid Vehicle Rotor Layout.....	162
<i>Table 4-3:</i> Stability of the Vehicle around the Pipe	204
<i>Table 4-4:</i> Image Resolution when Stationary around the Pipe	209
<i>Table 4-5:</i> Landing Location Deviation from Target	212
<i>Table 4-6:</i> NDE Features Visible in Figure 4-23(c).....	224

Abbreviations & Acronyms

2D	Two Dimensional
3D	Three Dimensional
ADC	Analogue to Digital Converter
ALARP	As Low As Reasonably Practicable
API	American Petroleum Institute
ARE	Algebraic Riccati Equation
ASL	Automated Systems Lab
ASME	American Society of Mechanical Engineers
ASTM	American Society for Testing and Materials
ATLAS	Association of Technical Lightning & Access Specialists
BEIS	Department for Business, Energy, and Industrial Strategy
BIM	Building Information Modelling
BPVC	Boiler and Pressure Vessel Code
BRIEF	Binary Robust Independent Elementary Features
BRISK	Binary Robust Invariant Scalable Key-points
BSC	Back-Stepping Control
BVLOS	Beyond Visual Line of Sight
CAA	Civil Aviation Authority
CCD	Charge Coupled Device
CMOS	Complementary Metal Oxide Semiconductor
CoM	Centre of Mass
CT	Computed Tomography
CTV	Crew Transfer Vessel
CUE	Centre for Ultrasonic Engineering
CUI	Corrosion Under Insulation
DoF	Degrees of Freedom
DSEAR	Dangerous Substances and Explosive Atmospheres Regulations
EKF	Extended Kalman Filter
EMAT	Electro Magnetic Acoustic Transducer

EPSRC	Engineering and Physical Sciences Research Council
ESC	Electronic Speed Controller
FAST	Features from Accelerated Segment Test
FMC	Full Matrix Capture
FPSO	Floating Production Storage and Offloading
FPV	First Person View
FRD	Front-Right-Down
FSSR	Floating Scale Surface Reconstruction
GDMR	Global Dense Multiscale Reconstruction
GNSS	Global Navigation Satellite System
GPR	Ground Penetrating Radar
HJB	Hamilton-Jacobi-Bellman
HSE	Health and Safety Executive
IAEA	International Atomic Energy Agency
iCASE	Industrial Cooperative Awards in Science and Technology
ICR	Instantaneous Centre of Rotation
IMU	Inertial Measurement Unit
IRATA	Industrial Rope Access Trade Association
ISO	International Organisation for Standardisation
LCoE	Levelized Cost of Energy
LIPA	Laser Induced Phased Array
Li-Po	Lithium Polymer
LQR	Linear Quadratic Regulator
MAE	Mean Absolute Error
MDT	Multidirectional Thrust
ME	Mean Error
MLE	Maximum Likelihood Estimation
MPC	Model Predictive Control
MVS	Multi View Stereoscopia
NDA	Nuclear Decommissioning Authority
NMPC	Non-Linear Model Predictive Control
OEUK	Offshore Energies UK

O&M	Operations and Maintenance
ONR	Office for Nuclear Regulation
ORB	Oriented FAST and Rotated BRIEF
OREC	Offshore Renewable Energy Catapult
PC	Personal Computer
PCI	Phase Coherence Imaging
PEC	Pulsed Eddy Current
PID	Proportional Integral Derivative
PIG	Pipeline Inspection Gauge
PSR	Poisson Surface Reconstruction
PVDF	Polyvinylidene Fluoride
PZT	Lead Zirconate Titanate
RANSAC	Random Sample Consensus
RCNDE	Research Centre in Non-Destructive Evaluation
RGB	Red Green Blue
RMS	Root Mean Square
ROS	Robotic Operating System
ROV	Remotely Operated Vehicle
RVI	Remote Visual Inspection
Rx	Receiver
SDF	Signed Distance Function
SDRE	State-Dependent Riccati Equation
SEARCH	Sensor Enabled Automation Robotics and Control Hub
SfM	Structure from Motion
SIFT	Scale Invariant Feature Transform
SLAM	Simultaneous Localisation and Mapping
SLERP	Spherical Linear Interpolation
SMC	Sliding Mode Control
SMVS	Shading Aware Multi-View Stereoscopia
SNR	Signal to Noise Ratio
SOV	Service Operation Vessel
SPE	Society of Petroleum Engineers

STEM	Science Technology Engineering & Mathematics
SURF	Speeded Up Robust Features
SVD	Singular Value Decomposition
TDA	Temporary Danger Area
TFM	Total Focussing Method
ToF	Time of Flight
Tx	Transmitter
UAS	Unmanned Aerial System
UAV	Unmanned Aerial Vehicle
UDT	Uni-Directional Thrust
USAF	United States Air Force
UTM	Unmanned Traffic Management
VGA	Variable Gain Amplifier
VO	Visual Odometry
VRIP	Volumetric Range Image Processing
VT	Vectored Thrust

CHAPTER 1

Introduction

1.1 Industrial Context & Motivation

1.1.1 UK Electricity Supply

The UK energy generation supply chain continues to evolve towards the goals of net zero by 2050 [1], entailing a shift in focus from conventional oil, gas and coal to emphasise renewable and low emissions strategies [2]. While much remains to be done, recent milestones are indicative of significant progress. Whereas 2009 saw renewable sources supply 25.97 % of UK electricity consumption, 2019 saw this figure double to 54.22 % [3]. More recently, in the year leading up to September 2022, a further 3.4 GW of renewable capacity was added UK-wide, of which 2.8 GW was new offshore wind [4]. Continuing forwards, Levelised Cost of Energy (LCoE) estimates, describing the cost per unit energy generated over the full lifetime of the plant, suggest wind turbines commissioned in 2025 will generate energy for 67 % to 54 % the cost of established gas turbine methods [5]. Accordingly, the UK intends to progress as a global leader in offshore wind [6], [7], expanding its proportional energy contribution amid the national portfolio. Moving beyond 2022, when 38.5 % of total UK generated energy came from gas, wind provided 26.8 %, and nuclear supplied 15.5 % [8], the ongoing strategy thereby aims to secure a robust energy supply from diverse but complementary sources while minimising costs and associated emissions [9].

To achieve this, the UK must then operate and maintain an extensive fleet of energy generation assets and infrastructure, all subject to wear and degradation over time through a variety of mechanisms. Regular in situ assessment is therefore vital to ensure continued function and satisfy legal requirements to demonstrate continued safe operation within relevant industrial standards [10], [11]. Such inspection can also identify factors decreasing operational efficiency, inform maintenance activity to

lower production costs, and prevent unanticipated failures, reducing forced downtime and lost production [12], [13]. However, detailed assessment of overall plant health is a highly time consuming, laborious, and repetitive process, often placing highly skilled technicians in hazardous environments.

Further logistical challenges are posed to deployment of effective NDE across the national energy fleet by skilled labour shortages and an ageing inspector workforce. This follows a larger trend shown across UK Science, Technology, Engineering and Mathematics (STEM) industries, with reports in 2021 suggesting that a shortfall of 173 000 workers, an average 10 unfilled roles per UK business, was costing the economy £1.5 billion per annum [14], [15]. At present, from January to March 2023, some 111 000 vacancies persist in the field of professional scientific and technical activities [16], indicating strong unfilled demand for STEM personnel. For NDE, the problem is exacerbated by the workforce age demographic. A review carried out in 2007 showed that 60 % of institutionally registered UK inspectors were over 50 years of age, making them the oldest cohort in Europe [17]. As such, a high rate of inspector retirement may be expected, depleting the collective skill level and experience of the remaining inspectors. Estimates from 2014 suggested approximately 900 new starts per year would be required to maintain population but, with only 500 per year being trained, a 10 % to 15 % NDE labour shortage could be anticipated going forwards [17].

Alongside increased training, systemic improvements to existing NDE practices are then necessary augment the capability of the existing workforce and meet this shortfall, while simultaneously addressing the specific challenges specific posed in maintaining the diverse range of national energy assets, as discussed below.

1.1.2 Wind Turbines

As a key component of the UK energy strategy, wind turbines represent a pertinent example. By functional necessity these assets are installed in remote locations with persistent winds for optimal power generation, but as such they are continuously exposed to a harsh environment. Beyond wear in the generator and ancillary electrical systems which may be addressed through classical machine condition monitoring

techniques [18], the most severe degradation is a direct result of this operating environment [19], [20]. The composite blade structures are predominantly affected, degrading performance or causing critical failure through various damage types, as depicted in Figure 1-1.

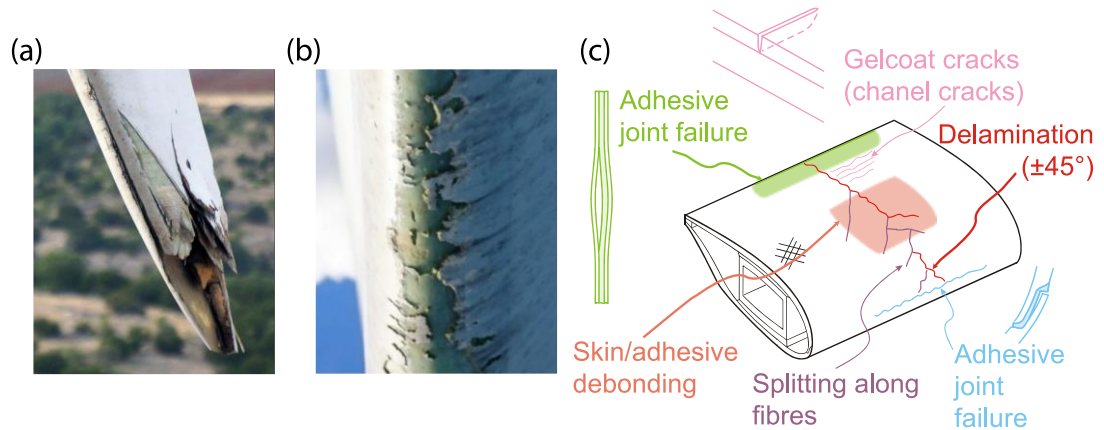


Figure 1-1: Examples of wind turbine blade damage. (a) Blade tip detachment following lightning strike [19]. (b) Leading edge erosion [21]. (c) Composite degradation mechanisms under loading [22].

Lightning strikes cause the most severe damage, to which wind turbines are especially vulnerable as tall isolated structures [23]. When struck, rapid heating and explosive vaporisation of any moisture in the blade structure can cause delamination, shell debonding and tip detachment as in Figure 1-1(a), each representing a critical failure mode. Lightning protection systems can substantially limit damage severity but remain imperfect, leading to more subtle damage within the structure. Similarly, impingement of rain droplets and particulate matter such as hail or sea spray on the leading edge of the rotating turbine blade causes pitting and erosion of the exterior coating and fibreglass composite [24], [25], as in Figure 1-1(b). This can increase aerodynamic drag by up to a factor of five and significantly reduce electrical output [19]. Without detection and remediation erosion can also allow water ingress to the composite core material, significantly weakening its structure. Repair of such defects is an invasive process, requiring the affected material to be ground out and replaced before the exterior surface can be resealed [26]. Fatigue associated with fluctuating wind gusts, transient behaviours outside of the designed load, and cyclical rotation effects can also lead to significant degradation [19], [27]. As in Figure 1-1(c), these mechanisms can introduce various structural defects, gradually expanding subsurface manufacturing

flaws, or other damage incurred during installation and operation, to eventually cause critical failure.

Access challenges make detection and monitoring of these structural defects non-trivial, with many critical damage features located within the material bulk and presenting no remotely visible indicators at the exterior surface. Wind turbine inspection and maintenance activity can then necessitate direct evaluation via rope access. Generally, such work is regarded as a hazardous procedure to be avoided whenever possible, with falls from height remaining the single biggest cause of UK workplace deaths during the year from March 2021 to 2022, with 29 reported instances nationwide [28]. However, there is often no alternative but for a highly skilled technician to perform the work necessary to retain asset function. This can see personnel suspended over 50 m above ground level and exposed to environmental risk factors including high or changeable wind speed and hypothermia. When forecasts or current circumstances exceed defined risk mitigation thresholds, work can often be prevented from starting or interrupted, with technicians kept on the ground until conditions improve [29]. The process is also time consuming and costly, requiring the turbine to be restarted and stopped multiple times to reposition each blade and enable abseiling down its length. Further, alternate crane access strategies can incur rental fees upwards of USD \$350 000 per week, as stated in a cost analysis of wind turbine failures from 2020 [30].

Challenges for offshore turbines are greater still, requiring Crew Transfer Vessels (CTVs) to transport technicians to the site with hire rates ranging from £2000 to £4000 per day in 2019 [31]. On top this, the charter period must include significant contingency to ensure work can be completed amid wind and wave height limits for safe crew transfer to the turbine [32]. Larger Service Operation Vessels (SOVs) may operate in rougher seas and provide on-site accommodation to reduce transit times but incur far higher costs, estimated to reach up to £60 000 per day in mid 2019 [33]. Wind turbine inspection thus presents a number of hazards to human technicians alongside substantial associated operation and maintenance costs, reportedly in the range of £70 000 to £80 000 annually per megawatt of offshore capacity as of 2019 [31].

While future technological improvements may reduce costs for a given offshore O&M

process, overall inspection spending is expected to rise markedly in the next 5-10 years. As comparatively new installations, wind turbines to date have benefited from minimal inspection requirements, but their UK population is now beginning to age past their 20 year design lifespan [34], [35], as shown in Figure 1-2(a). To maintain operability and cost-effective energy generation capacity, more frequent and in-depth inspection will become necessary, providing the actionable insight into their overall structural health necessary to target repairs and sentence assets for decommissioning or lifetime extension.

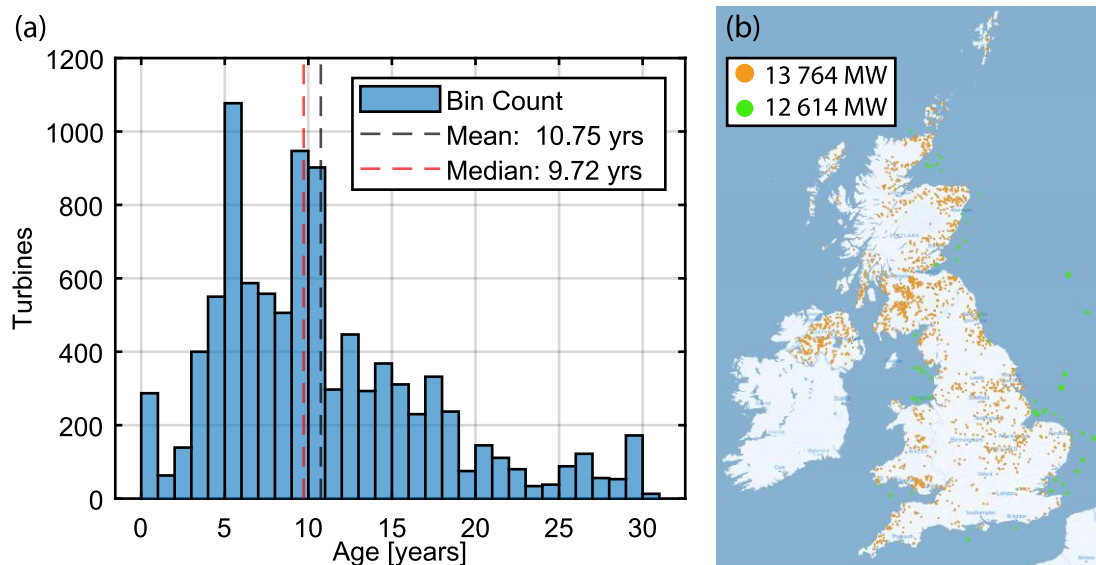


Figure 1-2: Age and geographic distribution of operational UK wind turbines. (a) Age histogram calculated based on when the wind farm became operational prior to December 2022. (b) Map of onshore (orange) and offshore (green) wind farm locations and total operational capacity. Data source: Renewable Energy Planning Database [36].

Moreover, a shifting proportion of offshore assets may further raise overall costs. Whereas current UK wind generation capacity is split approximately evenly between onshore and offshore sites, as shown in Figure 1-2(b) [36], installation of new assets to meet targets for 50 GW of offshore wind generation capacity by 2030 [9] will place a greater proportion far out at sea, entailing significant increase in crew transfer expenditure. Ageing of existing infrastructure alongside eventual degradation of newly installed wind turbines will therefore ensure the need for end of life and recommissioning inspection persists well into the future.

1.1.3 Nuclear Power Plants

A similar trend for lifetime extension is presented by the UK nuclear energy sector, supporting continuous low-carbon generation capacity while the country constructs new renewable assets and implements plans to extend the currently operable 5 883 MW of nuclear capacity to 24 GW before 2050 [9], [37]. As such, of the 8 UK nuclear plants operable in 2016, a total of 7 have seen operational extensions ranging between 3 and 10 years past their original design life, without which only one plant would have remained operational past 2023 [37].

Given the criticality of safe operations, this has entailed stringent inspection to ascertain current nuclear plant health status and to support repair and upgrade activity prolonging useful functionality [38], [39]. Inspection also forms an integral part of the decommissioning process which is conducted after a nuclear plant has reached its maximum operable lifespan to return the site to an agreed end state, fit for other purposes in wider society [40]–[42]. However, these assessments can present certain challenges both in live plant and post-operational settings, as highlighted below.




Figure 1-3: Examples of structural monitoring in nuclear applications. (a) The exterior of the reactor containment buildings, such as the DC Cook power plant domes [43]. (b) Example of concrete cracking and spalling found there [43].

A prominent example concerns monitoring for structural degradation, with a particular focus on the reactor containment building, as depicted in Figure 1-3(a). Under the standards set by ASME Boiler and Pressure Vessel Code (BPVC) Section XI [44], detailed visual inspection of active facilities must be completed at least every 10 years, meeting requirements to detect cracks, as in Figure 1-3(b), pop-outs, spalling, erosion,

corrosion or other features as small as 0.8 mm to ensure physical integrity. Facility monitoring requirements also extend into the safe enclosure phase of the nuclear decommissioning process, wherein the structure is safely held in state while sealed off from the environment. This occurs following removal of any major radiological contaminants and can span from months to decades to allow natural decay of any remaining short half-life radionuclides and so reduce waste products [40]–[42]. Effective care and monitoring then entails frequent but less intensive visual screening for signs of structural degradation, missing fixtures and damage to ventilation systems, deterioration of roofing, and evidence of animal life that may inadvertently spread contaminated material, typically performed every 4 – 6 weeks [45]. In both cases, attaining the necessary viewpoints for survey over the entire building exterior surface can entail extensive scaffolding, crane lifts, or rope access, incurring hazardous working at height. A manual inspection is then a highly time and cost intensive process, exacerbated by extensive facility scale and the need for repetition throughout the plant lifespan.

Further challenges are presented in decommissioning where insufficient knowledge of the layout, contents, or structural status of an encapsulated facility is available to effectively plan subsequent decontamination and remediation steps. This can arise due changes in decommissioning practices, infrastructure degradation during a protracted enclosure period, or when circumstances force sudden unanticipated facility closure.

A prominent example is given in the depleted fuel holding pool depicted in Figure 1-4 which required investigation following the Chernobyl accident in 1986 to confirm that it had never been used and was free of hazardous material [46]. In such cases, the requisite information often cannot be obtained from outside the sealed structure and direct survey is necessary. However, excessive residual contamination can make manual entry to ascertain facility status impossible without incurring personal health risk. The nuclear workforce therefore abides by legal limits establishing a maximum annual dose rate [47], under which safe systems of work can be defined following As Low As Reasonably Practicable (ALARP) guidance [48].



Figure 1-4: Example of reconnaissance to determine the state of an enclosed structure. Here exploration and monitoring of an unused holding pool building is conducted within the Chernobyl nuclear power plant [46].

Where maximum working duration set within these limits then restricts feasibility of in-person tasks, specialised Remotely Operated Vehicles (ROVs) may be used for exploration, inspection, and other basic manipulation [49]. However, as ground-based vehicles, their access can be denied by floor-level obstructions while their low viewpoint limits visibility of countertops or elevated structures within high ceiling rooms. Complete satisfaction of inspection requirements can thus require new technologies with greater mobility.

Radiation exposure also presents an obstacle to inspection processes within active nuclear power plants, where degradation arising as part of their normal function must be regularly monitored [44]. A notable example shared across the energy sector is the corrosion and erosion of the heat exchange pipework necessary to drive a steam turbine generator. This can occur via a number of mechanisms including: cavitation, flashing erosion, liquid impingement erosion, solid particle erosion, and flow accelerated corrosion, as in Figure 1-5 [50], [51].




Figure 1-5: Examples of flow accelerated corrosion. (a) Typical degradation pattern. Note the pipe wall thinning in the upper right corner. [52] (b) Failed section of the feedwater pipe from the Surry nuclear plant, USA [53].

Of these, flow accelerated corrosion, is regarded as one of the most destructive mechanisms affecting high-temperature carbon steel pipes due to its combined corrosive dissolution and mechanical erosion of the pipe wall [51]. In the 1980s, this led to a number of incidents such as at the nuclear power plant in Surry, USA [54]. There, flow accelerated corrosion caused sudden failure of the 18 inch (460 mm) diameter pipe to the main feedwater pump, resulting in ductile rupture and ejection of a 0.6 m by 1.2 m section of pipe wall, as depicted in Figure 1-5(b), and tripping the reactor into automatic shutdown. Consequently, modern standards emphasise regular pipework inspection [44], with particular focus on regions likely to experience flow accelerated corrosion under reliability-based monitoring [55]. Evaluation then draws upon multiple sensing modes, including radiographic, ultrasonic, eddy current, and visual testing, each supported by detailed standards of practice formed over decades of industrial expertise and codified in ASME BPVC Section V [56]. However, outside of specialised applications, the range of dexterous interactions required to capture measurement data across all points of inspection often requires a human technician.

Such manual inspection is a major undertaking, necessitating shutdown and depressurisation to ensure a safe working environment, as well as extensive insulation removal and subsequent reinstallation for access to the entirety of critical components [44]. At an electricity spot rate of £50/MWh, as was often the case before 2020 [57], such outage in an average UK nuclear power plant operating at net capacity of 1200 MW [37] would cost an estimated £1 440 000 per day in lost income alone. To minimise downtime, inspection is then conducted in tandem with other planned outages, flagging any non-critical flaws detected for further monitoring or repair during the next inspection and maintenance campaign [55]. This maximises cost

effectiveness but can incur significant delays and expense where unanticipated degradation is found and requires immediate repair. An ability to perform more frequent in-depth examination without incurring the costs of human presence is hence recognised as highly desirable within predictive maintenance strategies.

1.1.4 Oil & Gas

While the UK transitions to a renewable low-carbon energy model, the oil and gas sector continues to make significant contributions to current generation capacity and the national economy, providing an estimated 26 900 directly related jobs and a £31.1 billion gross value added nationwide during 2021 [58]. It also shares in many of the challenges encountered in other sectors such as managing and remediating the effects of corrosion, with total annual costs across all global industry estimated at USD \$1.4 trillion in 2016 [59].

This deterioration can be particularly severe in many structures employed by the oil and gas sector, including pipework, pressure vessels, above ground storage tanks and offshore platforms. Regular inspection, performed in satisfaction of applicable industrial standards, is then vital to ensure their continued structural integrity and safe functionality. Failure to do so can result in extensive unplanned outages, with lost production alone costing from USD \$340 000 to upwards of USD \$1.7 million per day for an average refinery in the USA, as estimated in 2021 [60]. However, the assessment process can often encounter challenges due to the nature of the asset.

Pipes, such as those depicted in Figure 1-6, frequently experience both interior and exterior corrosion. Internally, they see many of the same corrosion, erosion and cracking mechanisms as in a nuclear setting, with degradation augmented by further application specific effects like stress corrosion cracking and high temperature hydrogen attack [61]. Externally, vibrations from turbulent flow or nearby pumps can loosen fixtures and cause leaks. Pipes also commonly exhibit atmospheric corrosion and water ingress through a protective paint coating or thermal insulation. This can lead to can Corrosion Under Insulation (CUI), where moisture trapped at the pipe wall-insulation boundary causes pockets of accelerated corrosion, as in Figure 1-6(b).




Figure 1-6: Examples of corroded and degraded pipework within the oil and gas sector. (a) Elevated pipework with general atmospheric corrosion [62]. (b) Insulated pipe with cladding removed to show corrosion under insulation [63].

Undetected beneath the outer jacket, the CUI defect then grows and causes extensive damage [64], implicated as a root cause in over 20 % of major component failure incidents in EU oil and gas plants since 1984 [65].

Detailed pipe inspection procedures are thus defined under API 570 standards [66], completing regular in-service monitoring and extensive volumetric structural assessment at 5-10 year intervals, depending on risk status. Elevated pipework, however, presents access challenges, worsened by its routing through confined spaces above and around other infrastructure. Manual assessment at height can then only be conducted with resource intensive scaffolding installation or use of mobile cranes. Moreover, robotic solutions deployed inside the pipe can entail process outage, may be obstructed by features like narrow-bores or complex geometries, and can require installation of dedicated access ports [67], [68]. Minimally invasive remote in-service inspection from outside the pipe can thus be more attractive, provided it can operate around supports and other exterior surface furniture.

Many of the same considerations apply to pressure vessels and above ground storage tanks, such as those depicted in Figure 1-7, where inspection is conducted under ASME BPVC Sections V and VIII [56], [69], or API 653 standards [70], respectively.




Figure 1-7: Examples of steel containment structures within the oil and gas sector. (a) Pressure vessel. [71]. (b) Above ground storage tanks [72].

For storage tanks, this entails in-depth visual examination of the full exterior surface at least every 5 years, checking for coating or insulation flaws, leaks, and structural distortion. Accompanied by off-line internal visual inspection, extensive non-destructive measurements are also applied in a regular grid across the entire tank shell, a total surface area approaching 20 000 m² in the largest cases [73]. Amid localised corrosion and pitting mechanisms, this process ensures that the tank meets a minimum acceptable wall thickness throughout and is repeated every 5-15 years, depending on material loss rate [70]. Beyond this, pressure vessel inspection then introduces more stringent volumetric inspection requirements to detect cracks, blisters, embrittlement, metallurgical changes, and other structural flaws within the shell material. Regions near welds, their heat affected zones, previous repairs, and stress concentrators receive particular focus, each with dedicated evaluation processes. Inspection intervals are also more frequent, conducted as often as annually in high risk cases [56], [69].

Monitoring of storage tanks and pressure vessels is therefore a highly resource intensive process, often necessitating draining, transport and alternate storage of their contents. Further, vaporised trace elements of their contents may linger within these confined spaces, presenting inhalation hazards or creating the highly flammable and potentially explosive atmospheres referred to as ATEX conditions [74]. UK Dangerous Substances and Explosive Atmospheres Regulations (DSEAR) [75] then mandate extensive cleanout and venting periods, intrinsically-safe non-sparking equipment, development of “hot work” procedures, and specialised training to eliminate accidental ignition and mitigate other associated hazards ahead of human entry.

These challenges are accentuated aboard an offshore oil and gas platform where geographical isolation and abundant highly flammable products mean that accidents can quickly become major incidents before external assistance arrives. Alongside rapid corrosive degradation of steel structures due to the maritime environment, inspection must be more frequent than is conducted onshore and carries increased safety risks [76]. This is particularly true for components such as risers, umbilical cables, or platform support legs within the splash zone, where wave impacts and wet-dry cycling are constant and together form the most highly degressive environment aboard the platform [77]. Whereas inspection and repair activity above this region remains possible via rope access, as in Figure 1-8, the hazard profile is significantly increased by the potential of the waves immediately below to knock the inspector into contact with the structure or entangle them in their lines. For this reason, human access to the splash zone is almost impossible and strong preference is given to remote assessment solutions for regions above or below.




Figure 1-8: Example of work at height conducted from an offshore platform. Here rope-access is used to conduct inspection and repair of the platform legs, operating above the splash zone [78].

1.1.5 Remote Airborne Inspection

Considering the many hazards presented to inspectors and the subjective nature of human assessment, automated NDE procedures are noted to be of great interest across all fields of the energy sector. Amid “Industry 4.0” discussions, automation via mobile robotics offers increased inspection speed, more traceable quantitative profiles of asset health, and is less subjective to inspector opinion or fatigue [79], [80]. Under this drive and alongside other complimentary technologies [49], [68], multirotor UAV systems have become well established as practical tools for non-contact visual screening across energy generation applications [81]. They are now widely used survey remote, inaccessible, or otherwise hazardous structures such as wind turbines [82], nuclear sites [45], and oil & gas plants [83], with wider applications including solar panels [84], power lines [85], bridges [86] and quarries [87]. As such, they represent highly mobile platforms granting a unique viewpoint unattainable via terrestrial methods alongside numerous other benefits.

In the wind sector, UAV survey can categorise the presence of damage to schedule immediate intervention, plan subsequent maintenance, or entirely preclude significant rope access and unnecessary work at height. Versus the manual process, aerial inspection has been shown to increase the daily number of wind turbines assessed by an inspector from 2-5 to upwards of 12, each blade taking less than 10 minutes to examine [88]. In a recent case study from 2019, UAV inspection was thus identified as a means to reduce estimated average inspection costs for an 83 turbine, 300 MW, offshore wind farm from £125 000 per year to £79 000 per year, while also reducing the time between inspections from 5 years in the prior rope access scheme to 3 years using UAVs [89]. Furthermore, the UAV process significantly mitigated safety concerns associated with offshore work at height and the need to transfer crew to the turbine.

In nuclear applications, UAVs allow rapid, flexible, visual assessment of extensive sites amid potential radiological hazards to minimise unnecessary human entry. Their utility is exemplified through inspection of a North American nuclear power plant, covering six structures including the concrete dome exterior of a reactor containment

building to ASME BPVC Section XI standards [43], [90], [91]. There, in 2018, the UAV satisfied visual inspection requirements with 100 % surface coverage over a total area in excess of 25 000 m² in under 5 days. It thus offered an estimated cost reduction of over USD \$375 000, saving 66 % versus the standard process conducted using a crane-suspended man basket over a period of 2 weeks and also minimising work at height [43], [90], [91]. A further example highlights the radiation exposure safety benefits of UAV inspection while investigating a suspected valve leak at an elevated position within an active reactor building in 2017 [92]. In this case, the UAV gathered high resolution video sufficient to confirm absence of any seepage in under 10 min. By contrast, human entry would have required an 80 % reduction in reactor output, ramped over a total duration of 12 hours, before performing an assessment lasting 1-2 minutes during which time the inspector could receive up to 10 % of their annual dose limit [92]. Thus, utilising a UAV saved approximately USD \$456 000 in lost production [92], removed the need for human radiation exposure, and also yielded a more comprehensive inspection.

In oil and gas applications, the benefits of UAV assessment are demonstrated across multiple further use cases. For example, where inspection of the primary tank within a Floating Production Storage and Offloading (FPSO) vessel can see a four-person team spending 14 days within a confined space, UAV operations allow two inspectors to complete the process in under 4 days without entering the tank [93]. Similarly, where inspection of the underdeck of a jack-up platform can require a six-person team of rope access technicians working at height above the splash zone for almost 100 days, UAV inspection can deliver a full visual screening in just 3 days to inform subsequent highly-targeted manual follow-up [94]. Hazard exposure is then minimised while the inspection process and lost production costs can see a reduction on order of over £1 million, as in an industrial case-study from 2014 [94].

Clear and abundant evidence of the practical advantages offered by airborne inspection is thus apparent throughout the energy generation sector. However, these methods are almost exclusively non-contact visual screening strategies. To directly assess the more critical structural defects posed by fatigue, cracking, and corrosion, inspection must draw upon volumetric NDE methods [95]. Following initial UAV screening to identify

high-risk target locations, these tests must still be applied by human inspectors. Versus a purely manual inspection, such combined processes do increase overall structural health insight while cutting hazard exposure, asset downtime, and associated expenditure, but the remaining in-person component limits these benefits. Motivation is then presented for solutions where detailed volumetric assessments can be conducted less sparingly, granting asset managers data that could directly inform the most cost-effective repair and predictive maintenance strategies.

1.2 Problem Statement

UK power generation encompasses a number of technologies to ensure robustness and availability of the electricity supply, but their continued monitoring and upkeep can present a number of practical challenges. To combat the degradation experienced as part of the normal operation of these industrial structures and the accelerating factors of their environment or function, regular inspection must be conducted throughout their life cycle. Where human entry is required to ascertain structural health, this activity poses a number of risk factors:

- The newly aging wind sector presents a need for rope-access to monitor turbine blades, entailing dexterous work at height amid inhospitable weather conditions. Furthermore, offshore wind farms require this to occur above the open sea and also entail considerable expense to transport crew to the asset before inspection has even begun.
- Nuclear plant management and decommissioning can mirror work at height requirements, vital to provide full-coverage structural health assessment of the active facility and monitor the enclosed post-operational site. Here, however, the greatest challenges stem from minimising radiation exposure risks during manual access for essential maintenance or to plan and conduct decommissioned site remediation.
- Inspection of oil and gas structures encounters similar situational hazards including work at height and in confined spaces amid possible ATEX

conditions. Prevalence of steel structures subject to substantial interior and exterior corrosion can also involve exhaustive inspection campaigns across vast facilities, with such problems accentuated in offshore platforms.

In all cases, thorough inspection across the full extent of live assets can also necessitate lengthy periods of downtime, where the asset does not generate income. Personal safety and cost of inspection must then be delicately balanced against the value of the insight attained. Beyond this, labour shortages and an ageing population of domain specific experts can present further obstacles to conducting frequent manual NDE. Strong motivation has therefore accrued across the wind, nuclear, and oil and gas sectors for development of robotic or automated inspection strategies whereby the same high-quality data can be quickly gathered remotely, at higher density, and with reduced outage periods to maximise benefits of safe, cost-effective operations and management practices.

In this context, airborne inspection by UAV agents is shown to present tangible benefits. Their highly mobile flight and ability to hover in a static location make them particularly suited to applications involving observatory work at height, allowing rapid deployment to elevated targets and remote access of confined or hazardous spaces. The preliminary visual surface screening offered by UAV agents can then entirely preclude the risks of human access where the target condition is obvious.

However, commercial airborne inspection does not yet satisfy assessment standards where internal flaw status cannot be inferred by exterior visual examination. Manual follow-up must be conducted to deploy detailed volumetric assessment. Thus, hazards and costs associated with human entry cannot yet be completely eliminated. Full realisation of the potential benefits offered by remote airborne inspection therefore requires new technologies to extend the range of NDE processes deployable from a flying platform. These must mirror the dexterous interaction capability used by human technicians to position sensors capable of sub-surface and volumetric imaging in direct contact with a range of target industrial assets. Finally, as a nascent field of research, the resulting airborne inspection platforms and methodologies must demonstrate accurate and reliable measurement, providing fully quantified confidence bounds in order to meet or exceed industrial inspection standards.

1.3 Research Objectives

In light of the motivation from inspection challenges across the energy generation sector and the high-level problem statement above, the following research objectives are proposed:

- Identify in detail current trends presented by airborne inspection practice and related research, noting in particular the shortcomings of current systems and improvements desirable from an industrial perspective.
- Establish the technological factors preventing widespread adoption of contact-based inspection within airborne platforms.
- Review wider UAV and mobile robotics technologies to assess those demonstrating potential to resolve these challenges and develop novel airborne systems capable of effective contact-based asset inspection accordingly.
- Explore opportunities afforded by improved sensor technologies developed for established non-destructive evaluation processes and their deployment from airborne platforms.
- Perform detailed characterisation of the resultant systems so as to qualify their performance in the context of practical applications, supporting widespread industrial adoption and future development.

1.4 Contributions to Knowledge

This thesis develops two distinct UAV system archetypes for contact-based non-destructive evaluation of in situ structures. The contributions made by each towards the overall goal of improving airborne inspection capability are listed as follows.

1.4.1 System I: An Over-Actuated UAV for Dry-coupled Airborne Ultrasonic Thickness Measurement

In response to the challenges facing aerial contact-based ultrasonic NDE, a novel means for automation of the in situ inspection of industrial assets is developed. In particular, this work presents the following contributions:

- Collaborative development of a novel aerial inspection system based on a vehicle with dual-axis tilting propellers in a tricopter architecture for environmental interaction via dynamic thrust vectoring created by researchers from the Autonomous Systems Lab (ASL) of ETH Zurich and Voliro A.G.
- A modified hybrid position-force controller for omnidirectional aerial interaction and deployment of contact based deformable ultrasonic probes.
- Stable and repeatable probe positioning for aerial ultrasonic thickness point measurement of multidirectional targets without an embedded robotic arm or supporting frame.
- A new mode of aerial ultrasonic inspection for scenarios not suited to couplant gel employing a dry-coupled wheel probe.
- Coupling force adaptive ultrasonic signal processing and cross-sectional thickness visualization algorithms for presentation of UAV acquired dynamic scan data.
- A detailed and application focused quantification of the multi-modal aerial ultrasonic NDE positioning and measurement accuracy confidence bounds.

1.4.2 System II: A Surface Crawling UAV for Visual Inspection

Addressing the research objectives posed above, this thesis also proposes a novel UAV-crawler hybrid vehicle that may fly into surface contact with pipes or other cylindrical assets and then redirect its net thrust to support itself while crawling around the structure for full circumferential inspection. Developed from first principles to performance characterisation herein, this novel system presents the following contributions:

- Rigid body modelling of generic pipe interaction and circumnavigation using thrust-supported adhesion.
- Proposition of quantitative vehicle design validation metrics allowing generic feasibility assessment across target asset diameters prior to deployment.
- Formulation of the minimum energy propeller wrench setpoint for pipe interaction exploiting contact forces for additional stability.
- A generic, offline-optimization-based, feed-forward algorithm for interaction control using differential steering to enable 2D surface crawling and fly-crawl-fly inspection.
- Development, construction, and experimental evaluation of a prototype vehicle using this control method.
- Detailed experimental evaluation of remote visual inspection processes using the vehicle, including single image resolution and surface area inspection via radial orthoimage generation.

1.5 Thesis Outline

Beyond this introductory chapter, the remainder of this thesis is constructed as follows:

- **Chapter 2** identifies and examines key trends in remote airborne inspection supporting stated research objectives. This includes a review of the operating principles of remote visual inspection and ultrasonic NDE. The actuation and control of multirotor UAVs is also examined. These topics provide context for review of current industrial activity and published research findings combining the fields in practical application.
- **Chapter 3** develops a system for airborne ultrasonic thickness measurement. This integrates a dry-coupling ultrasonic wheel probe with an aerial manipulator UAV to address shortcomings shown in other contact-based inspection systems. A full system performance characterisation is conducted, profiling accuracy and repeatability for multiple inspection tasks representative of manual processes.
- **Chapter 4** introduces a novel vehicle for contact-based NDE which can access pipes and other cylindrical structures from flight before altering its net thrust output to support stable crawling around the surface in continuous contact. The design principals of the vehicle are examined, formulating a rigid-body model and proposing a quantitative design evaluation metric. Again, a full system performance characterisation is performed to quantify both platform stability and NDE sensory aspects in the context of the application.
- **Chapter 5** concludes the thesis. A summary of the prior chapters is presented alongside implications of key findings and suggested topics for future work.

1.6 Dissemination of Findings

The research conducted throughout this project has given rise to several publicly circulated materials. This section provides a record of these publications for further reference by the reader.

1.6.1 Arising from this Thesis

1.6.1.1 Journal Articles

- [J1] **R. Watson**, M. Kamel, D. Zhang, G. Dobie, C. MacLeod, S. G. Pierce, and J. Nieto, ‘Dry Coupled Ultrasonic Non-Destructive Evaluation Using an Over-Actuated Unmanned Aerial Vehicle’, *IEEE Transactions on Automation Science & Engineering (TASE)*., pp. 1–16, 2021, doi: [10.1109/TASE.2021.3094966](https://doi.org/10.1109/TASE.2021.3094966).
- [J2] **R. Watson**, T. Zhao, D. Zhang, G. Dobie, C. MacLeod, S. G. Pierce, G. Bolton, and A. Joly, ‘A Hybrid Unmanned Aerial Vehicle – Crawler for Full-Contact Airborne Pipe Inspection’, *IEEE Transactions on Mechatronics (T-Mech)*., [Reviewed & Under Revision]

1.6.1.2 Conference Proceedings

- [C1] **R. Watson**, S. G. Pierce, M. Kamel, D. Zhang, G. Dobie, C. MacLeod, and J. Nieto., ‘Deployment of Contact-Based Ultrasonic Thickness Measurements Using Over-Actuated UAVs’, in *European Workshop on Structural Health Monitoring*, Cham, 2021, pp. 683–694. doi: [10.1007/978-3-030-64594-6_66](https://doi.org/10.1007/978-3-030-64594-6_66).
- [C2] **R. Watson**, T. Zhao, D. Zhang, M. Kamel, C. MacLeod, G. Dobie, G. Bolton, A. Joly, S. G. Pierce, and J. Nieto., ‘Techniques for Contact-Based Structural Health Monitoring with Multicopter Unmanned Aerial Vehicles’, in *Proceedings of the 13th International Workshop on Structural Health Monitoring*, Stanford University, Stanford, CA, USA, Mar. 2022, pp. 21–28.

1.6.1.3 Conference Presentations

- [P1] **R. Watson**, M. Kamel, D. Zhang, C. MacLeod, G. Dobie, S. G. Pierce, and J. Nieto, ‘Dry-coupled ultrasonic thickness measurement deployed using an over-actuated unmanned aerial vehicle’, BINDT 58th Annual NDT Conference, Telford, UK, Sep. 03, 2019. [Online]. Available: <https://www.bindt.org/events/NDT-2019/dry-coupled-ultrasonic-thickness-measurement-deployed-using-an-over-actuated-unmanned-aerial-vehicle/>
- [P2] **R. Watson**, T. Zhao, D. Zhang, J. Cao, C. MacLeod, G. Dobie, and S. G. Pierce, ‘Advances in Contact-Based Non-Destructive Evaluation using Unmanned Aerial Vehicles’, BINDT Condition Monitoring Conference, Online, Jun. 16, 2021. [Online]. Available: <https://www.bindt.org/admin/Downloads/5A5%20-%20Watson.Robert.pdf>
- [P3] T. Zhao, **R. Watson**, D. Zhang, J. Cao, E. Mohseni, R. McMillan, C. MacLeod, G. Dobie, P. Crouzen, and C. Forrester, ‘Towards unmanned aerial vehicle-deployed pulsed eddy current for NDT’, BINDT Webinar Week, Online, Sep. 08, 2021. [Online]. Available: <https://www.bindt.org/events/ndt-2021-webinar-week/abstract-6a4/>
- [P4] **R. Watson**, T. Zhao, D. Zhang, G. Dobie, C. MacLeod, and S. G. Pierce, ‘Full-contact and immediate-proximity airborne inspection with a hybrid crawler-multirotor vehicle’, BINDT Condition Monitoring Conference, London, UK, Jun. 07, 2022. [Online]. Available: <https://www.bindt.org/events/cm-2022/abstract-216/>
- [P5] G. Dobie, **R. Watson**, T. Zhao, D. Zhang, C. MacLeod, and S. G. Pierce, ‘Full-Contact and Immediate-Proximity Airborne SHM with a Hybrid Crawler-Multirotor’, European Workshop for Structural Health Monitoring, Palermo, Italy, Jul. 04, 2022. [Online]. Available: https://www.ewshm2022.com/wp-content/uploads/2022/07/Technical-Program_04072022.pdf

1.6.2 Also by the Author

1.6.2.1 Journal Articles

- [J3] D. Zhang, **R. Watson**, G. Dobie, C. MacLeod, A. Khan, and S. G. Pierce, ‘Quantifying impacts on remote photogrammetric inspection using unmanned aerial vehicles’, *Engineering Structures*, vol. 209, p. 109940, Apr. 2020, doi: [10.1016/j.engstruct.2019.109940](https://doi.org/10.1016/j.engstruct.2019.109940).
- [J4] D. Zhang, **R. Watson**, C. MacLeod, G. Dobie, W. Galbraith, and S. G. Pierce, ‘Implementation and evaluation of an autonomous airborne ultrasound inspection system’, *Nondestructive Testing and Evaluation*, vol. 37, no. 1, pp. 1–21, Jan. 2022, doi: [10.1080/10589759.2021.1889546](https://doi.org/10.1080/10589759.2021.1889546).

1.6.2.2 Conference Proceedings

- [C3] D. Zhang, **R. Watson**, G. Dobie, C. MacLeod, and S. G. Pierce, ‘Autonomous Ultrasonic Inspection Using Unmanned Aerial Vehicle’, in *2018 IEEE International Ultrasonics Symposium (IUS)*, Oct. 2018, pp. 1–4. doi: [10.1109/ULTSYM.2018.8579727](https://doi.org/10.1109/ULTSYM.2018.8579727).
- [C4] D. Zhang, **R. Watson**, G. Dobie, C. MacLeod, D. Lines, W. Galbraith, C. Mineo, and S. G. Pierce, ‘Evaluation of Coded Excitations for Autonomous Airborne Ultrasonic Inspection’, in *2019 IEEE International Ultrasonics Symposium (IUS)*, Oct. 2019, pp. 1–4. doi: [10.1109/ULTSYM.2019.8926048](https://doi.org/10.1109/ULTSYM.2019.8926048).
- [C5] D. Zhang, **R. Watson**, J. Cao, T. Zhao, G. Dobie, C. MacLeod, and S. G. Pierce, ‘Dry-Coupled Airborne Ultrasonic Inspection Using Coded Excitation’, in *2020 IEEE International Ultrasonics Symposium (IUS)*, Las Vegas, USA, Sep. 2020, pp. 1–4. doi: [10.1109/IUS46767.2020.9251483](https://doi.org/10.1109/IUS46767.2020.9251483).

1.6.3 Further Output

In addition to the academic publications above, numerous other dissemination activities were conducted as part of this project. Notable items include:

- Public exhibition as a short-listed entry to the Images of Research photography competition promoting research within the University of Strathclyde, 1st May 2019.
- A poster and lightning talk given to leading academics at the “Aerial Interaction and Manipulation: Unsolved Challenges and Perspectives” workshop, hosted within the Robotics Science and Systems Conference, Germany, 23rd June 2019.
- A poster and oral presentation at the 2nd Scottish Ultrasound Annual Scientific Meeting, 28th Feb 2020, winning a “highly commended” prize for the poster based on votes cast by attending academics.
- Representing the University of Strathclyde as one of two PhD researchers selected to participate in the Nuclear Engineering theme of a virtual colloquium with the University of Waterloo, Canada, 12th November 2020.
- Flash-talk presentation to peers at the Research Centre in Non-Destructive Evaluation (RCNDE) Early Career Researchers event, 25th – 26th November 2020, receiving a 2nd place award based on votes cast by attendees.
- Presentation to 100+ industrial and academic RCNDE affiliates as part of the Annual Review event, 29th May 2021.
- Work featured in the “Sensor Enabled Automation” keynote speech given by Prof. S. Gareth Pierce at the Advances in Process Analytics and Control Technologies (APACT) conference, Chester, UK, 14th – 16th September 2022.

CHAPTER 2

A Review of Multicopter Unmanned Aerial Vehicles and Non-Destructive Evaluation

2.1 Introduction

Towards the research objectives concerning NDE of energy generation infrastructure, the work described herein will focus on the multicopter variety of Unmanned Aerial Vehicles (UAVs). Unlike fixed wing varieties, such aircraft present highly beneficial 3D mobility alongside the ability to retain stationary position in mid-air. As emphasised by the case-studies of Section 1.1.5, this airborne vantage point allows rapid visual screening of large areas and structures [96], minimising hazard exposure of skilled workers, while reducing asset downtime. UAV screening then permits more frequent, lower cost inspections, better informing predictive maintenance strategies [12], [97], [98], and so improving overall operational efficiency of industrial facilities.

To further realise these benefits and expand actionable insight, airborne NDE must now consider enhancements to current visual methods and draw upon other established sensing mechanisms that have yet to be fully leveraged in UAV deployment but offer the possibility for traceable, data-driven assessment [99], [100]. Therefore, the most promising technologies for airborne NDE are initially down-selected ahead of a more detailed investigation over the course of this thesis.

2.1.1 Common NDE Methods & UAV Compatibility

Beyond visual inspection, the most widely adopted modes of NDE may be broadly categorised as dye penetrant, magnetic flux leakage, radiography, eddy current techniques, and ultrasonic testing [99], [100]. These technologies also include a number of specialised sub-categories and are each continually evolving to better

address inspection challenges posed by the nature of the item under test and its common failure modes. The field of NDE is thus expansive, with a vast collection of sensors and procedures, but not all are intrinsically suitable for airborne deployment. Accordingly, UAV compatible techniques other than visual testing are identified by high-level consideration of their generic properties.

Of the common NDE methods, radiographic inspection can provide effective means to detect and assess defects within the microstructure of complex components. It is broadly applicable to metals, composites, and most other materials common to energy generation infrastructure and may be applied during manufacture or throughout the asset lifespan [101], [102]. Radiographic volumetric imaging is conducted by monitoring the transit of high energy photons through the target item, with discontinuities in the quantity of radiation received used to infer the presence of defects or incipient structural flaws within the bulk of the material. Further enhancements through Computed Tomography (CT) reduce sensitivity to defect orientation and provide a full 3D view of internal structure, but all forms of radiography share the same requirement for through transmission [99], [100]. Assuming both sides of a large target asset are accessible, the relative pose and alignment of both the radiation source and detector must then be precisely controlled and maintained to enable signal acquisition during each exposure, with further requirement to accurately determine their position relative to the part so as to combine measurements and generate a 3D reconstruction. This is highly challenging for UAV deployment given the motion intrinsic to the floating platform amid aerodynamic disturbances, uncertainty in its estimated pose, and the relative scale of defect features like cracks and corrosive pitting. The mass of a portable X-ray source can also range from around 5 kg to 25 kg [103], [104], entailing a large UAV with substantial payload capacity that may struggle to access obstructed airspace around the targeted industrial plant. Moreover, operations with an airborne X-ray or gamma source would raise significant additional safety concerns beyond typical in situ radiography. While it remains a highly discriminative NDE technique, able to detect a wide range of fine-scale defects throughout the targeted volume, these challenges impose marked restrictions on the practicality of radiography using current UAV systems.

Similarly, widely deployed methods such as dye penetrant and magnetic flux leakage represent excellent strategies to detect surface breaking cracks and structural flaws in individual components but are unlikely to prove effective when applied by UAV [99], [100]. Their execution requires multiple dextrous interactions with the part to prepare the surface, apply contrast agents, generate results, and clean up afterwards that are not possible with current UAV systems. In assessing only near-surface defects, their potential utility is then insufficient to warrant the requisite development effort.

Techniques using eddy currents and similar electromagnetic mechanisms [105] are far simpler to adapt to a robotic system. Their operation through insulative coatings and paint does not require an external coupling medium and entails minimal surface preparation, allowing direct, low-dexterity, deployment without subsequent clean up. They also present excellent sensitivity to fine-scale cracks and discontinuities [99], [100]. However, eddy current methods are incompatible with electrically non-conductive materials and their penetration depth again limits sensitivity beyond near-surface features. Measurements are also relative; in that they detect localised changes across the material but are unable to give the absolute measurements necessary for full evaluation relative to standardised acceptance criteria. This limits their applicability beyond specialised use cases, making other sensor modes more appealing to the nascent field of airborne inspection.

Among remaining established NDE methods, ultrasonic inspection represents a highly attractive candidate for airborne deployment. Using ultrasound, absolute measurements of defect size and position within the full volume of the target structure may be obtained, penetrating large depths or propagating out along its surface. The measurement process is also near instantaneous, lessening issues due to motion of the floating platform shown by other NDE modes. Assessment can be performed with single side access, removing transmitter and receiver alignment issues, and presents no significant hazards to nearby personnel. With the exception of highly scattering materials, ultrasound is also generically applicable across most structural components employed within the energy sector, with no restrictions based on electrical conductivity [99], [100]. Likewise, the lightweight and portable instrumentation

presents a minimal obstacle amid UAV payload restrictions. Airborne ultrasonic testing thus has the potential to be both highly informative and readily implementable.

There are, however, a number of key challenges relating to the deployment of contact-based NDE methods such as ultrasound that must be addressed prior to their widespread airborne adoption, as are examined in greater detail hereafter.

2.1.2 Chapter Structure

The remainder of this chapter is thus composed as follows. Section 2.2 presents and reviews background material concerning the principles of remote visual inspection and its application to 3D reconstruction via photogrammetry. Section 2.3 examines the fundamental operating mechanisms of ultrasonic inspection and briefly considers how these may be applied using a multirotor UAV. Subsequently, Section 2.4 examines how the effects of multiple propellers enable flight in multirotor UAVs. Informed by a rigid body model, common flight control structures are also investigated. A review of current UAV uses for airborne NDE then follows in Section 2.5, including an evaluation of opportunities for meaningful development identified in existing knowledge and practice. The chapter concludes with a final summary in Section 2.6.

2.2 Remote Visual Inspection & Photogrammetry

Visual inspection represents one of the most prolific forms of non-destructive testing. This is a direct consequence of its simplicity and ease of execution relative to the high value of insight offered. Where underlying condition may be inferred from exterior surface features, egregious damage or degradation may thus be quickly identified and addressed in components, machines and structures across industrial sectors [100]. As such, inability to conduct simple visual inspection can necessitate far more invasive and costly forms of NDE to meet operational safety standards [11], [80].

Accordingly, the key theoretical underpinnings of Remote Visual Inspection (RVI) are reviewed, providing context to the advantages and difficulties encountered during its deployment by UAV systems. These principals then form basis for extension to quantitative visual NDE through photogrammetric reconstruction, granting insight into

structural geometry, deformation, erosion, and underlying health beyond that attained through examination of visual information alone.

2.2.1 Remote Visual Inspection

While cost-effective and highly adaptable, the primary downside of manual visual inspection is its extreme subjectivity, exhibiting outcome variation between both different inspectors and multiple reviews by the same inspector. Attributed to their experience level, tiredness, distraction, light levels in the facility, and many other environmental factors, these variances are also often unrecorded where the only outcome of inspection is a qualitative pass-fail decision [100].

To address such shortcomings, images of the target object may be recorded and later accessed, verifying process standards and providing basis for comparative degradation tracking. This readily extends to scenarios where direct human access is not practical, such as amid work-at-height or radiological hazards, with images captured by robotic agents easily retained to provide fully traceable NDE documentation. Subsequent review by offsite experts in the optimum environment further mitigates subjective human variance, leading to improved inspection consistency and more reliable insight into structural health status [99].

Machine vision cameras then constitute the keystone technology of remote visual inspection. They enable low-cost digitised capture of an image of the environment or test piece for expert review, constituting a permanent record of inspection for reporting and assurance purposes. Furthermore, they also support a number of navigation strategies used in mobile robotics where absolute positioning, e.g. from a Global Navigation Satellite System (GNSS), is unavailable. Means to remotely deliver the robot and other NDE sensors to a target in situ location are then provided.

At a high level, operation is very simple: the camera focuses light from within its field of view onto a plane in its rear where the image may be recorded chemically, as in film photography, or digitised by an appropriate sensor, typically a Charge Coupled Device (CCD) or Complementary Metal Oxide Semiconductor (CMOS) [106]. This process

is most simply represented per the diagram of Figure 2-1, depicting the commonly used pinhole camera model.

Manipulation of the projected spatial geometry and the mathematical model, as detailed by Hartley and Zisserman [107], and Corke [108] and reviewed in Appendix A, then enables quantitative visual NDE. As such, images may be processed to track changes in appearance and extract sizing information from defect features or discontinuities within the scene. This provides a traceable and objective dataset, free from variation in human interpretation, and is the theoretical basis for many machine vision applications relevant to visual NDE and UAV inspection, as are examined below.

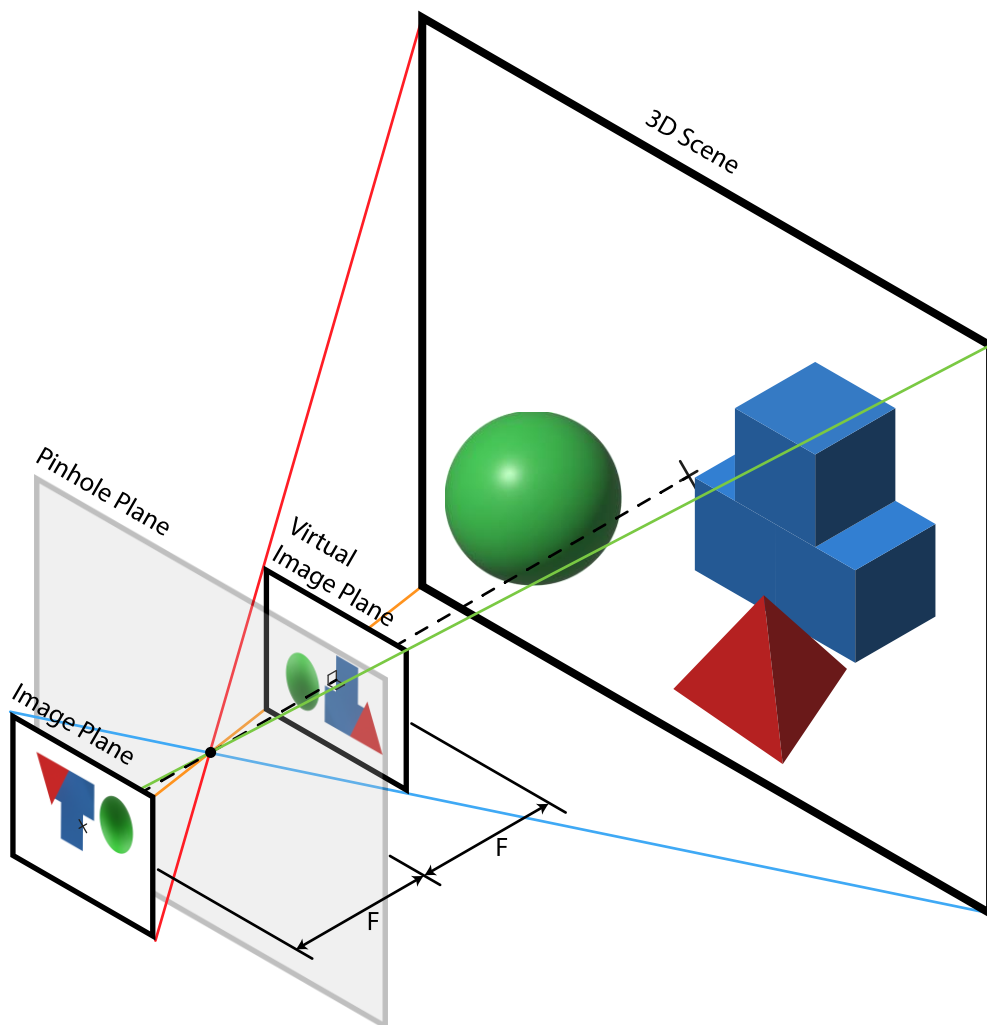


Figure 2-1: A diagram of the pinhole camera model. The 3D scene is projected to the image plane through the pinhole, which lies a focal length F behind the pinhole plane, but is inverted in the process. The virtual image plane exists at same distance in front of the pinhole plane and so depicts the non-inverted image.

2.2.2 3D Reconstruction & Photogrammetry

The extraction of geometric information describing a scene from its depiction in a number of images is referred to as photogrammetry. For NDE this application of machine vision focusses on using multiple images captured at different viewpoints to generate a photorealistic 3D model of the scene. These models may then be employed to compute the dimensions of reconstructed objects, as in stockpile volume monitoring [109], or track changes and deformation in large-scale physical structures, as with wind turbine blade erosion [110]. Digital reconstructions may also be used to spatially encode other NDE data, forming the basis of digital twin and Building Information Modelling (BIM) strategies [111], [112]. In comparison to RVI processes generating numerous isolated and self-similar images, photogrammetric reconstruction then enables assessment of the nature and severity of defect features with full location context via conventional imaging, fusion of multiple images into a single perspective, or other established NDE methods [113], [114].

Underpinning these applications, photogrammetry commonly implements a number of algorithms collectively referred to as Structure from Motion – Multi View Stereo (SfM-MVS) techniques. While highly related, these two algorithm classes have a number of key distinctions and each address a specific aspect of the reconstruction process, as indicated in the popular overview by Smith *et al* [115]. Further, as a major ongoing research theme within the field of computer science, the constituent steps of both the SfM and MVS stages may be performed via any of a number of competing methods, chosen according to the nature of the target, available sensor data, and the desired visual fidelity. A detailed evaluation of the leading approaches and general steps common across such processes is given in Appendix A.

In brief, SfM takes an unstructured set of images and identifies unique visual features in the scene, typically prominent edges, corners, surface textures, colour patches, or other high-contrast features. Each feature is then assigned a descriptor: a numerical vector with dimensions and composition varying according to the algorithm employed. Matching these between images and exploiting the projective sensor model, camera pose at each instant of image capture may be optimally estimated alongside a sparse point cloud representation of the target region, indicating the spatial position of each

feature. Utilising this output, MVS algorithms then create a more densely sampled representation of the target, populating the regions between feature points to make the scene recognisable to a human observer. This densified point cloud may subsequently be used to produce a detailed mesh representation, forming a contiguous shell of multiple interconnected faces that replicates the exterior surface geometry of the target. Transferring visual texture from the source images to the mesh then yields a final photorealistic output, allowing for remote observation of the scene with a flexible digitally rendered viewpoint that can be manipulated to remove degradative lighting or obstructive perspective effects [114], [116], [117].

2.2.3 Reconstruction Sensors & Software

Practical implementation of the SfM-MVS process can take a number of forms. Accordingly, key points regarding commercially available sensor hardware and reconstruction software are noted to provide context regarding both the design of the existing airborne inspection systems, reviewed in Section 2.5 below, and the larger findings of this thesis.

A range of common 3D imaging sensors, as surveyed by Giancola *et al* [118], may enable reconstruction using the operational theory detailed Appendix A. Monocular methods, employing a sequence of images captured by any single mobile camera, are most common with the lowest requirements for specialised hardware. Larger passive stereo cameras (e.g. the ZED 2i [119]) then present an alternative, synchronising image capture between two or more cameras with fixed relative position. In this case, the known transform between cameras provides a scale reference and can improve overall reconstruction accuracy. Beyond this, active RGB-D cameras (e.g. the intel D455 [120]), are capable of detecting both the visual Red-Green-Blue (RGB) colour spectrum and sensing depth via a projected infrared dot matrix or time-of-flight following an encoded light pulse. Reliance on matching unique features in the visible spectrum is then relaxed, allowing robust reconstruction of targets with uniform appearance and in poor lighting conditions. Each sensor family can then provide distinct advantages where its mass, power, and processing requirements are adequately supported by the UAV platform.

Likewise, SfM-MVS software has seen significant technology development since its inception and is now fully commercialised. Multiple toolchains are readily available [121], ranging from open source, research-grade, suites like COLMAP [122], Alice Vision [123], and openMVG [124], to paid commercial options such as Agisoft Metashape [125], Reality Capture [126], and Pix4D [127]. Their relative performance is well documented in literature with comparative review by Kingsland *et al* [128] finding Metashape to be a popular commercial choice, widely used as a standard in evaluation of other tools, whereas Stathopoulou *et al* [129] and Bianco *et al* [130] show COLMAP to perform consistently well amid the open source algorithms. As such, existing software suites may fully support the UAV inspection applications considered in this thesis.

A range of options therefore exist to deliver high-fidelity photogrammetry using data captured by a UAV or mobile robot, with the specific sensor family and SfM-MVS algorithm selected to best match the application context. Airborne photogrammetry then provides adaptable means to quickly record complex geometry over extended target areas with significantly reduced supporting hardware versus traditional surveying techniques [115].

2.2.4 Orthoimage Mapping

The photogrammetry reconstruction process outlined in Section 2.2.2 may also be utilised in production of more conventional 2D site maps, rendering an orthogonal view of the reconstructed mesh from vertically above the landscape, as in Figure 2-2.

Whilst also offering a top-down aerial view, these orthoimages differ from conventional aerial photography by covering a larger area with higher resolution than what is typically attainable from a single image captured by a UAV. At the expense of vertical surface information they are notably better suited to human analysis than the disjointed texture automatically computed as part of the mesh reconstruction, as shown in Figure A-8.



Figure 2-2: An example of an aerially captured orthoimage of a light industrial facility [131]. Minor distortions are visible at the image extremes owing to the repeating texture and incomplete geometry capture of the trees on the property boundary.

Moreover, orthoimages are free of any projective artefacts, arising from the camera model as detailed in Appendix A.1, instead appearing as though all contents are viewed from directly above. For example, where an aerial photograph of a house may see the building walls and roof, skewing its projected footprint, an orthoimage will see only the roof and so may accurately determine structure dimensions.

UAV orthoimages may therefore provide high-fidelity maps without incurring the time and expense of terrestrial survey or the low update frequency of satellite imaging. Accordingly, site-scale orthoimage generation is a common feature of the photogrammetry software packages examined in Section 2.2.3, enabling them to serve as highly effective tools for geolocation of visual and non-visual data. Numerous examples, as reviewed by Zhang *et al* [132], include addition of radar data for urban monitoring [133], and fusion with infrared or hyperspectral images for solar panel health survey [134] and precision agriculture [135]. Benefits of flattening the target surface image are also noted in more limited cases outside of existing UAV applications in the context of pipe internal surface assessment [114], [136], improving defect sizing and bestowing intrinsic localization benefits of 3D reconstruction within existing 2D NDE workflows. Orthophotos are thereby highlighted as a beneficial application of RVI and UAV technologies often overlooked in more general NDE.

2.2.5 Evaluation

In summary, following its identification as a key technology for modern UAV applications, the review above introduces the underlying principles of remote visual inspection. Utilising the pinhole camera model and projective geometry, such images captured by an airborne robot may be employed in a photogrammetric mesh reconstruction workflow, as fully detailed in Appendix A.2. Minimal barriers to practical deployment of this process are presented, with an abundance of compatible imaging hardware and open-source software offered commercially. Results then yield quantitative data, documenting structural status without the subjectivity of direct visual inspection while providing full spatial context to the reported features and discontinuities. Extension of UAV inspection to orthophoto imaging applications is also briefly reviewed, noting an underutilisation outside of geospatial survey despite benefits to defect sizing and visualisation processes. The rich complexity of the data gathered by these airborne evaluations and the insight thereby attainable then motivates continued development to further provide rapid cost-effective inspection.

2.3 Ultrasonic NDE

To aid contextual understanding of existing UAV inspection works and the research findings of this thesis, the fundamentals of ultrasonic inspection are now briefly reviewed. This includes basic wave modes, their propagation through multi-material structures, means for their generation and reception, and how these physical mechanisms may be utilised for sub-surface quantitative NDE. Where appropriate, context is provided to identify aspects directly relating to airborne deployment.

2.3.1 Wave Modes & Propagation Speed

At its most fundamental level, ultrasound is the propagation of energy through a material by means of a mechanical vibration, often in the form a short pulse, transmitted between adjacent matter particles via collisions and internal bonds. In this manner, ultrasound exhibits underlying physics common to acoustic [137] and geophysical applications [138]. The key distinction is then that the primary frequency of these ultrasonic waves lies above the range of human hearing, typically 1 Hz to

20 kHz [139]. For the common purposes of NDE using bulk waves capable of propagating in unbounded media to perform thickness measurement and internal imaging, ultrasound is often employed at frequencies in the range of 100 kHz to 50 MHz [140]. The extremes of this range typically represent more specialist applications, with the frequencies of 100 kHz to 900 kHz serving well for through-transmission assessment of highly attenuative or scattering materials such as concrete [141]. Frequencies over 100 MHz have also been used for medical imaging applications requiring detailed spatial resolution of biological microstructures [142]. Additionally, the lowest frequencies of ultrasound from 20 kHz to 30 kHz and up are widely employed by guided wave testing, using wave modes which rely on spatial boundaries to propagate over longer ranges and rapidly detect defects in large structures such as plates, pipes, beams and rails [143], [144].

In the frequency range around 5 MHz typical to bulk wave NDE of metallic industrial structures, ultrasonic wave packets are most commonly excited in one of two harmonic oscillation modes, forming compression waves or shear waves, as in Figure 2-3.




Figure 2-3: Common wave modes used in ultrasonic NDE shown by their distortion of a 2D grid. The energy propagation direction is indicated by the orange arrows. The direction of particle motion caused by the wave is indicated by the blue arrows.

A compression wave is characterised by alternating regions of rarefaction and contraction between the matter particles comprising the carrier material along its direction of propagation. Also referred to as a “longitudinal” wave, this mode of vibration is defined by a direction of propagation parallel to that of the vibration undergone by a particle within the wave.

A shear wave, by contrast, is characterised by a harmonic displacement of the particles in a direction perpendicular to that of the wave propagation. As such it may also be referred to as a “transverse” wave. Note, however, that in fluids the relatively unconstrained movement of their particles does not support significant shear stress, and they instead experience inelastic deformation. While highly viscous fluids will offer some opposition to this deformation and can permit transmission of very limited shear wave energy, fluids similar to water are commonly idealised as inviscid, with no ability to propagate shear waves [145].

Both wave modes depicted in Figure 2-3 exhibit a consistent relationship between the observed cycle frequency of particle movement, f , the physical distance between similar points within the harmonic cycle termed “wavelength”, λ , and the phase velocity, v , such that they satisfy the equation below.

$$v = f\lambda \quad (2-1)$$

This velocity, however, will vary with both the specific wave mode induced and the mechanical properties of the material through which it propagates. For an ideal fluid, this relationship is expressible using the bulk modulus of the material, K , which describes the resistance to change in volume of a material when subject to equal pressure on all sides [146]. The compression wave velocity, v_c , may then be given as below, where ρ is the material density [147].

$$v_c = \left(\frac{K}{\rho}\right)^{\frac{1}{2}} \quad (2-2)$$

In the case of an elastic solid, support for shear stress alters the expression for longitudinal velocity. Wave speed is then expressed as a function of: Young’s Elasticity Modulus, E , describing the ability of the material to withstand relative changes in length when under uniaxial load [148], and the Poisson Ratio, σ , describing the transverse contraction in a material with longitudinal elongation [149].

The compression wave speed for a given carrier material is then calculated as follows.

$$v_c = \left(\frac{E(1 - \sigma)}{\rho(1 + \sigma)(1 - 2\sigma)} \right)^{\frac{1}{2}} \quad (2-3)$$

Shear wave propagation velocity, v_s , in solid materials may be also expressed as a function of these parameters and the shear modulus, G , denoting the ability of the material to elastically resist transverse deformations and under shear loading [147].

$$v_s = \left(\frac{G}{\rho} \right)^{\frac{1}{2}} = \left(\frac{E}{2\rho(1 + \sigma)} \right)^{\frac{1}{2}} \quad (2-4)$$

A broad approximation is then that the shear mode propagates at half of the velocity of the compression wave [99].

These compression and shear wave modes are most commonly used for dimensional measurement and volumetric scanning within ultrasonic NDE, propagating through the cross-section of the structure as discussed in Section 2.3.4, and are accordingly prioritised in this thesis.

It is noteworthy, however, that inspection strategies may equally utilise more complex wave-modes such as surface or ‘‘Rayleigh’’ waves, where the material particles move in elliptical cycles with both longitudinal and transverse components relative to the wave propagation direction and wave speed is approximately $0.9v_s$ [140]. Complex plate or ‘‘Lamb’’ wave modes, where harmonic oscillations occur in materials of thickness less than one ultrasonic wavelength, are also frequently employed. The ability of these wave modes to propagate along the target structure beyond the region immediately beneath the transducer lends them great utility for large area screening. Rayleigh waves have shown success as a means to characterise plate surface cracking [150], while Lamb waves are well established as the basis for rapid assessment of large metallic plate structures using guided waves [151], with numerous other applications identified in the general review of guided wave ultrasonic testing method provided by Olisa *et al* [152]. Recent investigations into novel applications of lamb waves also include in process monitoring of plate welding [153], delamination detection in a

composite structure [154], and integrity screening of resistance seam welded containers of nuclear material with limited manual access [155]. Such guided wave NDE techniques may present great utility in remote UAV deployment following successful evaluation of the more fundamental methods considered here.

2.3.2 Reflection, Transmission & Refraction

When encountering a boundary between two dissimilar materials, ultrasonic waves of the modes considered above are subject to a number of physical effects. Fundamentally, a portion of the incident wave packet will be transmitted into the new material, while the remainder will be reflected at the interface and continue to propagate back through the original material. The mechanisms for this behaviour thus form an essential component of the functionality of ultrasonic NDE, warranting examination.

As such, it may be shown that the proportion of the wave following either path is dependent upon the relative acoustic impedance of each material, in satisfaction of boundary conditions stipulating continuity of the particle displacement and equilibrium of mechanical stress [156]. For a planar wavefront in an isotropic lossless material, this impedance is denoted Z with the units of $\text{kg/m}^2\text{s}$, equivalently termed Rayl, and expressed as the product of the material density and the propagation speed of the relevant wave mode.

$$Z = \rho v \quad (2-5)$$

The reflection and transmission coefficients are denoted R_{AB} and T_{AB} respectively, with the subscripts indicating the original and secondary materials encountered by the wave in the sequence matching its propagation. Using the acoustic impedance of each material, these provide scaling factors expressing the amplitude of both the reflected and transmitted waves relative to the incident wave. In the case of normally incident planar waves, the expressions are as follows.

$$R_{AB} = \frac{Z_A - Z_B}{Z_A + Z_B}, \quad T_{AB} = \frac{2Z_A}{Z_A + Z_B} \quad (2-6)$$

Further, the reflection and transmission coefficients are related by the equation below, allowing mutual evaluation where one is measured experimentally.

$$1 + R_{AB} = T_{AB} \quad (2-7)$$

As a consequence of (2-6), it may be observed that upon impinging a material of higher acoustic impedance, the reflection coefficient will become negative, modelling a phase shift of 180° in the reflected wave. A special case also exists where there is no secondary material and the initial material is surrounded by a vacuum. Here, Z_B is zero and pure reflection occurs, doubling the vibration amplitude experienced by a particle at the edge of the original material via in-phase superposition of the incident and reflected signals. No energy is transmitted out of the material in such cases.

A pertinent example of the effects of acoustic impedance on wave reflection and transmission is shown when attempting to propagate a longitudinal ultrasonic wave into steel, having acoustic impedance of 45.63 MRayl [140]. From air, with acoustic impedance of 415 Rayl [156], the amplitude transmission coefficient found via (2-6) is 1.82×10^{-5} , indicating near zero energy enters the target. From water, with acoustic impedance of 1.48 MRayl, this rises to 6.28×10^{-2} , while from glycerine, a common ultrasonic couplant with acoustic impedance of 2.42 MRayl, the transmission coefficient is 0.100 [140], vastly improving the propagation of ultrasound into the target and highlighting the need for impedance matching across material interfaces during NDE.

In the case of waves incident at an oblique angle, the boundary effects are more complex, as depicted in Figure 2-4. Therein, ultrasonic waves are labelled according to their wave mode (C being compressive, S being shear) and the numerical index of the material through which they propagate. At an oblique angle of incidence relative to the surface normal, θ_{CA} , the incoming compression wave is no longer reflected back along its original path but propagates at an equal angle of $\theta_{CA'}$ on the opposing side of the boundary normal direction. Accordingly, expressions for transmission and reflection coefficients are more complex than with normal incidence, but may similarly be computed via the boundary displacement consistency and stress equilibrium conditions applied in both directions normal and parallel to the interface [95], [156].




Figure 2-4: Ultrasonic wave reflection, refraction, and mode conversion at the boundary between a fluid and a solid material. In this case, $v_{CB} > v_{CA}$ causing the refracted path to bend away from the boundary normal direction.

Moreover, at non-normal incident angles, the arriving longitudinal wave is mode-converted at the interface, exciting both compression and shear wave modes. Additional propagation pathways are thus generated, with the changing speeds between materials and wave modes further influencing their direction. As a consequence of (2-1), where oscillation frequency is conserved, altered wave speed requires a compensatory change in wavelength to maintain phase consistency across the boundary. The wave path therefore experiences deflection from its incident path, termed “refraction”. The angular direction of each of the resultant transmitted, reflected, and refracted waves is then related by the Snell-Descartes law [99].

$$\frac{v_{CA}}{\sin(\theta_{CA})} = \frac{v_{CA'}}{\sin(\theta_{CA'})} = \frac{v_{SA}}{\sin(\theta_{SA})} = \frac{v_{CB}}{\sin(\theta_{CB})} = \frac{v_{SB}}{\sin(\theta_{SB})} \quad (2-8)$$

This expression confirms the reflection condition that $\theta_{CA} = \theta_{CA'}$ via the equality of the longitudinal velocity in both the incident and reflected waves. Extension beyond the liquid-solid interface of Figure 2-4 to include solid-solid interfaces is also possible. In such cases, the initial material permits further mode-conversion and a reflected shear mode wave, SA' , may appear in the original material at an angle closer to the interface normal.

A further consequence of (2-8) is found in the emergence of critical incidence angles, denoted θ^* , whereupon transit to a faster material sees the refracted wave propagate at 90° to the interface normal, traveling along the boundary surface as an evanescent wave. Such angles exist for both the compressive and shear wave modes and may be calculated using the appropriate velocities in the equation below.

$$\theta_A^* = \sin^{-1} \left(\frac{v_A}{v_B} \right) \quad (2-9)$$

Given omnidirectional coupling, (2-9) thereby sets theoretical limits for the incident angles at which ultrasonic waves may be transmitted into the test piece. For example, at a water-steel interface, the critical angle of longitudinal transmission is 14.4° . Above this, no compression wave can be transmitted into the sample and its use for inspection is prevented. Proceeding further, past the critical angle of transverse wave transmission at 27.0° , generation of a mode-converted shear wave is similarly inhibited. Where both compression and shear waves are evanescent, there is no steady state transmission of energy into the surface and all incident energy is reflected [156], rendering the test piece opaque to ultrasound. This therefore represents an important practical consideration for ultrasonic NDE activity outside of normal surface incidence.

2.3.3 Piezoelectric Transducers

Ultrasound may be generated via a range of physical mechanisms, including laser-based [157], electromagnetic [158], and capacitive methods [159] as further discussed in Appendix B.5. However, the most practiced approach makes use of crystalline elements exhibiting the piezoelectric effect [99], [160] and is examined below.

As subset of ferroelectric materials, these piezoelectric elements possess an electrical polarisation even in the absence of an external field owing to the arrangement of ions within their chemical structure. The distinction from their parent class is then that their crystalline structure will rapidly deform when subjected to an external electrical potential and, vice versa, generate a voltage when subject to mechanical deformation [160]. With appropriate mechanical coupling, this relationship can be exploited to both transmit and receive high frequency ultrasonic signals.

Using the piezoelectric effect, ultrasound is most simply generated by stimulation of transducer elements with a short duration, high-amplitude, negative voltage pulse [161]. This broadband excitation permits the element to naturally enter harmonic oscillation and generate ultrasonic waves matching its structural resonance mode, having a set vibration frequency determined by its mass, stiffness, and physical dimensions [160]. However, the bandwidth of the excitation pulse is far larger than that of the transducer and so only a fraction of the supplied energy is transmitted as ultrasound.

Alternatively, a tone-burst signal of the desired frequency may excite the piezoelectric element [162]. The driving frequency is then matched to the transducers natural resonance mode to maximise energy transfer. Doing so, a lower amplitude excitation may be used to build up the ultrasonic wave amplitude over multiple cycles. This can increase the total energy of the signal while reducing the need for high-voltage transmission electronics. A tone burst also provides narrow-band excitation with high modal purity, as is sought for guided wave testing [163]. Extension to frequency-swept chirp signals provides similar control of frequency content over a wider bandwidth, allowing the simultaneous capture of multiple frequency responses [163]. However, chirps and tone bursts incur an increased excitation time and longer transmitted signal versus simple pulses, impacting near surface detection and resolution of fine defect features in bulk wave applications without compensatory matched filtering [164].

A final category of ultrasonic stimulation is found in pulse compression schemes such as pseudorandom binary sequences, Golay codes, or other coded excitations [164]–[166]. These transmit a signal composed of a complex but distinctive sequence of binary $\{0,1\}$ or bipolar $\{-1,0 + 1\}$ bits, defined based on the ability of their auto-correlation function to approximate an impulse signal. Random noise will not share this property, thus cross-correlating the received ultrasonic signal with the transmitted sequence returns a signal with reduced noise amplitude relative to the standard pulse response. A single transmission can then provide an increased Signal to Noise Ratio (SNR), akin to averaging multiple discrete pulse excitations but without the successive propagation delays. This method for ultrasound generation, however, can increase complexity of the transmission electronics and adds significant signal processing

requirements. The desired excitation scheme, be this a pulse, tone-burst, chirp or encoded signal, must then be selected based on the application context, considering their relative benefits and implementation costs.

To form a practical ultrasonic probe, a piezoelectric element must be combined with suitable electronic instrumentation and a number of other mechanical components as depicted in Figure 2-5. These must function collectively to optimise signal transduction and support NDE activity.




Figure 2-5: Cut-away view of a single element ultrasonic probe with an added delay line. A simplified anatomy is shown to highlight the function of key labelled components. Note that here the ultrasound propagation path returns directly along its incident direction and is only depicted with a U-shape for visualisation purposes.

As indicated, a piezoelectric element forms the core of this single element probe. This may be composed of a number of materials, most commonly the piezoceramic lead zirconate titanate (PZT) and its derivatives, or the piezopolymer polyvinylidene fluoride (PVDF) [167]. Element thickness is typically under 0.5 mm for a 5 MHz probe and must meet tight tolerances to ensure generation of the desired ultrasonic output frequency [160]. Silver electrodes applied during manufacture by electrostatic deposition are used to apply the electric field across the element. These also connect it

to the external instrumentation via a selection of passive electrical components, designed to match the electrical impedance of the probe to the instrumentation and provide limited signal conditioning. Encasement of the active element in an isolation material then helps minimise acoustic noise from internal reverberations, while a metallic probe exterior acts in concert to limit electrical noise.

A second critical aspect of the probe design is found in the backing layer. This is located behind the active element and serves to absorb ultrasonic energy propagating in that direction. Accordingly, its acoustic impedance must be well matched to the active element to facilitate ingress of the incident ultrasonic wave to the material, where high acoustic attenuation or scattering features can then dissipate its energy and prevent repeat transmissions due to internal reflections [99]. The damping influence of the backing layer on the piezoelectric oscillation is also controlled to balance the length of the transmitted wave packet, and thus its spatial resolving capability, with the total transmitted energy according to the needs of the specific application.

A matching layer applied to the front of the active element must meet similar dimensional and impedance criteria to maximise energy transmission into the test piece. This is designed for $1/4\lambda$ thickness at the target ultrasound frequency, allowing additive superposition between the transmitted wave and those internally reflected by either of its interfaces. Acoustic impedance of the matching layer, Z_l , is then set by selecting a material with impedance equal to the geometric mean of the active element impedance, Z_p , and that of the subsequent ultrasonic medium, Z_m , as below [160]. The result is then ideal transmission such that the layer appears transparent to an incident wave.

$$Z_l = \sqrt{Z_p Z_m} \quad (2-10)$$

An optional delay line with a long traversal time may be included between matching layer and test piece to adapt the probe for use with especially thin materials. Its function is primarily to ensure an ultrasonic signal received from the test piece cannot reach the transducer piezoelectric element before it has returned to rest after initial excitation. Additionally, the delay line can act as a thermal barrier, preventing

depolarisation from heat transfer into the probe during applications such as in process inspection of hot metal pipes or tanks, or plastic extrusion and thermal moulding [168].

Instrumentation hardware then completes the functional ultrasonic probe. This is connected to the assembly of Figure 2-5 via shielded coaxial cable to minimise electrical noise and contains the signal transmission and reception electronics necessary to both drive the piezoelectric element and capture the ultrasonic readings. Generally, the transmitter side will consist of a high-voltage DC supply and a high-speed transistor or radio frequency power amplifier [161], [162] whereas the receiver side will include a low noise amplifier, followed by one or more variable gain amplifiers and an analogue to digital converter [169]. Both transmit and receive signal pathways are controlled by embedded logic hardware which will trigger transmission, buffer the recorded samples, perform low-level signal processing, and interface with a display, storage device, or external computer. Beyond these fundamental components, the design of the instrumentation is defined by the features necessary for its application. For example, instrumentation capable of coded excitation will include both a positive and negative voltage supply as well as additional switching components to generate the requisite bipolar signal, whereas instrumentation driving multiple piezoelectric elements may include multiple parallel transmission and reception pathways or use a multiplexer to share hardware between channels [166]. In the context of an embedded system for mobile robotics, instrumentation selection gives strong preference to hardware making the best use of restricted mass, power draw, and signal processing budgets while still meeting functional requirements.

For UAV applications, piezoelectric ultrasound probes can present certain advantages and drawbacks relative to the alternate electromagnetic and optical transduction methods discussed in Appendix B.5. Piezoelectric transducer hardware is typically smaller, lighter and more rugged than optical ultrasound instrumentation systems, better supporting integration amid the payload constraints and high vibration of a UAV platform, but optical methods can operate without physical contact and so preclude challenging airborne interactions. Likewise, generation of ultrasound within a piezoelectric element enables application across most engineering materials with adequate coupling. By contrast, electromagnetic ultrasound methods encounter

limitations in non-conductive materials such as the fibreglass composites of wind turbine blades but can provide intrinsic magnetic adhesion to steel structures to passively stabilise the probe. As the current most established means for bulk wave ultrasound transduction, with mature practical utilisation across multiple industrially qualified manual and robotic NDE processes, significant support thus exists for the use of piezoelectric ultrasound transducers during initial investigation into quantitative airborne NDE. It should be noted, however, that opportunities remain for future exploration of the relative merits of other prominent transduction methods following establishment of the technology.

2.3.4 Thickness Measurement & Flaw Detection

In monitoring structural degradation incurred by various corrosion mechanisms, and tracking the formation or growth of cracks over the useful lifespan of a component, simple ultrasonic measurements using single element transducers have seen wide adoption across industrial sectors and practices [100], [102]. These non-invasive assessments can be performed both in a dedicated facility and in situ, with access to the structure surface from either one or both sides as shown in Figure 2-6.




Figure 2-6: Single-element ultrasound probe configurations for inspection. The blue arrows show the path of the transmitted ultrasound through the target material.

In practice, either the pitch-catch or pulse-echo configuration may be employed to determine both the part thickness and the presence of discontinuities indicative of structural defects. Of these, the pitch-catch case offers the simplest signal interpretation: knowing the transit time between ultrasound pulse generation and reception when the probes are in direct contact with each other, any additional time

correlates directly to the thickness of the test piece via the propagation velocity. Similarly, in the presence of an internal flaw, a component of the ultrasonic wave packet will be obstructed and the received signal amplitude thereby reduced [100].

However, the pitch-catch scheme requires that a second probe be accurately positioned on the rear-surface opposite the transmitter and any motion be synchronised between the two for the duration of the NDE process. Each probe must also be consistently coupled to the sample throughout to accurately detect internal microstructure. These aspects can prove problematic in manual deployments, thus its applicability to UAV utilisation is markedly limited.

Instead, using a pulse-echo configuration enables application where the rear of the structure is inaccessible and relaxes transducer positioning criteria at the expense of more complex signal interpretation. As such, the mechanism for inferring test object geometry and the presence of material is considered in greater detail via Figure 2-7.

This diagram plots the propagation of a normally incident ultrasonic compression wave pulse by its position in a single spatial dimension as a function of time. As such, the speed of the wave per (2-3) is indicated by the line gradient, with slower propagation taking longer to travel the same distance and so appearing with a steeper slope. This also shows how the pulse is split by reflection and transmission at each of the material interfaces per (2-6), and so permits the path of any echo in the characteristic A-scan to be visually interpreted by following its propagation through the graph. Accordingly, the strong reflections of multiple back-wall echoes are easily distinguished, aiding thickness measurement. The skill required to interpret A-scans for flaw detection is also highlighted in consideration of the multiple reflection and transmission permutations shown by waves interacting with the defect. While the defect is clearly shown in the waveform before arrival of the first backwall echo (*BI*), thereafter signal propagation becomes extremely complex, with multiple low-amplitude echoes challenging manual interpretation.

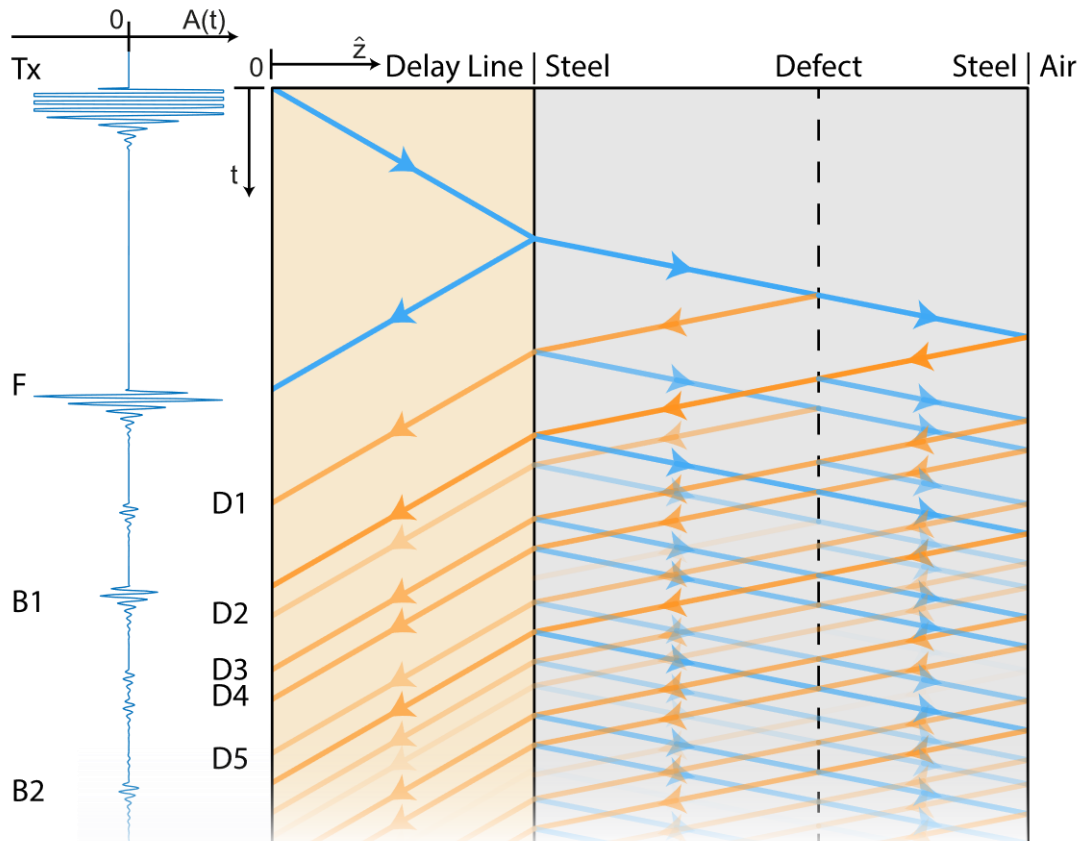


Figure 2-7: A bounce diagram showing ultrasonic wave propagation through a sample containing a defect. The signal trace gives the waveform visible to a “pulse-echo” probe, assuming an ideal matching layer and no reflections from the piezoelectric element itself. Transparency indicates the relative amplitude of each signal pathway. Blue rays are in phase with the original transmission, orange rays have encountered a 180° phase shift during reflection. **Echo Key:** T_x , Transmitted pulse. F , front-wall echo. B , back-wall echo. D , an echo that has been reflected by the defect.

Understanding these principles, pulse-echo thickness measurement may be conducted in one of three modes [170], [171] differentiated by the features of the A-scan waveform they use to determine the transit Time of Flight (ToF). These are outlined as follows, referencing the signal trace of Figure 2-7. Note that a similar approach may be employed to estimate the depth of any sub-surface defect, following identification of the appropriate echo signal.

- **Mode 1:** For probes without a delay line, measure the time from excitation (T_x) to the first backwall echo ($B1$) and subtract an offset time calibrated to account for the matching layer and couplant thickness.

- **Mode 2:** For probes with a delay line, measure the time from the front wall echo at the end of the delay line (F) to the first backwall echo ($B1$). No compensation is necessary.
- **Mode 3:** For probes with or without a delay line, measure the time between repeat backwall echoes ($B1, B2, \dots$ etc.). No compensation is necessary.

Of these three thickness measurement methods, Mode 1 is regarded as the most general, used unless the conditions are such that the others are better suited. Mode 2 can improve minimum thickness resolution over Mode 1 owing to its pairing with a delay line, also showing utility for measurement of sharp radiuses or in confined spaces and in immersion testing, where the water surrounding the component forms the delay line. Mode 3 requires multiple clear back-wall echoes and so is only suited to non-scattering, low-attenuation materials, but also provides the most accurate ToF and so is favoured where accuracy and minimum thickness resolution are critical [170]. Its use, however, is not typically recommended under current standard practice when measuring thin samples with a dual-element probe [171], [172]. This is due to the separation of the transmit and receive elements adding significant length to the ultrasonic propagation path beyond a direct thickness transit, as is examined in greater detail in Appendix B.2.

From these measurements, material thickness, d , is then determined by scaling the ToF to transit the material, T , by the wave propagation speed, v , generally of the longitudinal mode, using the formula,

$$d = \frac{vT}{n} \quad (2-11)$$

Where n is the number of transits made across the thickness of the material of interest. Typically $n = 2$ for pulse-echo thickness measurement, corresponding to travel from the front wall to the back-wall, subsequent reflection, and return to the front wall. This equation thus assumes a one or more direct transits of the material thickness, not considering the effects of complex geometry or any separation between discrete transmit and receive elements.

As one of the most fundamental applications of ultrasonic testing, this permits a broad range of sub-surface measurement without mechanically damaging target structures and so is highly beneficial to structural health monitoring processes and remaining lifespan estimation. Further extensions to produce cross-sectional or volumetric images of internal microstructure, as outlined in Appendix B, also enable detailed monitoring of incipient flaws within target assets, positioning ultrasound as an integral component of modern NDE [95], [99]. Significant applicability to the monitoring of the targeted defect modes described in Section 1.1 is then recognised, constituting a meaningful expansion to airborne inspection capability.

2.3.5 Evaluation

In summary of the review above, the fundamentals of ultrasonic inspection are established in terms of the physical mechanisms of wave propagation and how these may be generated using conventional piezoelectric transducers. Basic principles governing thickness measurement and flaw detection are also introduced, supporting detection of the corrosive, erosive, and cracking defects common across energy generation infrastructure. The capability offered by existing manual ultrasonic inspection processes is thus made clear.

It is notable then that such methods are also well suited to robotic application, with examples at the point of manufacture including deployment via industrial manipulator arms for high throughput quality assurance of composite aircraft wings [173] and wind turbine blades [174], while in situ applications may be supported by magnetic crawlers [175]. Motivations for the selection of ultrasound for further development as a means remote airborne NDE are thereby fully supported.

2.4 Multirotor UAVs

In contrast to fixed wing vehicles of similar scale, multirotor UAVs provide the capability to remain in flight without continuous forward motion. This is highly advantageous for remote monitoring tasks, increasing the precision with which a target area on a structure exterior may be accessed while removing the challenges of data capture posed by continuous dynamic motion.

This section therefore reviews the accepted operational principles of a conventional quadrotor UAV. To this end, the physical behaviour of a single propeller actuator is first introduced, drawing from established aerodynamic theory. The collective function of multiple propellers mounted to a multirotor airframe is then considered via a vector-based framework, including a full dynamic model for vehicle behaviour in free flight. An overview of common control architectures underpinned by this model follows, with a brief note on other relevant trends shown by contemporary UAVs concluding the section.

Note that mathematical analysis within this section employs two coordinate frames:

- $\{W\}$, the world frame which remains static relative to the environment.
- $\{B\}$, the body frame which is fixed to the UAV at its Centre of Mass (CoM) and so translates and rotates relative to $\{W\}$ during flight.

2.4.1 Propeller Thrust Generation

Prior to the initial development of multirotor UAVs, the basic principles of propeller thrust generation were well established in manned aeroplanes and helicopters, permitting an actuation model for an individual propeller to be obtained directly from aerodynamics textbooks [176]–[178]. Therein, thrust behaviour is derived via blade element momentum theory, classifying the rotor as an infinitesimal thickness actuator disk, and applying the Bernoulli equation across its steady-state pressure differential.

Examined within $\{B\}$, as depicted in Figure 2-8, this approach yields a simple expression for static propeller behaviour in terms of two vectors denoting thrust, $\mathbf{F}_{p,i}$, and aerodynamic drag torque, $\boldsymbol{\tau}_{d,i}$, aligned with the rotor axis, $\hat{\mathbf{a}}_i$. Note that a trailing numerical subscript may be later used to refer to a specific rotor, but is replaced by i in this generic case.




Figure 2-8: A diagram of the primary forces and torques exerted on the UAV body by a single propeller in static conditions, arising from the aerodynamic effects of its rotation. This generic rotor has position \mathbf{p}_i relative to the body frame $\{B\}$.

Accordingly, thrust varies as a function of propeller rotational speed, Ω_i , the density of air, ρ , the area, A , swept by the propeller of radius r , and the propeller lift force coefficient, C_F , itself dictated by its aerofoil cross-section and pitch angle relative to the motor axis.

$$\mathbf{F}_{p,i} = C_F \rho A r^2 \Omega_i^2 \cdot \hat{\mathbf{a}}_i = C_T \Omega_i^2 \cdot \hat{\mathbf{a}}_i \quad (2-12)$$

Recognising that the propeller aspects are consistent, and assuming only very minor relative change in air density due to environmental conditions, a simplified expression of (2-12) may group these parameters using a single thrust coefficient, C_T , as above. When mounted to an airframe with the rotor centred at position \mathbf{p}_i relative to the CoM and origin of $\{B\}$, the axial thrust will then generate a torque given as follows.

$$\boldsymbol{\tau}_{m,i} = \mathbf{p}_i \times \mathbf{F}_{p,i} \quad (2-13)$$

The propeller aerofoil will also generate an aerodynamic drag torque, expressed similarly to (2-12) using the aerofoil drag coefficient, C_q , and its lumped value, C_M . This drag torque acts about the rotor axis in the spin direction opposing Ω_i , denoted $s_i \in \{-1, +1\}$ per the right hand rule.

$$\boldsymbol{\tau}_{d,i} = s_i C_q \rho A r^3 \Omega_i^2 \cdot \hat{\mathbf{a}}_i = s_i C_M \Omega_i^2 \cdot \hat{\mathbf{a}}_i \quad (2-14)$$

Thus, summation yields the collective propeller torque acting on the UAV airframe.

$$\boldsymbol{\tau}_{p,i} = \boldsymbol{\tau}_{m,i} + \boldsymbol{\tau}_{d,i} \quad (2-15)$$

Equations (2-12) and (2-15) thereby describe action of the propeller in an idealised static atmosphere. However, as first observed in helicopter blades [176]–[178] then recognised in the context of UAVs by Pounds *et al* [179], [180], motion of the surrounding air across the rotors may introduce disturbance from the above behaviour. This change in airflow can occur equally due to environmental wind or motion of the UAV, leading most prominently to rotor flapping, a phenomenon intrinsic to all rotor craft.

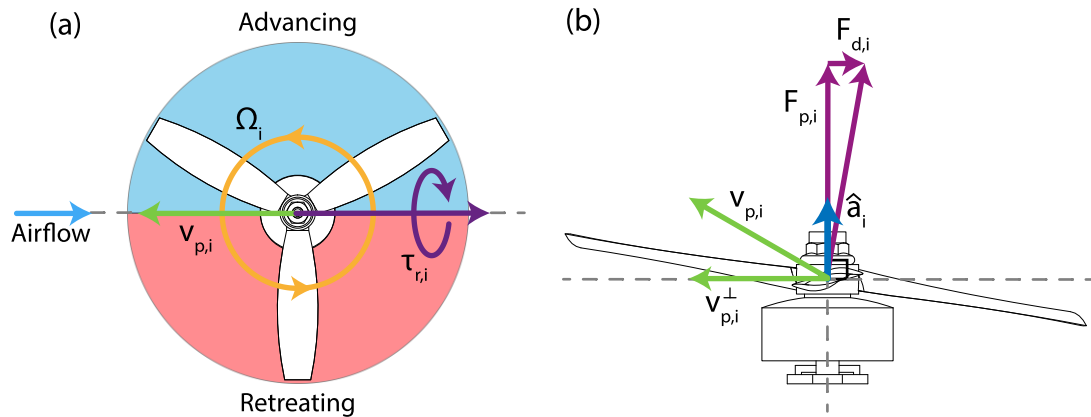


Figure 2-9: Secondary propeller aerodynamic effects arising from rotor flapping. (a) Generation of roll torque, $\tau_{r,i}$. (b) Generation of rotor drag force, $F_{d,i}$.

Occurring cyclically, as depicted in Figure 2-9, rotor flap sees the propeller blades rise and fall relative to the axis-normal plane during each rotation. It occurs because, when translating through the air, a spinning non-rigid propeller blade advancing into the airflow will possess higher relative tip velocity and thus increased lift versus the retreating blade. This results in a “roll torque” acting about the direction of the apparent wind, also causing out-of-plane deflection in the blade. Notably, this roll torque may be negated by two or more propellers of opposing spin within a multirotor UAV, provided their rotation speeds and encountered airflow are similar.

Effects due to blade deflection, however, will remain and may be further considered by representation as a torsional spring system. Accordingly, it may be shown that a

rotor blade follows driven damped harmonic oscillation and the natural frequency of its flapping equals propeller spin speed Ω_i . Its deflection response will thus lag the advancing and receding lift changes by 90° , causing the blade tip to rise to peak height when pointing into the airflow then fall to its nadir when in the position opposite, as in Figure 2-9(b). Because propeller thrust acts in the direction normal to the rotor plane, rotor flap thereby generates a thrust component acting to arrest the propeller translation relative to the surrounding air. This effect is known as “rotor drag”, with a full derivation available in popular helicopter textbooks [176]–[178].

The net effects of rotor flapping are succinctly expressed by Furrer *et al* [181], [182] as below, in terms of a drag force, $\mathbf{F}_{d,i}$, and rolling torque, $\boldsymbol{\tau}_{r,i}$, as functions of the propeller centre velocity, $\mathbf{v}_{p,i}$, having a component, $\mathbf{v}_{p,i}^\perp$, perpendicular to the rotor axis.

$$\mathbf{F}_{d,i} = -\Omega_i C_D \cdot \mathbf{v}_{p,i}^\perp \quad (2-16)$$

$$\boldsymbol{\tau}_{r,i} = -\Omega_i C_R \cdot \mathbf{v}_{p,i}^\perp \quad (2-17)$$

Positive constants C_D and C_R are the rotor drag constant and rolling torque constant, respectively. In addition to depending upon the aerodynamic parameters that influence C_T and C_M , these will also vary with stiffness and other mechanical properties of the rotor, with stiffer rotors taking larger values of C_R but experiencing less deflection and so having lower C_D .

The effects represented in (2-16) and (2-17) are included in high-fidelity simulation of multirotor UAVs [108], [181] but are often neglected for control design in settings outside of highly dynamic trajectory tracking where \mathbf{v}_p is small relative to tip velocity [183]–[185]. Accordingly, they are regarded as minor unmodelled disturbances in applications such as airborne inspection, where slow regular motion better accommodates data capture.

2.4.2 Collective Actuation

A quadcopter multirotor possesses a total of four propeller actuators, as show in Figure 2-10. Each acts independently on the airframe as described above, with their net effect used to stabilise the vehicle attitude and position and thereby govern flight.




Figure 2-10: Quadcopter actuation and free-body diagram. Propeller pairs with consistent spin direction are grouped by colour. Those coloured in orange spin clockwise, with counter-clockwise drag torque, and vice versa for blue.

The collective behaviour of these propellers is expressed as a net thrust, F_p , and torque, τ_p , calculated by summation of the individual rotor contributions in the body frame $\{B\}$, for ease of computation. Note that it is conventional to employ a Front-Right-Down (FRD) sense for the x-y-z directions of the $\{B\}$ frame, with the $\{W\}$ frame following a similar North-East-Down scheme per ISO 1151 standards [186]. As such F_p required for stable horizontal hover will point in the negative z-axis of both frames.

In a conventional quadrotor all propeller axes are therefore parallel to the negative $\{B\}$ z-axis, granting Uni-Directional Thrust (UDT) for maximum hover efficiency but

permitting direct translation control in only one direction. As such, in this specific case, \mathbf{F}_p is calculated by summation over a single axis, substituting (2-12) to grant a function of each rotor speed as follows.

$${}^B\mathbf{F}_p = \begin{bmatrix} 0 \\ 0 \\ -\sum_{i=1}^4 F_{p,i} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -C_T \sum_{i=1}^4 \Omega_i^2 \end{bmatrix} \quad (2-18)$$

Translation in $\{W\}$ is then achieved by manoeuvring the airframe to point the negative $\{B\}$ z-axis, per the FRD convention, in the desired direction of travel.

Such orientation changes are performed by manipulating the net torque, $\boldsymbol{\tau}_p$, which may be generated about each of the $\{B\}$ x-y-z axes to grant full control of airframe rotation under the quadrotor layout depicted in Figure 2-10. This behaviour is examined on an axial basis to derive the components of the $\boldsymbol{\tau}_p$ vector in terms of the dominant aerodynamic effects. Secondary effects such as rotor drag and roll torque are omitted in light of the slow-moving nature of the application and common practice.

Acting around the $\{B\}$ x-axis, the signed scalar torque component, $\tau_{p,x}$, is generated via the moment of the propeller thrusts on either side of this axis, causing the craft to roll through angular displacement ϕ . As depicted in Figure 2-10, propellers 2 and 3 will generate positive torque, while 1 and 4 contribute opposing negative torque per (2-13).

Given that the quadrotor propellers are centred symmetrically on the $\{B\}$ x-y plane, forming the vertices of a square each a distance p from the origin, it may be shown that $\boldsymbol{\tau}_{m,i}$ acts with a moment arm length $\frac{p}{\sqrt{2}}$ from the x-axis. Thereby, $\tau_{p,x}$ is given by

$$\tau_{p,x} = \frac{p}{\sqrt{2}} (-F_{p,1} + F_{p,2} + F_{p,3} - F_{p,4}) \quad (2-19)$$

Substitution of (2-12) then yields a function of Ω_i

$$\tau_{p,x} = \frac{p}{\sqrt{2}} C_T (-\Omega_1^2 + \Omega_2^2 + \Omega_3^2 - \Omega_4^2) \quad (2-20)$$

A similar analysis applies to net torque component $\tau_{p,y}$, acting about the $\{B\}$ y-axis to cause pitch through angular displacement θ . In this case propellers 1 and 3 contribute positive torque via the expression below.

$$\tau_{p,y} = \frac{p}{\sqrt{2}} C_T (\Omega_1^2 - \Omega_2^2 + \Omega_3^2 - \Omega_4^2) \quad (2-21)$$

In contrast, the torque component about the $\{B\}$ z-axis, denoted $\tau_{p,z}$, is not influenced by $\tau_{m,i}$ but is instead dependent on the rotor drag torque, $\tau_{d,i}$, as in (2-14). As the rotor axes are all antiparallel to the $\{B\}$ z-axis and diagonally opposing rotors share spin direction, $\tau_{p,z}$ may be generated independently of x-y torque components by altering the distribution of total thrust among these diagonal pairs. The resulting torque is expressed as follows.

$$\tau_{p,z} = C_M (-\Omega_1^2 - \Omega_2^2 + \Omega_3^2 + \Omega_4^2) \quad (2-22)$$

The full force-torque actuation capability of a quadrotor is thereby defined by combination of (2-18), (2-20), (2-21) and (2-22). Taken together, these expressions form an actuation matrix, \mathbf{A}_p , concisely expressing the net wrench of the propellers, $[\mathbf{F}_p^T \quad \boldsymbol{\tau}_p^T]^T$, as a function of the column vector containing each squared motor speed, $\boldsymbol{\Omega}^2 = [\Omega_1^2 \quad \Omega_2^2 \quad \Omega_3^2 \quad \Omega_4^2]^T$, as follows

$$\begin{bmatrix} \mathbf{F}_p \\ \boldsymbol{\tau}_p \end{bmatrix} = \mathbf{A}_p \boldsymbol{\Omega}^2 \quad (2-23)$$

Where this actuation matrix is composed from the sub-matrices specific to the force and torque components \mathbf{A}_F and \mathbf{A}_τ , respectively.

$$\mathbf{A}_p = \begin{bmatrix} \mathbf{A}_F \\ \mathbf{A}_\tau \end{bmatrix} \quad (2-24)$$

$$\mathbf{A}_F = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -C_T & -C_T & -C_T & -C_T \end{bmatrix} \quad (2-25)$$

$$\mathbf{A}_\tau = \begin{bmatrix} -\frac{p}{\sqrt{2}}C_T & \frac{p}{\sqrt{2}}C_T & \frac{p}{\sqrt{2}}C_T & -\frac{p}{\sqrt{2}}C_T \\ \frac{p}{\sqrt{2}}C_T & -\frac{p}{\sqrt{2}}C_T & \frac{p}{\sqrt{2}}C_T & -\frac{p}{\sqrt{2}}C_T \\ -C_M & -C_M & C_M & C_M \end{bmatrix} \quad (2-26)$$

From the structure of \mathbf{A}_p , it is apparent that the UAV may exert control on only four of the six Degrees of Freedom (DoF) in which it operates, corresponding to full 3D orientation control but only single DoF translation, as depicted in Figure 2-11. The conventional quadcopter is therefore termed “under-actuated”, and so must complete a yaw, pitch or roll manoeuvre to enable arbitrary translation in $\{W\}$. Note also that, as C_M is typically much smaller than the $\frac{p}{\sqrt{2}}C_T$ coefficient, control authority is weakest in the yaw axis. However, governance of quadcopter position by pitch and roll ensures this is of minor consequence to free-flight applications.

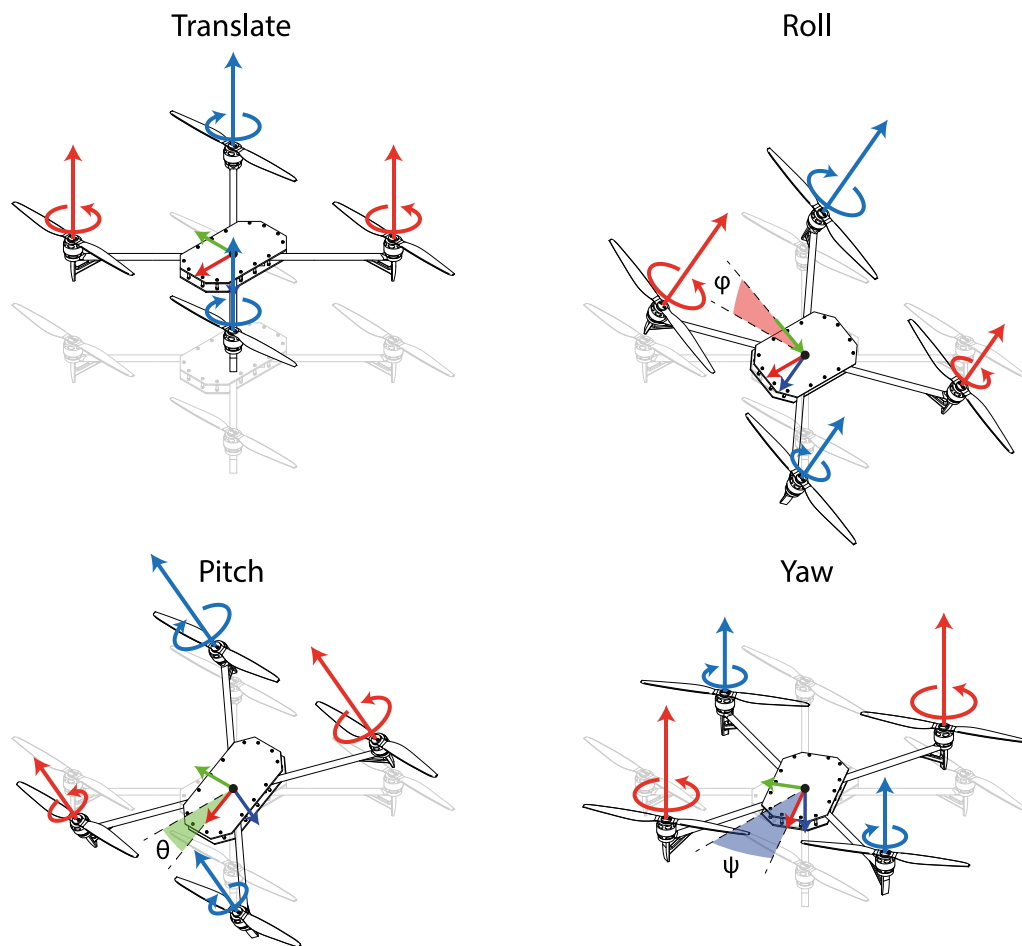


Figure 2-11: Axes of motion for a conventional quadrotor UAV shown relative to a static hover pose in grey. By altering the relative thrust of each propeller, indicated by arrow size, it may directly translate in one axis while rotating about a further three in the roll, pitch, and yaw directions. Angular displacement is coloured according to its rotation axis.

2.4.3 Free-Flight Model

With the actuation principals of a multirotor established, a full model of the system must now consider motion relative to its operating environment, incorporating propeller effects on both 3D translation and rotation. This is considered via two components: mathematical representation of UAV pose, and rigid body dynamics.

Note that that the notational framework used to encode motion during modelling is also employed when examining the recorded trajectory of a physical craft.

2.4.3.1 Expressing UAV Pose

To this end, instantaneous UAV pose, given by the body coordinate frame $\{B\}$, must be expressed relative to $\{W\}$. For pure linear translation this is easily achieved using a position vector, such as ${}^W_B\mathbf{O}$ defining the origin of $\{B\}$ with respect to $\{W\}$.

Rotation may then be represented by a number of methods. Under Euler's Rotation Theorem [187], any arbitrary orientation may be represented by three parameters. Consequently, any two coordinate frames may be related by a series of no more than three sequential rotations about a set of orthogonal axes, where no two sequential rotations use the same axis, e.g. XYZ or XYX. The degrees of rotation about each axis are then referred to as Euler angles. Any sequence of Euler may equivalently express a given orientation, but the specific rotation necessary about each axis can alter in value dependent on the sequence employed [108].

For aeronautics, a ZYX sequence is conventional [108], describing aircraft yaw, ψ , pitch, θ , and roll, ϕ , as depicted in Figure 2-10, according to the right-hand-grip convention. Starting with a $\{B\}$ frame aligned to $\{W\}$, the UAV pose may then be intuitively constructed by three rotations applied to $\{B\}$ in ZYX sequence. This starts with a yaw rotation by ψ about the UAV body frame z-axis to set a global heading, followed by a pitch rotation by θ about the $\{B\}$ y-axis resulting from the previous step to set elevation above horizontal, then finally a roll rotation by ϕ about the new x-axis [108] to yield the final orientation. As such, this aeronautics convention represents a sequence of intrinsic rotations, where each transform is applied about the axes resulting from the previous step. The same resultant pose can equivalently be expressed by an extrinsic rotation sequence, applying the same component rotations about the fixed world axes in the reverse order. Note also that, as this ZYX convention uses all axes, it is categorised as a Tait-Bryan or Cardanian sequence, differing from the Eulerian sequences which non-sequentially repeat an axis (e.g. ZXZ).

Mathematically, the yaw-pitch-roll orientation may be written as a rotation matrix, expressing the orientation of $\{B\}$ relative to $\{W\}$ [108]:

$$\begin{aligned} {}^W\mathbf{R} &= \mathbf{R}_Z(\psi) \mathbf{R}_Y(\theta) \mathbf{R}_X(\phi) \\ {}^W\mathbf{R} &= \begin{bmatrix} c_\psi & -s_\psi & 0 \\ s_\psi & c_\psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c_\theta & 0 & s_\theta \\ 0 & 1 & 0 \\ -s_\theta & 0 & c_\theta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_\phi & -s_\phi \\ 0 & s_\phi & c_\phi \end{bmatrix} \\ {}^W\mathbf{R} &= \begin{bmatrix} c_\psi c_\theta & c_\psi s_\theta s_\phi - c_\phi s_\psi & s_\psi s_\phi + c_\psi c_\theta s_\theta \\ c_\theta s_\psi & c_\psi c_\phi + s_\psi s_\theta s_\phi & c_\phi s_\psi s_\theta - c_\psi s_\phi \\ -s_\theta & c_\theta s_\phi & c_\theta c_\phi \end{bmatrix} \end{aligned} \quad (2-27)$$

Where c_x and s_x represent $\cos(x)$ and $\sin(x)$, respectively. As such, the rotation matrix above may be used to convert the direction of a unit vector in the body frame, ${}^B\hat{\mathbf{u}}$, to the same direction expressed relative to the world frame, ${}^W\hat{\mathbf{u}}$.

$${}^W\hat{\mathbf{u}} = {}^W\mathbf{R} {}^B\hat{\mathbf{u}} \quad (2-28)$$

The relationship given in (2-27) may also be inverted to permit human interpretation of orientations expressed by a rotation matrix. However, issues arise in certain orientations where the rotation encounters a singularity and there is no unique Euler angle representation. This phenomenon is termed ‘‘gimbal lock’’ after the behaviour of a mechanical gyroscope, occurring where two of its concentric rings align and it loses one degree of rotational freedom [108], [188]. The ZYX convention is chosen such that this scenario is rare in flight, occurring only at pitch of $\pm 90^\circ$, where ability to discern yaw from roll is lost and any UAV model or control system based on Euler angles will behave erratically.

While the rotation matrix itself remains valid during gimbal lock, it can be impractical to work with in a real-time embedded environment owing to the nine elements of a 3x3 rotation matrix. This redundancy versus Euler’s Rotation theorem increases processing and storage requirements as well as sensitivity to numeric error accumulation, causing distortion on application of the matrix that is challenging to undo [189].

Alternately, quaternions may be used to represent orientation, as proposed by Hamilton over 150 years ago [190]. These are four-dimensional tuples, formed of a real scalar and three orthogonal complex numbers, denoted \mathbf{i} , \mathbf{j} , \mathbf{k} . By definition, the complex numbers are such that

$$\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1 \quad (2-29)$$

The four elements forming a quaternion, \mathbf{q} , may then be thought of as a scalar, q_0 , accompanied by a complex vector, \mathbf{q}_v , and represented as follows.

$$\mathbf{q} = q_0 + q_1\mathbf{i} + q_2\mathbf{j} + q_3\mathbf{k} = \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{bmatrix} = \begin{bmatrix} q_0 \\ \mathbf{q}_v \end{bmatrix} \quad (2-30)$$

Rotations are described via unit quaternions, having magnitude $\|\mathbf{q}\| = 1$, and applied via the quaternion product operator, as is widely detailed in robotics literature [108], [191], [192]. Interpretation of the result may be aided by expression in axis-angle form as below, where unit vector, $\hat{\mathbf{u}}$, defines the axis of rotation and, θ , the angular displacement around it.

$$\begin{bmatrix} q_0 \\ \mathbf{q}_v \end{bmatrix} = \begin{bmatrix} \cos(\theta/2) \\ \hat{\mathbf{u}} \sin(\theta/2) \end{bmatrix} \quad (2-31)$$

As such, orientation represented using quaternions uses only one redundant parameter beyond the requirements of Euler's Rotation theorem, thereby avoiding the singularities of Euler angles, but also offering a significant reduction in storage and computation of chained operations versus rotation matrices. Accumulated numerical error may also be mitigated by renormalising to unit magnitude, a much simpler operation than the rotation matrix equivalent [189].

This thesis therefore makes use of roll-pitch-yaw angles, rotation matrices and quaternion representations, leveraging the advantageous properties of each according to the situation.

2.4.3.2 Rigid Body Dynamics

Employing these representations of pose, flight of a multirotor UAV may be modelled by the behaviour of a single rigid body acted upon by the propeller actuators as modelled in Section 2.4.1. It is then possible to compute the full UAV translation and rotation trajectory response to both its own thrust and external forces or torques.

Accordingly, equations of motion describing the flight of the rigid UAV body may be derived through established classical mechanics, documented in detail by modern textbooks [193], [194]. Approaches may equivalently involve Lagrangian dynamics, where the energy of the system is considered in generalised coordinates, or Newton-Euler dynamics, where actions of force and torque vectors are examined directly under conservation of momentum, typically in Cartesian coordinates. Having been used by Altug *et al* to derive one of the earliest explicit models of a quadrotor UAV [195], the Newton-Euler method was similarly adopted in other seminal works by Hamel *et al* [196], Pounds *et al* [197], [179], [198], and Bouabdallah *et al* [199], [200]. It remains more popular than the Lagrange approach shown by Castillo *et al* [201], [202] around the same time. Despite otherwise comparable control performance, this may be attributed to the additional insight and capability for expansion offered by explicit modelling of wrench vectors with the Newton-Euler approach [193], [203].

Accordingly, the Newton-Euler dynamics of a multirotor UAV are presented via the canonical equation below, where all parameters are defined relative to the UAV body affixed frame $\{B\}$, as indicated by the leading superscript.

$$\begin{bmatrix} m\mathbf{I}_{3 \times 3} & 0 \\ 0 & {}^B\mathbf{J} \end{bmatrix} \begin{bmatrix} {}^B\dot{\mathbf{v}} \\ {}^B\dot{\boldsymbol{\omega}} \end{bmatrix} + \begin{bmatrix} {}^B\boldsymbol{\omega} \times m^B\mathbf{v} \\ {}^B\boldsymbol{\omega} \times {}^B\mathbf{J}^B\boldsymbol{\omega} \end{bmatrix} = \begin{bmatrix} \sum {}^B\mathbf{F} \\ \sum {}^B\boldsymbol{\tau} \end{bmatrix} \quad (2-32)$$

Here, m is the UAV mass, $\mathbf{I}_{3 \times 3}$ is an identity matrix, \mathbf{v} represents the translational velocity, $\boldsymbol{\omega}$ is the rotational velocity, and \mathbf{J} is the moment of inertia tensor. Note that whereas \mathbf{J} varies with the rigid body pose when calculated in $\{W\}$, its expression in the $\{B\}$ frame is invariant, simplifying evaluation of rotational dynamics significantly. Finally, the stacked net force and torque vectors, $\Sigma\mathbf{F}$ and $\Sigma\boldsymbol{\tau}$, respectively, together describe the total effect of all wrench acting on the UAV airframe.

In free flight, this will include the total propeller wrench, as given in (2-23), and the weight of the UAV, $m\mathbf{g}$. Note that the gravity vector, \mathbf{g} , lies parallel to the $\{W\}$ z-axis per the NED coordinate convention, so is expressed in $\{B\}$ using the rotation matrix giving the current pose relative to $\{W\}$, ${}^W_B\mathbf{R}$, as follows

$$\begin{bmatrix} \sum {}^B\mathbf{F} \\ \sum {}^B\boldsymbol{\tau} \end{bmatrix} = \mathbf{A}_p\boldsymbol{\Omega}^2 + m {}^W_B\mathbf{R}^T {}^W\mathbf{g} \quad (2-33)$$

Note that this expression can be extended to explicitly include other influences on the UAV such as the rotor drag and propeller roll torque described in Section 2.4.1, secondary aerodynamic effects like wind and air resistance [181], [183], or contact forces arising from environmental interaction [204].

The full equations of motion for the UAV trajectory may then be found by substitution of (2-33) into (2-32), re-arranging to extract acceleration terms $\dot{\mathbf{v}}$ and $\dot{\boldsymbol{\omega}}$, and solving the resultant differential equations for translation and rotation in the world frame.

Such equations of motion now form the basis of a range of UAV modelling and simulation software as surveyed by Ebeid *et al* [205], including free open-source packages such as Gazebo [206], JMAVSim [207], or AirSim [208], and commercial products like MATLAB/Simulink [209], CoppeliaSim [210], and Nvidia Isaac Sim [211]. These software tools offer excellent support for development of novel UAV behaviours and control algorithms without the risks or expense of physical hardware.

2.4.4 Flight Control Structures

A final critical component of UAV operation is found in the control system. As a multirotor is not intrinsically stabilised by aeromechanical means, active control is necessary to maintain orientation and altitude or track a desired trajectory. The overall purpose of the control structure is then to determine the propeller motor speeds necessary to perform the requested action in light of the difference between the current UAV state and its desired value, as illustrated in Figure 2-12.

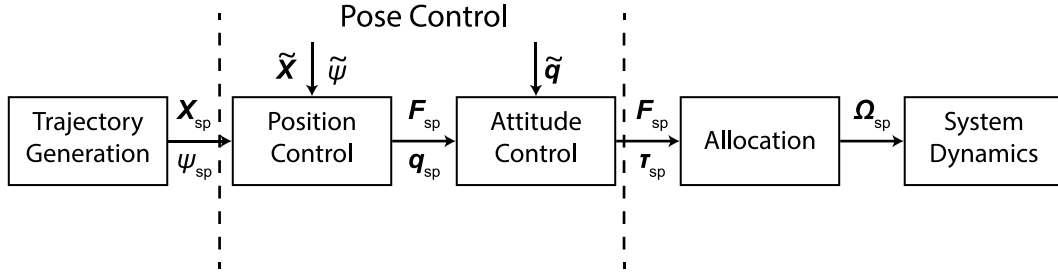


Figure 2-12: Generic multirotor nested flight control architecture. Position setpoint, \mathbf{X}_{sp} , and yaw setpoint, ψ_{sp} , are first set by the trajectory generation module. Pose control then computes the desired thrust, \mathbf{F}_{sp} , attitude quaternion, \mathbf{q}_{sp} , and torque setpoint, $\boldsymbol{\tau}_{sp}$ based on feedback of the current state. Note the attitude control sub-module is typically aware of the desired thrust but will not alter its value. The allocation module assigns the desired wrench to propeller motor speeds, $\boldsymbol{\Omega}_{sp}$, which act within the system dynamics to move the UAV. Estimated state parameters are indicated by a \sim decorator.

Within this structure, the flight trajectory may be manually set by pilot control input or generated automatically by a higher-level path-planning algorithm for processes like area grid survey. For inspection tasks around complex 3D geometries such algorithms can take many forms, utilising both prior knowledge of the asset or pure exploratory search methods as reviewed in detail by Aggarwal and Kumar [212].

The pose control algorithm then determines the action to be taken to follow this trajectory. Typically, a nested loop structure is used, with the inner loop governing the attitude dynamics while the outer loop regulates position. This assumes that the inner loop is exponentially stable, having higher bandwidth and faster dynamics than the outer loop, as is commonly the case for multirotor UAVs. The outer loop may then be designed and tuned independently, simplifying the overall control algorithm significantly [213]. In modern works, the controller structure will often be independent of the specific propeller layout and return a generic force-torque vector output.

Since the establishment of the field in the early 2000s [195]–[197], [199]–[201], a number of control methodologies have been adopted to regulate the inflight pose of multirotor UAVs. In popular approaches, these range from comparatively simple Proportional Integral Derivative (PID) control algorithms, through model-based optimisation methods employing a Linear Quadratic Regulator (LQR) or Model Predictive Control (MPC), to well-established non-linear control structures like Backstepping Control (BSC) or Sliding Mode Control (SMC). With abundant

architectural permutations inside these categories and beyond, an extensive range of multirotor control strategies now exist, as highlighted in literature surveys of the topic [214]–[217]. Accordingly, further details of the operating principles and pertinent functional attributes associated with the more popular methods are provided in Appendix C, as is summarised below in Table 2-1.

TABLE 2-1: COMPARISON OF COMMON UAV CONTROL METHODS

Method	Advantages	Disadvantages
PID	<ul style="list-style-type: none"> • Lowest compute complexity. • Good pose tracking performance. • Reduced sensitivity to model uncertainty. 	<ul style="list-style-type: none"> • Implementations often linearise behaviour. • Limited performance in aggressive manoeuvres.
LQR	<ul style="list-style-type: none"> • Algorithmic design process. • Jointly optimises tracking error and control effort. 	<ul style="list-style-type: none"> • Linearised system model. • Sensitive to model uncertainty and large disturbances. • Can exhibit steady state error.
SMC	<ul style="list-style-type: none"> • Non-linear system control. • Good compromise between tracking performance and robustness. • Simple robust structure. 	<ul style="list-style-type: none"> • Chattering behaviour around setpoint. • Can excite airframe structural resonance. • Also entails excessive energy consumption.
BSC	<ul style="list-style-type: none"> • Non-linear system control. • Well adapted to cascading structures. • Highly responsive. 	<ul style="list-style-type: none"> • Large control signals and actuator saturation. • Can exhibit steady state error. • Loses performance with disturbances and system uncertainty.
MPC	<ul style="list-style-type: none"> • Linear and non-linear implementations. • Explicit handling of actuator dynamics and limits. • Best in class performance with state prediction. 	<ul style="list-style-type: none"> • Performance dependent on model accuracy. • Exceptionally high compute requirements. • Feasible control solution not guaranteed.

While specific implementations of each architecture may prove advantageous in certain scenarios, given the typical sensitivity of NDE sensors to changes in their pose during data capture, an airborne inspection control system should generally be designed to prioritise effective position regulation and disturbance rejection over aggressive dynamic trajectory tracking. This can make the chattering associated with SMC problematic, while the steady state error shown in LQR or BSC control can

negatively influence precise sensor placement. Based on Table 2-1, a strong case is then made for a PID or MPC strategy, with PID preferred where onboard computing power is limited.

The final stage of the control structure then concerns allocation of the desired wrench output to individual propeller motor speeds. This can entail a number of strategies based on the actuator effectiveness matrix, \mathbf{A}_p (2-23), as utilised throughout early works on quadrotor modelling [199]–[202].

Most commonly, the Moore-Penrose pseudo-inverse is used to form a constant mixing matrix. Calculated using Singular Value Decomposition (SVD), this may provide an inverse of the non-square \mathbf{A}_p matrices common to multirotors. Applied to a target wrench vector, this grants a least-squares solution which correlates well with the combination of individual propeller thrust requiring minimal total power [218]. Other options based on real-time optimisation algorithms improve upon this by explicit handling of per-motor thrust limits and other constraints not representable within the actuation matrix to better meet the desired wrench at the expense of significant additional computing power [219].

The overall multirotor UAV control structure of Figure 2-12 is thereby formed of a number of constituent parts, with multiple approaches to the core pose control algorithm established and continuing to evolve in modern times. Key principles of each component have been recognised so as to inform developments in the remainder of this thesis.

2.4.5 Evaluation

With the research objectives of this thesis, the effective dynamic system modelling and control of multirotor platforms is identified as a topic of great significance for the successful deployment of inspection modalities beyond established, manually piloted, visual methods. From the review of multirotor UAV platforms and their underlying principles conducted above, it is clear that the fundamental system is well understood in modern research. Actuator models drawing on mature aerodynamic theory established decades prior in the context of helicopters [176]–[178] have been widely

employed in both simulation [205] and control research [214]–[216]. Primary effects of thrust generation and drag torque are ubiquitous across these works, whereas secondary effects such as roll torque and drag force are discretionally included based on the dynamic performance requirements for the application. While essential for high performance trajectory tracking and accurate simulation [108], [181], [220], they are often omitted for slow moving applications [183]–[185], more relevant to NDE deployment. The means by which these propellers can collectively effect rotation and translation in an underactuated UDT multirotor craft are likewise well documented in literature [108], [183]. These principals now serve as the basis for ongoing research into novel UAV configurations, as will be examined in Section 2.5, below.

From the findings of Section 2.4.4, it is also apparent that the field of multirotor UAV flight control represents a developed body of knowledge that may be gainfully employed in pursuit of NDE sensor delivery objectives. Literature presents numerous successful applications of established control structures to multirotor platforms including PID, LQR, SMC, BSC, and MPC approaches. Thus, multirotor flight control may be regarded as mature to an application level while ongoing research seeks further performance enhancements.

Finally, it is worth noting the availability and capability of numerous open-source flight control platforms such as ArduPilot [221], Px4 [222] and Betaflight [223], with a detailed comparison provided by Ebeid *et al* [205]. Fundamentally based on PID algorithms, these provide an extensive codebase incorporating additional features such as multi-sensor data fusion for robust pose estimations, safety features like geofencing and return to home on signal loss, and ready integration with onboard computers via protocols like MAVLink [224] and ROS [225]. Extensive testing and technical support for each of these platforms is available through the UAV community, providing great benefit to practical research and development, with the ecosystem around Px4 in particular supporting multiple academic works and commercial products [222]. Modification of these control platforms is therefore well suited to serve as the basis for further development in this thesis.

2.5 Airborne Non-Destructive Evaluation

Following the establishment of the operating principles of visual inspection and 3D reconstruction in Section 2.2, ultrasonics in Section 2.3, and conventional multirotor UAVs in Section 2.4, their combined utilisation for airborne inspection is now considered through contemporary academic literature and publicised industrial developments.

This review is further subdivided into a number of categories according to the nature of the UAV employed. Consideration is given first to multirotor UAVs performing NDE from free flight, then to those which may interact with a target structure while fully supporting their own weight, and lastly to those vehicles engaging in deeper contact where the target may partially or fully support their weight. Examples of effective implementation of a quantitative airborne NDE are examined in each case to identify trends and inform development objectives.

2.5.1 Non-Contact Systems

As introduced in Section 2.1, utilisation of UAV agents for remote visual inspection has become commonplace. Commercial entities offering UAV related services within the UK now number over 500, ranging from a majority of small sole-trader enterprises, to well-established businesses servicing large industrial clients managing thousands of assets, including international energy providers [226]. Technologies supporting visual image data capture in free flight can therefore be deemed functionally mature. Accordingly, current development efforts regarding visual and other non-contact airborne NDE among leading service providers and research organisation investigate refinements to improve process efficiency and expand the level of structural health insight offered.

To this end, amid the themes of traceability and digitisation expressed by the general drive for Industry 4.0 [79], leading UAV inspection providers, such as Cyberhawk [227], Skyspecs [228], and Field Group [229], have moved to provide online reporting. As such, they utilise proprietary software to collate, analyse, and document UAV inspection images via a cloud server, permitting images captured in situ to be

immediately uploaded for analysis by experts in an off-site centralised location. In addition to the operational efficiency improvements due to reduced travel of key personnel, such software suites provide the ability to add further metadata to conventional inspection reports. Features such as the visual linking of images to their location within a 360° photo or 3D reconstruction, as depicted in Figure 2-13, provide intuitive navigation that helps inspectors and operators assess defect severity from the context of its surroundings and reduce incidents of confusion due to structural self-similarity. Multiple images of the same defect may also be automatically tagged and grouped to provide redundant views of the same region, offering robustness against artefacts due to viewpoint, transient occlusions, or suboptimal camera configuration and historical tracking across multiple inspections [227]–[229]. Such software is expected to increase clarity of reporting and actionable insight, informing more cost-effective data driven operations and maintenance strategies.

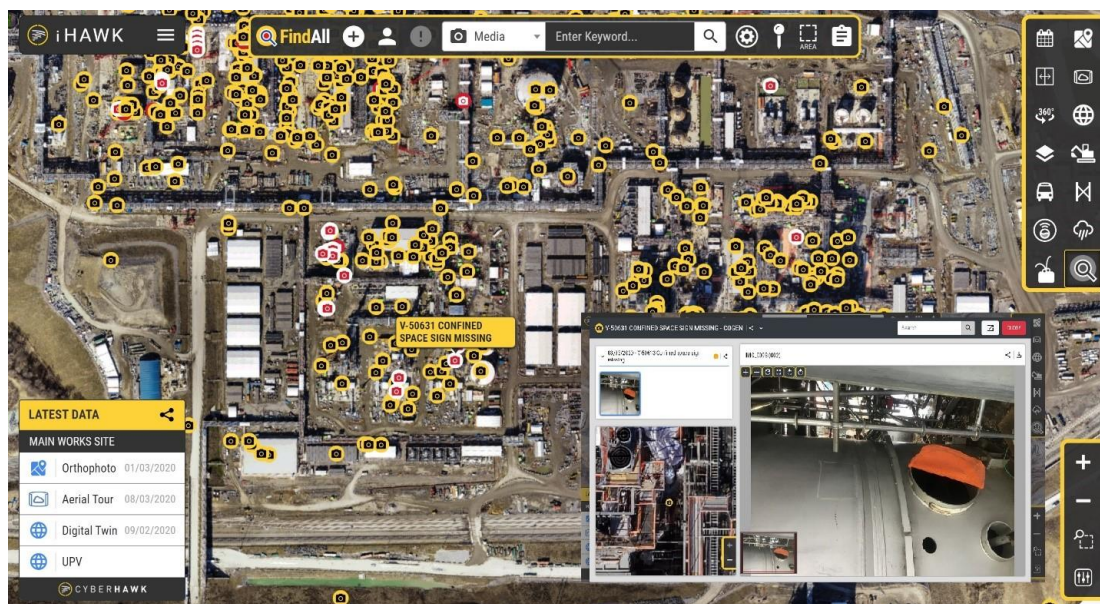


Figure 2-13: A screenshot from the Cyberhawk iHawk inspection reporting and analysis software [227]. Key points are marked on the orthoimage map, highlighting points of interest alongside their current status and date of last inspection. The option to access deeper inspection data is also presented.

However, analysis of thousands of highly similar inspection images is still required within this digitised workflow. This remains a bottleneck owing to the limitations of human inspectors, as discussed in Section 2.2. Use of machine vision and artificial

intelligence to ease the workload has therefore become a prominent a research theme concerning UAV visual inspection.

Numerous algorithms and approaches have been developed to this end, as are reviewed in depth by Czimmermann *et al* [230]. Recent advances with thematic relevance have shown ability to autonomously identify visible cracks in wind turbine blades, as exemplified by Wang *et al* [82]. Therein cracks are successfully detected and segmented down to length of 6.54 cm and width of 1.55 cm at a 95 % true-positive rate, as shown in Figure 2-14(a). Modern neural-net approaches show further advancement, as in comparative assessment versus classical edge-detection methods by Dorafshan *et al* [231], depicted in Figure 2-14(b). Applied to concrete structures, as are common to infrastructure in the energy and civil sectors, neural net results exceeded 97 % detection accuracy, detecting cracks down to 0.04 mm width in high resolution images, and requiring 50 % less computation than classical edge detection methods. However, the safety-critical nature of NDE prevents full reliance on unsupervised defect identification. In practice, current algorithms thus only offer a limited time saving, acting to eliminate obvious defect free image regions by highlighting salient features potentially indicative of defects for final review by a human inspector. Their benefit is thus found in allowing the inspector to focus their attention on the borderline cases with more subtle underlying damage.




Figure 2-14: Results of automated defect detection in UAV inspection images. (a) Wang et al [82] detect cracks in a wind turbine blade, as indicated by blue boxes. Segmentation on the pixel level is shown by the yellow polygons in the inset zoomed images. (b) Dorafshan et al [231] segment hairline cracks in concrete using a neural net, indicated in red.

Outside of these current efforts to improve RVI process efficiency, UAVs also see widespread utilisation as data capture agents within the photogrammetric reconstruction process, described in detail in Appendix A. Adding quantitative insight to asset geometry beyond what is available in pure remote viewing modes, numerous research works have, therefore, evaluated this inspection process for health monitoring of large-scale structures.

In surveying a timber frame bridge, Khaloo *et al* [232] demonstrate that UAV photogrammetry may generate high density point clouds suitable for photorealistic representation and inspection, as shown in Figure 2-15. Moreover, in comparison to terrestrial laser scanning, these showed consistently increased sample density and coverage, with reduced data capture time and no requirement for work at height. Spatial noise was approximately 2.5 times higher than was estimated in the laser reconstruction but, at a mean deviation of 4.8 mm, the data remained more than sufficient for purpose.




Figure 2-15: UAV-based photogrammetric inspection of a timber frame bridge by Khaloo *et al* [232]. (a) Photograph of a timber joint. (b) High density point cloud reconstruction of the same region.

Similarly, Zhao *et al* [233] examine UAV photogrammetry for routine monitoring and emergency assessment of dams, characterising the influence of ground control points on reconstruction. They find that photogrammetry is able to retain strong relative spatial consistency in regions not constrained by absolute reference dimensions, observing a median absolute deviation of under 2.9 cm between a point cloud with full

reference coverage and one with 44 % fewer control points grouped to omit the lower half of the dam entirely. Additionally, therein, photogrammetry showed error under 1.6 cm versus the physical distances between sufficiently dense control points measured using GNSS. Damage to the concrete structure such as spalling, cracking and precipitate formation were readily detected and sized within the reconstruction, highlighting the value of the UAV survey methodology.

Factors influencing photogrammetric accuracy are further investigated by Zhang *et al* [234] for reconstructive NDE of wind turbine blades, a particular challenge owing to their expansive, low-textured, uniform colour surfaces. In comparison with a metrology grade scanning system, analysis of the reconstruction from UAV images highlights potential degradation by poor environmental lighting, under-exposing the subject and masking visual features. Compensation with electronic gain can add false features due to sensor noise, whereas increasing exposure time then limits fidelity by increasing vulnerability to UAV motion during image capture, as shown in Figure 2-16. This indicates existence of a reduced maximum image quality in the absence of an additional light source. Camera standoff from the surface is a further critical influence, with closer proximity enhancing resolution and visibility of fine-scale unique features essential for photogrammetry and NDE but also increasing likelihood of a dangerous collision due to near-surface aerodynamic effects. Moreover, when using a fixed focal length camera, deviation from the desired standoff distance negatively influences reconstruction. The target surface is displaced from the focal point and further blurring introduced, with particularly high standoff sensitivity caused by the small depth of field accompanying large lens apertures used to counteract motion blur. Understanding and practical mitigation of these highly interrelated effects arising from standard camera parameters and the UAV motion profile are thus shown to be particularly valuable to continued development of quantitative visual processes for airborne NDE [234].




Figure 2-16: Investigation into factors affecting UAV photogrammetric reconstruction quality of a wind turbine blade by Zhang *et al* [234]. A colour-mapped mesh shows deviation from the ground truth with: (a) baseline configuration, (b) 33 % ambient light, (c) 33 % ambient light and 2x exposure time.

Accordingly, improvements to aerial non-contact data capture represent another significant theme in ongoing research. For visual NDE, the objective is to maintain consistent surface-relative standoff while ensuring complete coverage and sufficient image overlap: aspects critical to consistent high-fidelity assessment across the target [234]. However, incomplete knowledge of the exact position and geometry of complex industrial structures often prevents offline generation of a flightpath fully satisfying such requirements. A range of adaptive methods have thus been developed to overcome this issue with varying levels of autonomy.

The most manual inspection approach requires significant skill from the pilot, splitting their attention between avoiding surface collisions while also ensuring complete visual coverage of the target exterior. Assistive systems, such as that developed by Omari *et al* [235], help reduce this cognitive load, holding UAV pose in the absence of pilot input and locking its path to a consistent surface standoff using real-time geometry estimation by a stereo camera pair, as in Figure 2-17. Control in the application critical flight dimension is thus automated, simplifying the piloting problem while retaining the full flexibility of a manual flight in the other axes.




Figure 2-17: A pilot assistance algorithm for photogrammetric UAV inspection by Omari *et al* [235]. (a) The flight path, in red, taken with constant standoff relative to the building façade, represented as an occupancy map. (b) A photograph of the building façade. (c) The post-processed façade reconstruction.

Another strategy is to introduce adaptive correction to a prescribed generic path formed using approximately known asset geometry, as used by Zhang *et al* [234] to maintain constant standoff while encircling the wind turbine blade via feedback from a planar lidar. Beneficially, this removes requirements for the pilot to track surface coverage and improves consistency across repeat inspections. Quenzel *et al* [236] demonstrate a similar strategy for photogrammetric inspection inside a decommissioned industrial chimney, as in Figure 2-18. They start with a conical helix path planned from its coarse dimensions then apply ad hoc standoff corrections based on feedback from a combined lidar and visual inertial stereo Simultaneous Localisation And Mapping (SLAM) system. They also expand assessment practicality, using a preliminary coarse reconstruction of the chimney interior and unrolling the surface texture for on-site expert review. Defects or visual anomalies are then targeted for immediate follow up imaging ahead of final full-scale reconstruction, improving inspection quality without costly return visits to the site.

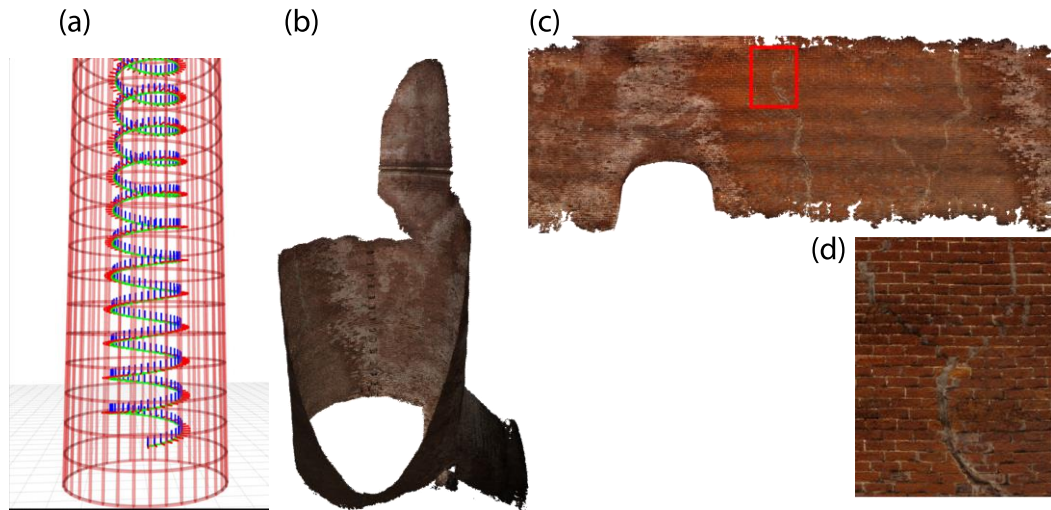


Figure 2-18: Automated photogrammetric inspection of an industrial chimney by Quenzel *et al* [236]. (a) Flightpath planned with coarse geometry knowledge. (b) Textured mesh reconstruction of the chimney interior. (c) Unrolled surface texture. (d) Inset of region in red box showing crack formation.

Overall, the highest level of non-contact inspection autonomy is shown by systems capable of independently exploring 3D surfaces within bounded regions of fully unknown environments, such as that by Bircher *et al* [237]. Utilising a receding horizon sampling of possible future waypoints in a random tree alongside visual inertial stereo mapping, they demonstrate real-time adaptive path-planning. The UAV can explore complex environments, such as around bridge trusses and scaffolding, without human intervention and then be directed to inspect a target in detail, autonomously generating a full coverage flight path. In this case, the operator is minimally involved with piloting and may fully focus on the inspection data, identifying and categorising defects as observed in real-time. NDE process efficiency is thereby improved, with reduced requirements for specialised UAV skillsets and a compressed end-to-end timespan.

With thorough validation of its core reconstruction ability to NDE standards well established and numerous developments readily supporting practical deployment, UAV photogrammetry is now seeing frequent adoption as a spatial reference for other NDE data. Reconstruction then expands their context and aids navigation, as with the commercial reporting software [227]–[229]. A prime example of this trend is given by Biscarini *et al* [111] who fuse infrared imagery and Ground Penetrating Radar (GPR)

with a UAV photogrammetric reconstruction of the Ponte Lucano, a roman bridge. Doing so: the photogrammetric mesh provides a photorealistic record of inaccessible and previously undocumented bridge geometry, thermographic imaging helps identify areas of water retaining minerals and biofouling, and GPR is used to detect and assess internal buried structures. Together the combined data produce a robust documentation of the current bridge status and ongoing degradation mechanisms, highly valuable for historical preservation efforts.

A further example of meaningful data fusion using UAV photogrammetry is given by Connor *et al* [238]. In this case a UAV is outfitted with both a camera and a scintillator to detect and quantify radiological contamination. Flying over a storage site for topsoil removed from the vicinity of the Fukushima nuclear plant, this system produced a geo-referenced mesh reconstruction that allowed operators to quickly identify where contaminated rainwater was pooling in the surrounding environment. Prompt remediation of the plume to the right of Figure 2-19(a) was then possible. Status tracking through repeat assessments, as in Figure 2-19(b), also yielded valuable insight regarding the overall reduction in contamination without requiring manual access to precarious locations atop the store, saving time and operator radiation exposure.




Figure 2-19: Fusion of radiation dosimetry survey data with photogrammetric mesh of a contaminated topsoil store near the site of the Fukushima power plant by Connor *et al* [238]. (a) Initial survey: October 2016. (b) Follow-up survey: Jun 2017, decay corrected to match initial survey. Note that the contamination plume in the top left of the follow-up survey is attributed to edge-effects of un-remediated ground outside the site.

Other non-contact airborne sensing approaches outside of the visible spectrum have also experienced significant recent development. Pertinent examples include the performance of non-contact acoustic inspection by Sugimoto *et al* [239]. In this system, a mock section of a concrete building façade is irradiated with acoustic

pressure waves, having frequency from 0.3 kHz to 4 kHz and transmitted from a large UAV-mounted planar source. Operating similarly to a “tap test”, differences in the localised surface vibration response are detected from the ground by a laser doppler vibrometer. These then correspond to areas detached from the substrate or cavities in the façade, which may be accurately detected and sized. This represents a novel means of screening large exterior surface areas of civil structures but its reliance on resonance behaviour may pose issues regarding penetration depth in thicker structures.

Limited non-contact volumetric sensing from a UAV may also be accomplished via Ground Penetrating Radar (GPR). This sees existing application in monitoring of reinforced concrete structures, often with manual data collection as by Dinh *et al* [240], but has also been adapted for aerial deployment. Massaro *et al* [241], mount a GPR unit beneath a UAV to check bridges for subsurface cavities and water ingress that may cause significant structural damage if left uncorrected. Similar technology has also been examined for hydrology survey, as by Valence *et al* [242], and detection of unexploded buried ordnance, as by Bähneemann *et al* [243], leveraging the mobility of a UAV platform for significant utility beyond the scope of NDE.

2.5.2 Self-Supported Contact Systems

Despite the many benefits offered by contact-based NDE, the mature state of multirotor flight control, and the high-mobility aerial viewpoint UAVs offer, systems combining these aspects are not yet widely deployed in standard industrial practice.

This may be attributed to fundamental limitations in the positioning accuracy, precision, and stability of environmental interactions performed by a conventional under-actuated quadcopter. As detailed in Section 2.4.2, a UDT multirotor craft requires reorientation of its body to oppose disturbances and exert contact forces. This limits dynamic responsiveness, and may displace the probe relative to the surface, negatively impacting NDE sensor coupling and measurement quality.

To bypass this issue, the small number of commercial systems attempting airborne NDE using UDT multirotors typically rely on large outriggers or magnetic adhesion, as depicted in Figure 2-20.




Figure 2-20: Commercial UAV ultrasonic inspection platforms. (a) Texo-DSI [244]. (b) Tritex Multigauge 6500 [245]. (c) Air Control Entech [246]. (d) Tera Inspectioneering [247].

Such approaches have some merit. Using a platform with eight rotors in the standard quadcopter layout, Texo-DSI were able to demonstrate the first successful commercial ultrasonic thickness inspection from an airborne robot in September 2017 [248]. While the technology was in its infancy, with significant piloting skill required to perform the interaction and outriggers later added to improve stability, their demonstration prompted further technology development by competing stakeholders. This included the UAV of Figure 2-20(b), created by Tritex NDT [245] to fully integrate a custom-designed ultrasound probe, couplant gel dispenser, and instrumentation unit with a high-payload craft. Others tackled the stability problem, with the Air Control Entech platform [246] using an array of electromagnets to affix itself to the metallic target structure before precisely deploying ultrasound NDE with a lightweight onboard manipulator arm shown in Figure 2-20(c). Similarly, Terra Inspectioneering [247] developed a platform including an extended guard frame that is used with the end effector to form a tripod against the target surface, as shown in Figure 2-20(d). Under continuous static thrust, this enables passive rejection of aerodynamic disturbances while small unpowered wheels also permit limited rolling motion for in-contact repositioning. It should be noted, however, that the mechanical structures added to

these vehicles entail significant additional mass, yielding large UAVs with restricted structural access in the face of further degraded under-actuated manoeuvrability and amplified snag-risk when operating within crowded industrial airspace.

Strategies for bypassing the problems of under-actuated environmental interaction have also seen substantial examination in recent academic literature. Accordingly, systems with proven ability to perform meaningful NDE are highlighted in Figure 2-21.

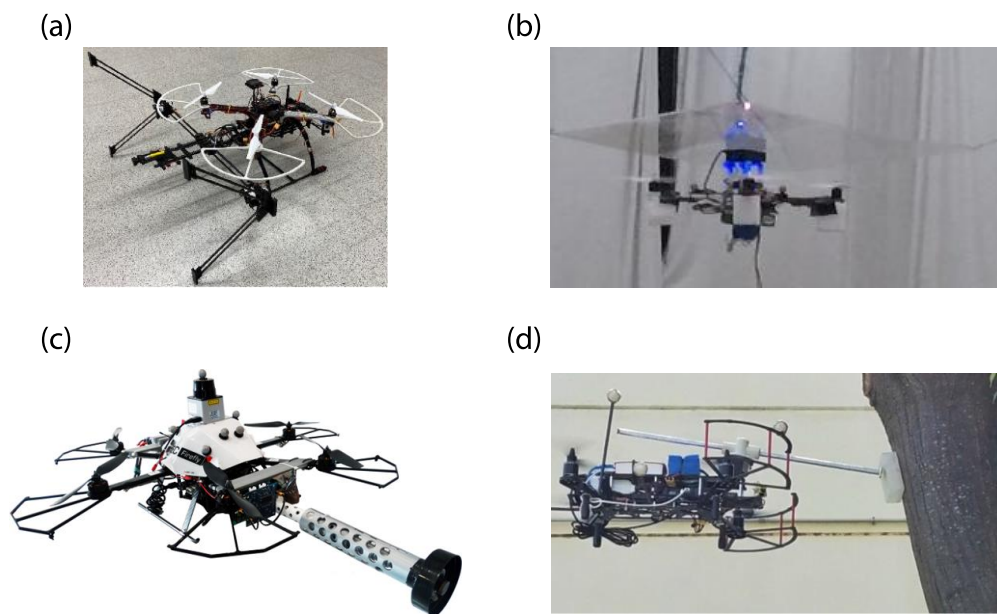


Figure 2-21: Under-actuated multirotors for contact inspection. (a) The quadrotor by Gonzalez de Santos *et al* [249]. (b) The quadcopter by Kocer *et al* [250]. (c) The hexacopter by Zhang *et al* [251]. (d) A UAV by Hamaza *et al* [252] for sensor installation. Note that while (a) (b) and (c) all perform ultrasonic thickness measurement using embedded probes, (d) may be designed to install a discrete sensor package granting similar capability.

Among such developments, ultrasonic measurement of a vertical metal plate is conducted by González de Santos *et al* [253] using a pair of laser range sensors to govern approach angle and ensure smooth transition into contact, with support and stabilisation from a large, mechanically damped outrigger. Updates to this system [249], shown in Figure 2-21(a), apply horizontal thrust into the plate via two additional propellers on the rear of the vehicle, but do not grant full-actuation. In Figure 2-21(b), Kocer *et al* [250] demonstrate a proof of concept controller for ultrasonic thickness measurement of the underside of a horizontal acrylic plate using a standard quadcopter

governed by a nonlinear MPC developed to estimate and control interaction forces in real time. Zhang *et al* [251] also conduct ultrasonic measurement using the UAV shown in Figure 2-21(c) via short-duration contact with a vertical plate under autonomous control with standoff distance and yaw angle feedback from a planar lidar sensor.

More novel airborne NDE strategies involve using a 1 DoF manipulator to perform a “tap-test” for voids in concrete bridge material [254], or to place and retrieve an embedded structural health monitoring sensor using an active single DoF manipulator that may impart impulse forces or dynamically alter its compliance as detailed by Hamaza *et al* [252] and shown in Figure 2-21(d). Such advanced control strategies thereby aid piloting and sensor deployment repeatability versus the more manual and mechanical approaches taken by the systems of Figure 2-20. However, all of these multirotor systems remain dynamically constrained by the negative aspects of under-actuation.

Accordingly, to improve the efficacy of UAV environmental contact, ongoing research now addresses the dynamics of the underlying multirotor platform. By including additional actuators, the system may generate both 3D force and 3D torque, and so directly influence its state across the six dimensions of translational and rotational freedom in which it operates. The UAV then becomes fully or over-actuated and gains the ability to regulate contact without complex whole-body manoeuvres.

There are two popular strategies to achieve this effect, as depicted in Figure 2-22. In a Multidirectional Thrust (MDT) system, a number of propellers are arranged about the airframe with fixed orientation spanning 3D thrust space. The speed of each propeller is then controlled to generate the desired wrench by thrust superposition at the cost of efficiency [255]. The resulting system is highly dynamic, with responsiveness limited only by the propeller motor throttle response. Alternately, in a Vectored Thrust (VT) system, propellers are actively re-oriented to directly generate the desired wrench without inefficient cross-cancelation, offering higher peak force than MDT and optionally returning to a more efficient UDT layout in free flight [256]. However, in this case, dynamic responsiveness is limited by the mechanical reorientation of the propellers and so is typically slower than an MDT approach.

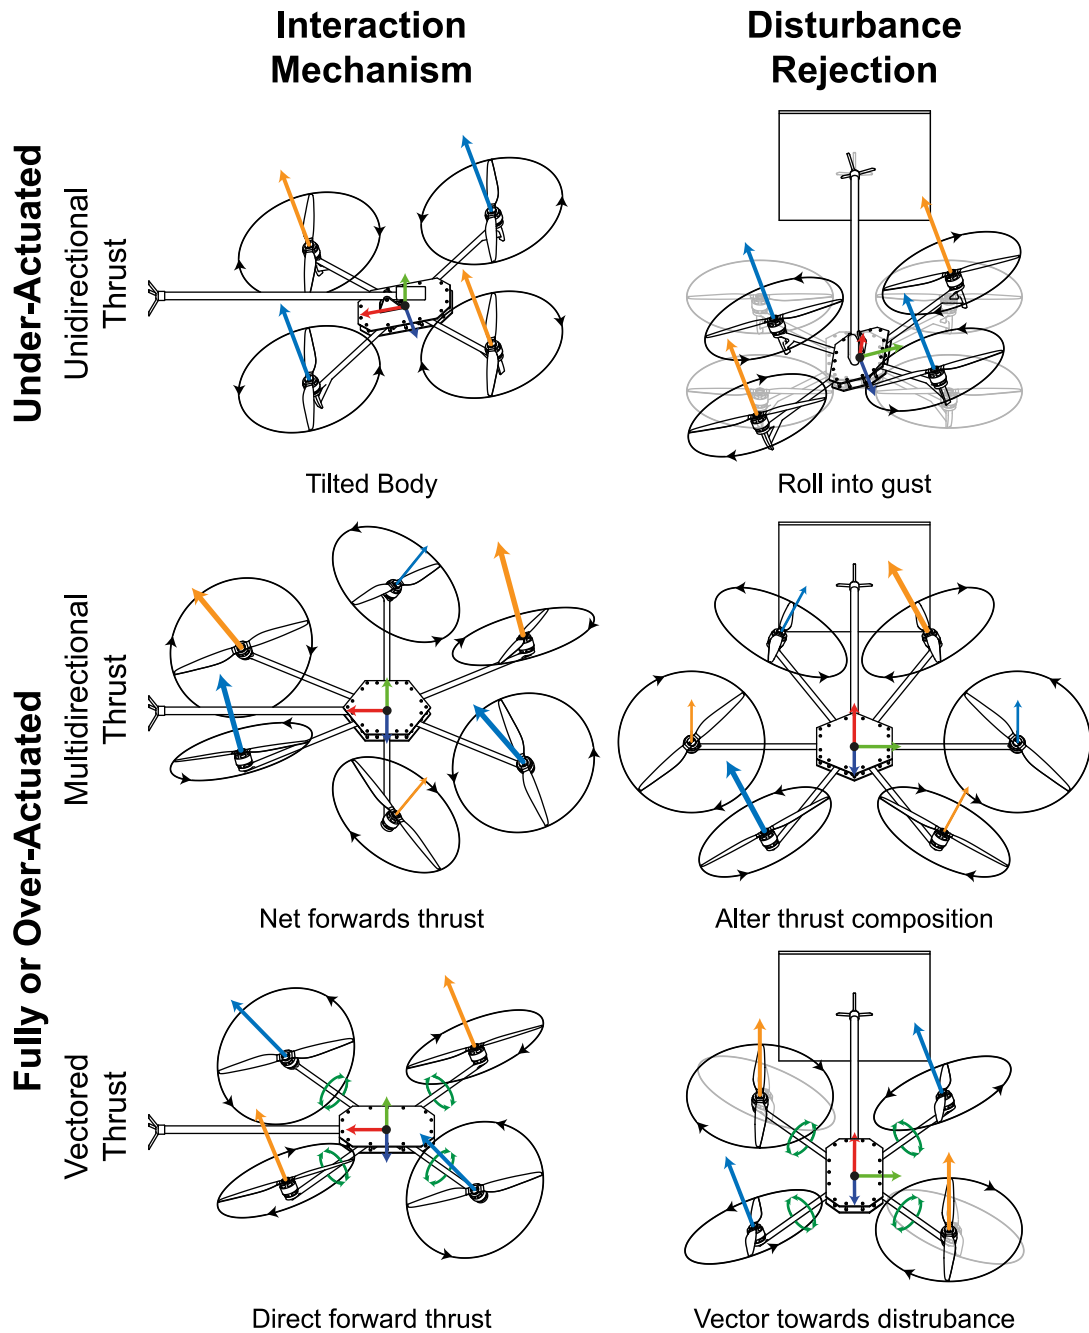


Figure 2-22: A comparison of multirotor configurations for environmental interaction. Propeller relative thrust is indicated by arrow size. The illustrated disturbance response corresponds to rejection of a gust of wind from the left of the image. Undisturbed craft position is shown in grey.

UAVs of either architecture can successfully enter repeatable stable interactions with their environment, rejecting aerodynamic effects without disturbing a surface-contacting end effector, and, thereby, function effectively as aerial manipulators, as reviewed in detail by Ollero *et al* [257]. Hence, characteristics supportive of

quantitative, contact-based, multi-modal and in situ NDE are identified and examined in a selection of recent works, with highlights depicted in Figure 2-23.




Figure 2-23: Aerial Manipulator UAV platforms. (a) The AeroX by Trujillo *et al* [258] performing ultrasonic inspection. (b) The Voliro Hexacopter by Bodie *et al* [259] performing half-cell potential measurement. (c) The OTHex by Tognon *et al* [260] performing eddy current NDT. (d) An aerial manipulator by Sanchez-Cuevas *et al* [261] in contact with a bridge underside. Of these (a), (c) and (d) use MDT to achieve full actuation, with additional degrees of freedom provided by an embedded manipulator arm. Only (b) uses VT, granting over-actuation for omnidirectional flight and force exertion with a simple rigid effector.

Most prominently, the AeroX octocopter detailed by Trujillo *et al* [258] uses multidirectional thrust superposition and a gel-coupled, dual-element, ultrasonic probe mounted at the end of a 6 DoF actuated manipulator arm to inspect oil refinery storage tanks and pipe sections, as in Figure 2-23(a). The large 1.7 m by 2.3 m UAV with a maximum take-off weight of 25 kg captured thickness measurement to the satisfaction of an NDE inspector but no quantitative claim is made to its accuracy.

Tognon *et al* [260] present an alternate configuration in the OTHex multidirectional-thrust hexacopter. As in Figure 2-23(c), this uses a 2 DoF manipulator arm to

successfully localise weld material on a horizontal pipe section with an eddy current probe. They make no comment regarding assessment of weld integrity. As with the AeroX, the onboard manipulator here enables omnidirectional surface contact while the UAV retains a horizontal attitude. In Figure 2-23(d), Sanchez-Cuevas *et al* [261] use a similar fully-actuated hexacopter to access the underside of a concrete bridge, maintaining contact using an outrigger frame atop the UAV, but again present no evaluation of NDE data.

Building on previous work [185], using a hexacopter layout with 12 propellers Bodie *et al* [259], [262] generate 6 DoF pose-decoupled force and torque that may be omnidirectionally vectored relative to the craft body. Whereas the other platforms may also generate 3D force, here the VT architecture allows the UAV to exert force greater than the craft mass in all directions about its body and thereby fly in non-horizontal orientations. They inspect a reinforced concrete block at 5 cm intervals using a half-cell potential mapping circuit grounded to the sample, as in Figure 2-23(b).

Whereas these aerial manipulator systems represent an excellent capability to deploy NDE at height and in scenarios of restricted manual access, it is noted that the long-reach hovering contact by rigid arm [259], [263] or serial manipulator [258], [260], requires a supportive frame [261], or necessitates strong interaction feedback and control [258]–[260], [262], particularly during any dynamic scanning motions. Moreover, the resulting large, complex, UAVs require a 1-2 m target clearance radius which could pose problems in confined industrial applications.

2.5.3 Externally Supported Contact Systems

Where the geometry of the asset under inspection is non-planar, highly textured, or supportive of other adhesion mechanisms, opportunities exist to enter much more stable contact at closer proximity than demonstrated by the systems above. In such cases, the structure itself can partially support a robotic vehicle through multiple points of interaction via mechanisms including friction, grasping, magnetism, and vacuum-based adhesion, as reviewed by Schmidt and Berns [264]. Potential improvements may then be expected to platform disturbance rejection and station-holding energy

requirements, allowing longer deployment duration for contact-based sensing and thereby improving structural health insight offered by airborne NDE.

Whereas the aerial manipulator craft discussed in Section 2.5.2 may provide highly-flexible generic inspection access without supportive structural features, exploration of adhesion mechanisms may present alternate benefits to airborne NDE processes where such platforms cannot operate, for example due to limited available access space around elevated pipework. Therefore, pursuant to the research objectives stated in Chapter 1, contemporary systems capable of contact-based inspection around pipes, cylindrical tanks, and pressure vessels in confined spaces are examined.




Figure 2-24: Crawler robots used for commercial NDE. (a) The Eddyfi Scorpion 2 [265] deploying ultrasonic thickness mapping. (b) The Mistras ART crawler [266] deploying tangential X-ray imaging.

Prime example of such systems is given by crawler robots like the commercial Scorpion 2 [265] and Mistras ART [266] platforms depicted in Figure 2-24. By utilising magnetic tyres to adhere to the vertical walls of petrochemical tanks or other large industrial structures, the Scorpion 2 provides excellent stability and motion uniformity to deploy ultrasonic area thickness mapping without manual rope-access. As power and data transfer are conducted by a long umbilical cable for unlimited operational duration, care must be taken to avoid its snagging. Designed for operation around insulated pipework, and so denied magnetic adhesion, the Mistras ART crawler instead conducts inspection from a position resting atop the pipe. Unlike in-pipe NDE robots, as reviewed in [67], it therefore does not require asset shutdown and draining before use, minimising lost production during inspection. Stability of the ART crawler is passively governed by its weight and support from the pipe below, aided by position correction using its motorised tracks. Enclosing the pipe diameter, this crawler then

performs tangential x-ray inspection of the pipe beneath outer insulation. Its applications, however, are restricted by the ability of the sensor to enclose larger diameters and the support for its weight offered by small-bore pipework.

A smaller form factor is presented by other pipe-external robotic systems, such as the C-clamp style crawlers with minimal radial clearance requirements developed by Chatzakos *et al* [267] and Choi *et al* [268], as in Figure 2-25. These almost fully enclose the target circumference for high-stability sensor positioning via rigid mechanical grasping. However, they cannot autonomously disengage and re-enter contact to pass circumferential obstacles like flanges, supports, or sharp bends, posing challenges to practical deployment. Khan *et al* [269] propose a means to step-over these obstacles using a walking robot with two magnetic feet, but such a strategy prevents application to insulated, composite, or otherwise non-magnetic structures.




Figure 2-25: Pipe exterior crawlers from recent literature. (a) Split ring crawler for nuclear feeder pipe inspection by Choi *et al* [268]. (b) C-clamping omnidirectional pipe crawler by Chatzakos *et al* [267]. (c) Bipedal walking crawler with magnetic feet by Khan *et al* [269].

By contrast, a UAV may bypass all obstructions along the asset surface, flying directly between the points of deployment and inspection. Small form factor craft engaging in strong environmental interactions then present the opportunity for stable, short-reach, sensor deployment while operating in confined spaces.

In work by David & Zarrouk [270] shown in Figure 2-26(a), a wall-crawling UAV is supported via friction from surface-normal thrust created by reversing its propellers and an articulated airframe that provides limited grasping outside of flight. However,

transition from flight to wall-crawling is not shown and high energy usage precludes applications with low-friction surfaces or large vehicles. A mix of friction and vectored-thrust weight support may relax these issues, as shown by Myeong & Myung [271] in Figure 2-26(b), but will still require power similar to that of free flight. In this case, thrust is also vectored in only 1 DoF, preventing full-actuation and limiting support of the vehicle when it is not oriented to point up the wall. Moreover, neither work [270], [271] examines adhesion to curved or non-planar surfaces, to the underside of overhanging surfaces, or the application of NDE.




Figure 2-26: Wall-crawling UAVs using active thrust-based support. (a) The quadcopter by David & Zarrouk [270] also uses the propeller motors to drive its wheels. (b) Another quadcopter by Myeong & Myung [271] uses vectored thrust support and passively rotating compliant wheels to drive up a wall.

Alternatively, passive perching strategies, as reviewed by Meng *et al* [272], may allow a vehicle to maintain surface-contacting position while minimizing its energy usage. Examples include deposition of a detachable Electro Magnetic Acoustic Transducer (EMAT) sensor package atop a horizontal steel sheet and a 6-inch diameter pipe section by a standard quadcopter, as reported by Jarvis *et al* [273] and shown in Figure 2-27(a), utilising the magnetic adhesion of the EMAT to stabilise its position. However, in this case, successful probe deployment to the plate and pipe targets is limited to below 65 % and 60 %, respectively, by the UAV momentum and pose stability during initial contact. Others mechanically support the UAV beneath the pipe, as shown by Yu *et al* [274] in Figure 2-27(b), and Zhang *et al* [275], or by having the vehicle rest atop it, as performed by Garcia Rubiales *et al* [276], Lopez-Lora *et al* [277] and Cacace *et al* [278] in Figure 2-27(c) and (d). These strategies typically encompass the pipe with a gripper, either rigid [274], [275], [277] or conformable [276], or deploy NDE via an

onboard robotic arm while resting passively atop the pipe surface [277], [278]. Such support mechanisms significantly enhance stability and operational duration versus a free-flying UAV but in some cases cannot support vehicle movement while in contact with the surface [274], [275], restricting potential NDE applications. Others exhibiting locomotion use an inchworm mechanism [276] or driven wheels [277], [278] to dynamically investigate suspected defects. However, arm mass, reach and gripper diameter limit scalability to large diameter assets as these vehicles lack suitable adhesion mechanisms to support their weight in positions other than atop the pipe.

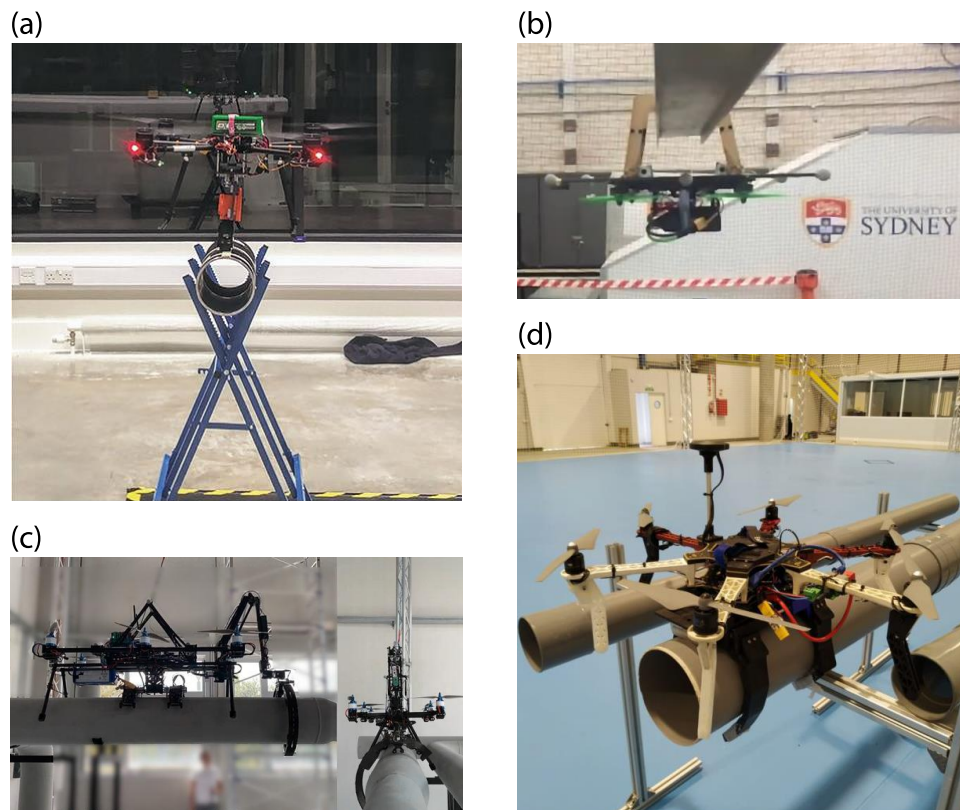


Figure 2-27: UAVs for passive perching and inspection on pipe-like structures. (a) Deposition of a discrete EMAT sensor package using a quadcopter by Jarvis *et al* [273]. (b) A quadcopter by Yu *et al* [274] flips upside down and grasps a beam. (c) A quadrotor with a specialised robotic effector by Cacace *et al* [278]. (d) A compliant gripper beneath the hexacopter used for stabilisation atop a pipe by Garcia Rubiales *et al* [276]. Note that (a), (c) and (d) are capable of ultrasonic NDE.

A more generic, paradigm is thus for the entire UAV to crawl about the asset surface. In the context of pipework, tanks, pressure vessels and other cylindrical infrastructure, preliminary works representing the current state of the art [279], [280] address control algorithms for non-flight-capable mock hardware and are depicted in Figure 2-28.

Zhao *et al* [279] use Non-Linear Model Predictive Control (NMPC) of wheel torque to stabilise a cart in the top $\pm 20^\circ$ of a pipe but omit any adhesion force necessary for circumnavigation. Nekoo *et al* [280] examine a benchtop test rig with variable-pitch propellers, regulating radial and angular position via the State-Dependent Riccati Equation (SDRE), but omitting direct consideration of beneficial contact forces. The potential benefits of inspection via aerial access and full-contact crawling have therefore not yet been fully explored.




Figure 2-28: Evaluation platforms for full-circumferential pipe access by a UAV. The mechatronic mock-up by Nekoo *et al* [280] using variable pitch propellers to control circumferential position. (b) The cart by Zhao *et al* [279] using lateral thrust and wheel holding torque to maintain position atop a pipe.

2.5.4 Other Technology Trends

Beyond these larger themes of ongoing research, multirotor UAVs have evolved to incorporate a number of supporting technologies with distinct benefit to commercial airborne inspection applications. Key themes identified in technologies relating to the design and operation of UAV systems employing conventional UDT layout are therefore briefly examined below.

2.5.4.1 Collision Tolerant Platforms

Airborne inspection within a facility often entails navigation of crowded airspace with numerous obstacles preventing direct access to the asset of interest. Consequently,

UAV platforms have been developed to incorporate protection against mechanical damage to flight critical hardware caused by unintended environmental contact.

While approaches can be as simple as installing rotor guards, these fail to protect the UAV against mechanical impingements from directions outside the primary rotor plane. For operation amid complex industrial structures, full body protection as depicted in Figure 2-29, is most advantageous.




Figure 2-29: Collision tolerant UAV platforms for confined space inspection. (a) The Flyability Elios 3 [281]. (b) The RMF-OWL by De Petris *et al* [282]. (c) The Cleo Robotics Dronut X1 [283].

A highly prominent example is the Elios 3 platform, developed commercially by Flyability [281] and widely recognised as the system responsible for establishing the interest of energy asset operators in remote confined space inspection using UAV agents. Its enclosure by a spheroid carbon fibre spaceframe, capable of limited rotation independent of the body, permits it to fly into light, unregulated, contact with the target asset to stabilise close-proximity visual inspection using an onboard camera and lighting array. This has previously enabled inspection of an underground tank room within a nuclear facility [92] and within the lower levels of a wind turbine monopile [284] without direct human presence. Further, variants outfitted with lidar and radiation sensors are now deployable for contaminant mapping within nuclear sites, providing a highly novel commercial capability to improve inspection safety and reduce costs [281].

A similar rotor protection strategy is shown recently by De Petris *et al* [282] in their RMF-OWL platform. In this case the quadrotor UAV is enclosed by a carbon-foam composite frame rigidly affixed to its body and capable of sustaining collision forces up to 128 N and 520 N in its weakest and strongest axes, respectively. The RMF-OWL is designed to support autonomous exploration and mapping of mining tunnels using

its onboard lidar and machine vision camera, operating collaboratively with other robots as a scouting agent to help secure victory for its team in the DARPA Subterranean Challenge [285]. Industrial utility of confined environments operations is thereby highlighted.

The final example of collision tolerant UAV design in Figure 2-29 is given by the Cleo Robotics Dronut X1 [283]. This comparatively unique UAV fully encloses its pair of coaxial rotors within its body for maximum environmental protection while actively redirecting their thrust to translate horizontally. Unlike the other craft discussed, the Dronut X1 thus prevents small protruding structures from entering the rotors by slipping between the structural members of a spaceframe. It then offers safest-in-class operation in confined spaces and in direct collaboration with human operators, a capability highly beneficial to any contact based airborne inspection system.

2.5.4.2 Operation Beyond Visual Line of Sight

Within the objective of removing human experts from dull, dirty, or dangerous inspection scenarios, a significant desire for UAV operation Beyond Visual Line of Sight (BVLOS) has emerged. This promises to enable fully remote inspection by experts located offsite, piloting the UAV from a centralised location to minimise the required human presence, associated hazard exposure, and travel costs. They also see application in point-to-point light cargo transport where conventional infrastructure is less cost-effective, such as rural postal delivery [286].

While BVLOS is similar in principal to modern First-Person View (FPV) platforms, allowing video to be streamed to a monitor or headset worn by the pilot, such equipment is deemed to impede pilot situational awareness and, under current UK laws [287], [288], FPV aircraft must be operated with a spotter assisting an on-site pilot. True BVLOS systems are designed to operate without any direct human presence and so must incorporate aspects of environmental mapping and autonomous obstacle avoidance, also demonstrating capability to safely and deterministically manage loss or interruption of operator control and interact with other air traffic.

To date only a very small number of UK organisations have met legal requirements to operate BVLOS inspection systems, with sees.ai [289] being the first granted

permission for routine operations at select sites in April 2020. This may appear unusual given the advanced state of UAV control, environmental mapping, and obstacle avoidance technologies but the limiting factor lies outside of these aspects.

The challenge is that, to operate safely, UAVs must be able to detect and avoid other aircraft as a pilot would when flying within visual sight. Doing so remotely, however, requires mutually compatible Unmanned Traffic Management (UTM) systems, continuously reporting UAV position, speed, expected operating bounds and pilot details, while also integrating with the flight controller to act on the same data received from nearby aircraft. As yet, however, specification for such systems has not been standardised or supported by centralised infrastructure. Without UTM, BVLOS operations are typically segregated to Temporary Danger Areas (TDAs), which must be manually registered with the Civil Aviation Authority (CAA), last for 90 days, are non-trivial to re-establish once expired, and restrict any access by others wishing to use the airspace for the duration [288], [290].

This approach is obviously problematic to operate the scale required by regular UAV inspections. UAVs capable of BVLOS operation are therefore identified as a key trend to monitor as national policy and UTM infrastructure develop to remove these barriers to widespread adoption.

2.5.4.3 Site Residency

As an extension of BVLOS systems, recent developments have examined the benefits of full integration of a UAV agent within facility infrastructure.

Providing a permanent presence on site, resident UAVs conduct both regularly scheduled surveys and responsive ad hoc deployments, such as quickly assessing unscheduled breakdowns or security incidents. Referred to as “drone in a box” systems, these aircraft are designed to reside in a specialised storage container built into the wider plant, to which they autonomously return after each mission to recharge and upload data. The container is environmentally sealed to protect the UAV from harsh conditions, opening and closing automatically to allow access, and so enabling system operation with zero human presence.

The underlying technology for these systems is, again, well established across the fields of UAV control, environmental mapping, and navigation for obstacle avoidance. As such, a number of example systems already exist in the commercial marketplace, as shown in Figure 2-30.



Figure 2-30: Examples of commercial site-resident UAV systems. (a) The Skydio Dock [291]. (b) The DJI Dock [292]. (c) The Percepto Base [293].

Where in-person access presents a challenge, these systems offer particularly high value to operations and management activities, providing immediate and fully remote monitoring that is generally applicable to multiple assets and may function across an extensive geographic area. Once the larger challenges associated with BVLOS flights are resolved, resident UAVs may therefore be expected to gain prominence for applications such as offshore NDE, where transfer of specialist personnel to the site can cost thousands before work even starts [31].

2.5.5 Evaluation

In light of the reviewed works, a number of key themes in ongoing research and development efforts are identified, as are gaps where further investigation could directly progress the ultimate goals of airborne NDE.

It is clear that simple visual screening using UAVs is well established and fully supported in practice by a number of commercial service providers [227]–[229]. As such, innovation within this theme has focused on refinements to inspection workflow efficiency. Reflecting the wider trends of Industry 4.0 [79], UAV inspection providers have embraced online reporting to streamline their assessment and data management processes, while also improving the insight available to asset operator clients through

location metadata and temporal comparisons. Integration of automated visual analysis [230], particularly deep-learning methods for defect classification, detection and segmentation [231], within this framework may further reduce turnaround time. However, the quantity of sensitive NDE data required for training, opacity of decisions regarding safety critical defects, and limited accountability of neural networks have led to hesitant adoption within assessment standards, limiting expression of their benefits.

Contrastingly, industrial appetite for methods offering additional insight using existing visual inspection hardware has been far higher. UAV photogrammetric inspection is a prominent example. It presents well comprehended underlying theory as established in Appendix A, characterised behaviour in the face of non-ideal practical influences [234], and a range of applications demonstrated in structural geometry monitoring [115], [232], [233]. In this regard, its UAV adoption is widespread, recognising the benefits to NDE afforded by the provision of viewpoints inaccessible to terrestrial survey, the reduced time and equipment costs versus conventional laser scanning, and the complimentary nature of photogrammetric localisation and other non-visual survey data [111], [238].

Nevertheless, amid wider ongoing refinement of the underlying photogrammetric algorithms to augment robustness and overall fidelity, select developmental gaps remain within the context of airborne inspection. In spite of work establishing close-proximity photogrammetry as a way to increase target surface resolution and visibility of unique visual features, directly improving reconstruction accuracy and utility for NDE [234], typical UAV operations are still conducted at standoff on the order of 1 m and above due to safety concerns associated with near structure aerodynamic effects [235]–[237]. Specifically, no consideration has been given towards conducting targeted visual inspection or photogrammetry from a position of airborne surface contact.

Such target contact has recently become significantly more practicable with the advent of fully and over-actuated aerial manipulator platforms [257]. Their potential for direct disturbance rejection and contact-based stabilisation of a camera relative to the surface may then significantly reduce issues of motion blur during image capture. This, in turn,

could enable close or immediate proximity imaging without the degradative effects of focal blurring otherwise encountered due to high sensitivity to standoff variation [234].

Incorporating improvements at the system level, contact-based visual inspection then presents a potential means to target fine-scale surface breaking defects with extreme resolution, beyond what has been attained through adaptive path planning and free-flight control theory [235]–[237]. This is envisioned as a means to improve insight at targeted locations around the structure where degradation is statistically most likely or is insufficiently assessed by broader screening. Redundant inspection of known defect free regions may thereby be avoided while retaining compatibility with current visual inspection practices and improvements made to workflow efficiency.

As established in current literature, airborne interaction may also be exploited to deliver volumetric NDE capable of assessing indicators of target health beyond the exterior surface [258]–[260]. Again, improvement is shown over measurements using existing quadcopter platforms; restricted in terms of contact stability and repeatability due to the dynamics of the underactuated UAV platform. Removing reliance on magnetic adhesion [246], [273], large outrigger supports [244], [247], [249], [261], or target alignment with the hovering thrust direction shown by UDT systems [250] thus presents more generic applicability. The small number of aerial manipulator systems demonstrating NDE to date are then proven to be capable of directly counteracting disturbances with responsive MDT or VT methods and offer more stable probe deployment to a range of surface orientations via the additional dexterity of an embedded robotic effector [258], [260] or omnidirectional flight [259].

However, as a nascent research theme in itself, aerial manipulator literature often focusses on design and control of the underlying robotic platforms, paying little attention to the NDE application or qualification of obtained inspection results. Furthermore, extant systems have not yet matched the range of NDE techniques and applications demonstrated by manual contact-based processes [99], [257], instead favouring generic transducers that lack the benefits of more specialised equipment developed within the NDE sector.

Accordingly, a gap can be shown in literature for new works concerning additional NDE modes and conducting full characterisation of the effects of a floating platform on transducer function. Alongside detailed examination of the causal mechanisms for any degradation to inspection findings, such research is expected to be highly beneficial for establishing confidence in the ability of airborne NDE to satisfy industrial standards. Strong potential is also presented to yield novel solutions that considerably improve the level of attainable structural health insight, further reduce hazardous in-person access and work at height, and minimise cost-to-inspect for energy sector assets.

A final development opportunity is presented in recognition of the widespread adoption of surface crawler NDE robots. Utilising passive and low-energy mechanisms like magnetism, or active mechanisms like vacuum adhesion in conjunction with a tethered power supply, these securely adhere to the surface of their target structure, granting best-in-class sensor stability and motion control with near unlimited operational duration [264]. Partial adoption of these desirable features within UAV platforms has been examined for sustained presence through low-energy perching, whereby the UAV is supported by the target structure [272]. These systems also have the benefit of flight access, bypassing many obstructions encountered by crawlers. However, perching strategies reliant on grasping or resting atop certain structures can preclude access to large regions of the target not reachable from a point of passive support. An embedded robotic effector can ameliorate sensor placement access challenges, but mass and reach limits then restrict scalability.

A gap in literature thus exists for a system better solving these issues. By exploiting passive support where possible but reverting to a self-supported active adhesion mode where necessary, a robot may access the full structure exterior. Retaining the strong contact offered by surface crawlers in combination with the mobility of a multirotor UAV, a hybrid system could then flexibly deliver NDE sensors across the target via a fly-crawl-fly process with excellent stability during interaction. As no instance of a UAV using generic adhesion to omnidirectional surfaces to enable contact-based NDE has been documented in previous research, this would also represent significant novelty.

Addressing this and the other gaps in existing research, future airborne NDE technologies therefore show significant potential to further support predictive maintenance strategies, cut inspector hazard exposure, and eliminate unnecessary human intervention. This in turn can lower running costs across applicable energy sector assets and improve their operational lifespan.

2.6 Summary

This chapter has presented a detailed review of the background theory associated with airborne non-destructive evaluation alongside its application in recent literature and commercial activity.

Commencing with established NDE methods recognised as having strong applicability to aerial deployment, specific attention is given to photogrammetric reconstruction, and single element ultrasonic methods. Supported by Appendix A and Appendix B, these sections provide the theoretical background necessary to understand their operation and identify factors likely to influence their performance when utilised aboard a UAV. Where photogrammetry is noted to extend existing RVI methods to yield a cost-effective system for reconstruction of large-scale asset exterior geometry, ultrasound provides a generally applicable means to detect fine-scale flaws and degradation within the volume of a target structure. Together, these otherwise distinct approaches may therefore yield mutually complementary advantages and meaningfully enhance the level of information attainable using UAV platforms.

Following this, the operating principles of multirotor UAVs were examined. This provided a breakdown of the force-torque actuation of individual propellers and how their combined aerodynamic effects may be modelled and manipulated to grant flight. Trends in UAV control were also examined to identify beneficial features exhibited by each general class of control structure for NDE applications, with detailed assessment in Appendix C. Among those structures considered, model predictive control strategies were noted to grant superior tracking performance and flexibility to encompass a range of system architectures at the expense of significant complexity. However, accurate modelling of the system and time-varying external disturbances

required for effective real-time optimisation introduces many practical challenges. By contrast, modern PID methods also indicate strong trajectory tracking with good robustness to system uncertainty and minimal complexity. Further, they do not exhibit the non-ideal behaviours such as chattering, steady state offsets, or control signal saturation artefacts that limit the applicability of other methods to the NDE application. Availability of multiple well-supported open-source flight-control packages based on PID also makes them an attractive candidate to support prototype UAV system development.

Lastly, multiple gaps in current literature and industrial practice are identified for further investigation. It is noted that that contact-based airborne inspection remains a nascent area of study, with many works focussing on the development of aerial manipulator platforms from a robotics perspective but giving minimal consideration as to the practicality of the NDE method they deploy. Similarly, most systems propose means for stabilising contact while the main body of the vehicle maintains standoff from the surface. Opportunities to better stabilise sensor deployment via the multi-point close-range contact demonstrated by mobile robotic crawlers are therefore overlooked. Investigation of these aspects within the context of photogrammetric reconstruction and ultrasonic inspection may therefore yield results with strong practical applicability beyond current commercial offerings.

This thesis examines UAV systems with the potential to address such technology gaps. An aerial manipulator with omnidirectional thrust-vectoring capability is augmented and further developed to address probe stability issues and enable ultrasonic thickness measurement across a wide range of industrial asset geometries. Thereafter, a second vehicle designed from first principles within this project is used to explore multi-point contact for close-range remote visual inspection of pipe and other cylindrical structures. Focussed consideration is given to initial demonstration and quantitative characterisation of both novel means for contact-based airborne NDE so as to establish confidence bounds and support future adoption of the emerging technologies.

CHAPTER 3

Dry-Coupled Airborne Ultrasonic Non-Destructive Evaluation

3.1 Introduction

Review of existing commercial practice and current development trends has shown that multirotor Unmanned Aerial Vehicles (UAVs) are deployed primarily in non-contact, surface-imaging inspection processes [81]. As examined in Section 2.5.1, their small size and mobility requires minimal supporting infrastructure and permits rapid visual survey of assets with challenging manual access. A pertinent example is noted in the case of wind turbines, where inspection rates have increased from between two and five per day for manual processes to upwards of twelve using UAVs [88]. In nuclear sites UAVs have also enabled savings of over 300 working hours across a single ASME BPVC Section XI inspection campaign, reducing the time to complete the process from six weeks to two and requiring only one qualified inspector instead of three [294].

However, UAV adoption for other NDE processes employing contact-based sensors such as ultrasound, eddy current, or electromagnetic testing has been significantly more tentative. This runs counter to the advantages such quantitative NDE modes offer over simple pass-fail visual inspection based on subjective expert opinion, and their ability to inform degradation forecasting and predictive maintenance strategies [12], [97], [98]. Beyond this, contact-based NDE modes can provide increased insight into sub-surface structural health and show mature usage across a range of industrial processes, forming a significant proportion of global NDE activity [97], [295], [296].

Within the context of the energy sector inspection scenarios examined previously in Section 1.1 and the existing means for contact-based sub-surface NDE, the potential advantages of an airborne system capable of deploying ultrasonic measurement are

readily apparent. Ultrasound can assess the full volume of many target structures with single side access, in many cases requires minimal surface preparation and post-inspection clean up, is applicable to a range of common construction materials and defect mechanisms, and does not entail the extensive safety precautions of radiation-based imaging. The discrete lightweight hardware is also well suited to UAV integration.

Furthermore, potential applications for a system capable of aerial ultrasonic thickness measurement and corrosion monitoring are common in the presence of in situ human access challenges. Specific examples include assessment of large areas of petrochemical storage tank walls which can be corroded by their contents [102], with decontamination of their caustic liquid contents for manned NDE potentially entailing downtime upwards of 100 days [297] and hazardous rafting operations [298]. Within power generation and chemical processing plants, pipe networks with limited accessibility are subject to destructive effects such as flow accelerated corrosion [51]. Monitoring corrosion in steel chimney stacks [299] and regions of offshore oil and gas platforms above the waterline [300] can entail hazardous roped access. Similarly, hydrogen sulphide gas build-up can pose hazards to manned entry for internal corrosion monitoring above the waterline in the monopile of offshore wind platforms [301]. A clear value-case may thereby be formed for remote, airborne, ultrasonic NDE within each of these inspection processes; minimising manned entry in accordance with “as low as reasonably practicable” (ALARP) guidance [48], reducing asset downtime, and increasing inspection cost effectiveness. Similarly, the enhanced utilisation of robotics and automation may help alleviate NDE skilled-labour shortages [302]. These factors thereby drive significant industrial appetite for the technology.

In spite of these benefits, obstacles to uptake stem from two closely related technical issues: the NDE process requirements for repeatable, fine-scale sensor positioning to monitor degradation over time; and sustained transducer coupling with the target structure to allow undisturbed measurement signal capture. Successful resolution of these problems is necessary to deliver operations and maintenance cost reduction benefits similar to, or greater than, the levels witnessed following widespread adoption of UAVs for free-flying visual inspection at height [303].

This chapter therefore investigates an integrated solution, drawing upon recent advancements in the capability of multirotor aerial platforms to enter and maintain stable environmental contact during flight, and examining ultrasonic transducers with less stringent coupling requirements to better support their flexible application.

3.1.1 Dry-Coupling Ultrasonic Probes

Commonly, ultrasonic inspection relies on water or glycerine-based gel couplant that displaces the highly reflective air layer between the probe and target material, better matching their acoustic impedances and enabling ultrasonic wave transmission. By contrast, the less conventional dry-coupling strategies investigated here use pressure applied to a deformable material at the probe-target interface to mechanically expunge the air layer and conform to the surface microstructure, as in Figure 3-1. Such materials also offer a more compliant mechanical interaction with the target surface, potentially improving tolerance to misalignment and small position changes versus the hard-faced probes introduced in Section 2.3.3.




Figure 3-1: Couplant material fills microstructure at the probe-object interface, minimizing impedance changes. Arrows show the ultrasonic signal propagation, with relative amplitude indicated by thickness.

Their usage is established in conventional inspections when liquid couplant would contaminate or weaken target structures [304], water bath immersion is inappropriate, or subsequent clean-up operations are impractical [305]. Previous work has also seen ultrasonic probes utilising dry-coupling material deployed from static ground-affixed manipulator arm robots for composite structure inspection [173]. Advancements in material science have supported their extension to dry-coupling, solid core, wheel

probes, smaller and lighter than initial fluid core designs, for area scanning applications [306].

Within the context of aerial NDE, operation without an onboard gel reservoir will reduce payload and extend flight times, also eliminating centre of mass variation when dispensing gel. Further, there is no need to continually replenish a gel supply during extended area scanning. This makes dry-coupled ultrasound an advantageous candidate for airborne deployment of quantitative sub-surface NDE, warranting detailed investigation.

3.1.2 System Conceptualisation

As examined in Section 2.5, there are a number of strategies for aerial interaction supporting contact-based NDE. Of those, the systems in recently published literature addressing airborne ultrasonic inspection are summarised in Table 3-1. This comparison is then used to identify opportunities within the state of the art that may be directly addressed via the system developed herein.

TABLE 3-1: COMPARISON OF AIRBORNE ULTRASONIC NDE SYSTEMS

Ref.	[258]	[273]	[249]	[250]	[251]	Herein
Rotor Layout	Octo	Quad	Quad +2	Quad	Hex	Tri
Actuation	Over	Under	Under	Under	Under	Over
Interaction^a	MDT	RU	DP	RU	RU	VT
Accessible Surfaces	Omni	Floor, Pipe	Wall	Ceiling	Wall	Omni
Effector^b	Arm A	Grip R	Frame C	Rod C	Rod C	Rod R
Probe Elements	Dual	EMAT	Single	Single	Dual	Dual, Wheel
Couplant^c	Gel D	EM	Gel D	Gel M	Gel M	Dry
NDE Type	Point, Scan	Point	Point	Point	Point	Point, Scan

Key:

^a UAV interaction force generation method i.e. *MDT*: Multidirectional Thrust, *DP*: Dedicated Propellers, *RU*: Reoriented unidirectional body thrust, *VT*: Vectored Thrust.

^b UAV effector structure type i.e. *A*: Actuated, *C*: Complaint, *R*: Rigid.

^c Couplant mechanism i.e. *D*: Automatically dispensed gel, *M*: Manually applied gel, *EM*: Electromagnetic induction.

From this table, it is clear that the systems performing ultrasonic inspection commonly do so using underactuated quadrotor UAVs [249]–[251], [273]. As such, they are restricted to single orientation interactions. Stabilisation of the transducer via magnetism [273], a supporting frame [249], or by enforcing an interaction direction aligned with the unidirectional thrust [250] improves their measurement capability, but their utility as workers for generic aerial NDE processes therefore remains fundamentally limited given the range of geometry presented by in situ industrial structures.

In this context, employing an over-actuated platform presents a significant opportunity to improve repeatability of the interaction against a range of target structures while maintaining prolonged stable contact in the presence of near-surface aerodynamic influences. Moreover, a base platform capable of generating omnidirectional, null-torque, force exceeding its weight in all directions about one or more axes of rotation presents further advantage. Such a craft may align a probe rigidly affixed to its body with the target surface by hovering in the required orientation, precluding the need for additional degrees of freedom offered by an embedded robotic manipulator, as used in [258]. Platform size, and mass may thus be significantly reduced. Similarly, the control complexity introduced by a large robotic manipulator via its variable centre of mass and reaction forces during articulation is largely precluded.

Regarding the ultrasonic hardware, piezoelectric transducers are favoured in the majority of cases examined in Table 3-1 [249]–[251], [258] owing to their position as a mature, de facto standard technology with generic compatibility to a range of materials. The flexibility of EMAT excitation frequency and their generation of ultrasound directly within the target material makes them a popular choice for guided wave applications, with growing interest in other use cases [151], [307], [308]. However, their intrinsic magnetic stabilisation and electromagnetic transduction mechanism is incompatible with electrically non-conductive composite materials, such as those common to wind turbines. This, combined with the mass of permanent magnet EMAT and significant signal conditioning requirements [309], reduces potential utility within the scope of this thesis. Of the systems employing piezoelectric ultrasound, all require the application of couplant gel to ensure transduction. Whereas some include

the additional payload of an automated dispenser [249], [258], others notably rely on a manual couplant application process [250], [251], limiting their practical usage to laboratory feasibility studies. Measurements are also predominantly static in nature [249]–[251], [273], with a requirement to disengage and re-enter contact markedly slowing any potential inspection of the large surface areas common to energy sector assets.

A system incorporating a dry-coupling wheel probe thereby presents significant opportunity to expand the capabilities offered by airborne ultrasonic inspection. Alongside the signal coupling advantages indicated above, intrinsic applicability to dynamic scanning along the target surface may also be meaningfully exploited to improve sample density versus repeat point inspection. This may then streamline the inspection process and augment wall thinning feature localisation efforts across extensive asset surfaces. The novel combination of an over-actuated vector thrust UAV and dry-coupled wheel probe therefore presents several potential advantages ahead of full investigation.

3.1.3 Chapter Structure

The remainder of this chapter is constructed as follows. An overview of the proposed UAV inspection system and its operating principles is provided through Section 3.2. In Section 3.3 the experimental methodology used to demonstrate and characterise system capability is described. Section 3.4 provides an analysis of the results of this study, accompanied by further discussion in Section 3.5. Finally, conclusions drawn from this work are provided in Section 3.6.

3.2 System Overview

3.2.1 Voliro Tri-Copter

Here, the Voliro Tricopter UAV [310] is examined as a means to perform contact NDE. This aerial platform is specifically designed for remote airborne physical interaction, with key features detailed as follows.

3.2.1.1 Aerial Platform Hardware

The aerial robot used in this chapter is an omnidirectional multirotor capable of generating 6 DoF wrench. The UAV is composed of two main thrusters capable of thrust vectoring by turning around two axes, and a tail thruster capable of inverting the force direction as shown in Figure 3-2.



Figure 3-2: The Voliro tricopter aerial manipulation platform. Annotations show the body and inertial world coordinate frames and axes of actuation granting 6 DoF wrench exertion. Both propeller arms rotate independently. Bi-propeller assemblies rotate in one further axis about the tip of these arms, as indicated by blue and orange arrows.

Ability to safely maintain surface contact is an important consideration during aerial inspection. As such, the Voliro Tricopter is rated for flight in windspeeds up to 12 m/s. Optional rotor guards may be attached, protruding beyond the rotor swept volume in the body x-axis. These protect the target from rotor collision during severe disturbances and are employed during empirical assessments conducted here. Further, the 6 DoF wrench exertion capability enables the UAV to enter a configuration well suited to physical interaction, wherein external disturbances may be directly opposed while maintaining static and reliable sensor placement. Up to 30 N net force may be exerted in all directions in addition to supporting the total system mass of 4 kg.

The UAV is equipped with a Pixhawk flight controller for sensor interface, state estimation and motor control signal generation. An intel NUC i7 core computer

running the Robot Operating System (ROS) is embedded within the platform to enable autonomous features and data recording. Communication between the NUC computer and Pixhawk flight controller is enabled via a serial data port.

This network supports data from various onboard and offboard position estimation technologies for both indoor and outdoor deployment. In addition to common Global Navigation Satellite Systems (GNSS), examples include ground based optical motion capture systems (such as Vicon [311]), laser tracking via reflective prism [312], and visual odometry [313]. The most appropriate positioning technology may be selected based on its operating characteristics and the intended use case.

3.2.1.2 Control Structure

This section describes the baseline controller for the Voliro tricopter platform, used to autonomously maintain inflight pose and generate interaction force as set by the pilot. This is similar to the architecture described in [185], so is presented in summary for system context alongside modifications in support of environmental interaction using the bi-axial tilting propeller tricopter platform. Figure 3-3 gives a diagram of the controller structure, leveraging the decoupling of position and attitude dynamics intrinsic to an omnidirectional platform.

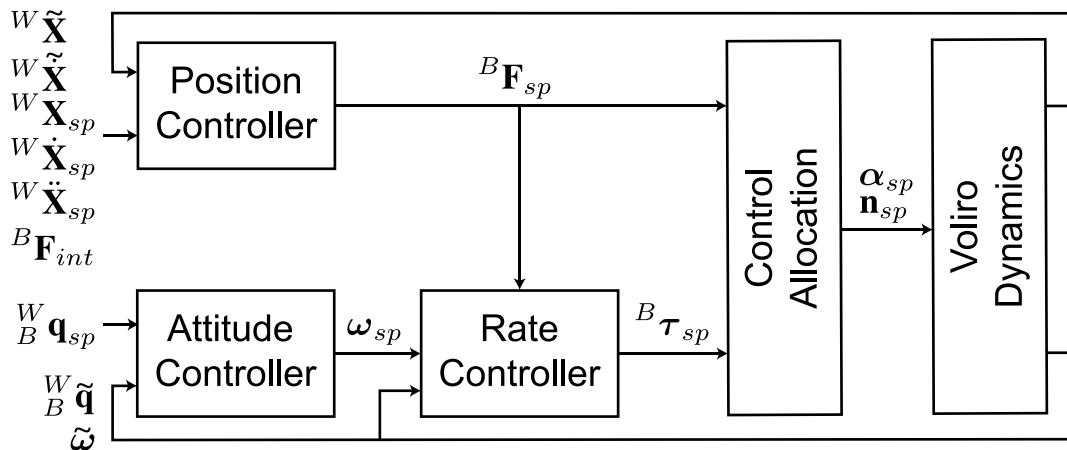


Figure 3-3: A block diagram of the Voliro control structure used to maintain omnidirectional stability and perform environmental interaction. Note the position control loop directly generates 3D desired force in this VT platform.

Here two coordinate frames are used to describe the controller function, as depicted in Figure 3-2:

- $\{W\}$, the inertial world frame, used to describe the global position of the UAV.
- $\{B\}$, the UAV affixed body frame, describing actuation and moment of inertia.

Note that a bold typeface indicates a vector or matrix parameter unless otherwise stated. The reference frame in which a parameter is examined is denoted using left-hand superscript. An estimate of vector parameter, such as UAV position, \mathbf{X} , is denoted with the tilde symbol as $\tilde{\mathbf{X}}$, to avoid confusion with the unit vector notation used elsewhere. The desired setpoint value is \mathbf{X}_{sp} and its time derivative is $\dot{\mathbf{X}}$. The rotation matrix expressing orientation of $\{B\}$ relative to $\{W\}$ is denoted ${}^W_B\mathbf{R}$ with transpose ${}^W_B\mathbf{R}^T$, whereas the quaternion expression is ${}^W_B\mathbf{q}$ with complex conjugate ${}^W_B\mathbf{q}^*$.

Here, as in the generic structure of Figure 2-12, the position controller forms the outermost loop. This generates desired force commands, ${}^B\mathbf{F}_{sp}$, based on a PID evaluation of position error in $\{W\}$, ${}^W\mathbf{X}_e$, using gain vectors $\mathbf{K}_{pos,P}$, $\mathbf{K}_{pos,I}$, $\mathbf{K}_{pos,D}$ which may contain different weights for each x-y-z component. Additive terms then balance UAV weight, $m\mathbf{g}$, and feedforward force for the desired acceleration, ${}^W\ddot{\mathbf{X}}_{sp}$. Desired interaction force is included as ${}^B\mathbf{F}_{int}$, a vector aligned with the x-axis of $\{B\}$ having zero magnitude during free flight. The result is expressed in $\{B\}$ via the inverse of the rotation matrix describing instantaneous UAV attitude.

$${}^W\mathbf{X}_e = {}^W\mathbf{X}_e - {}^W\tilde{\mathbf{X}} \quad (3-1)$$

$$\begin{aligned} {}^B\mathbf{F}_{sp} = {}^W_B\mathbf{R}^T \left(\mathbf{K}_{pos,P} {}^W\mathbf{X}_e + \mathbf{K}_{pos,I} \int {}^W\mathbf{X}_e dt + \mathbf{K}_{pos,D} {}^W\dot{\mathbf{X}}_e + m\mathbf{g} \right. \\ \left. + m {}^W\ddot{\mathbf{X}}_{sp} \right) + {}^B\mathbf{F}_{int} \end{aligned} \quad (3-2)$$

When interaction is requested, the controller transitions to a flight-mode in support of force application. While maintaining position in the other axes, ${}^B\mathbf{F}_{sp}$ in the $\{B\}$ x-axis is then augmented by ${}^B\mathbf{F}_{int}$ at a magnitude meeting the specific requirements for deployment of the contact-based NDE sensor. Known platform geometry and propeller

thrust characteristics enable open-loop force and torque generation with sufficient accuracy to meet application demands. This control block thereby serves as a hybrid force-position controller.

The attitude controller consists of two cascaded blocks. The first computes the desired body rotation rate, $\boldsymbol{\omega}_{sp}$, in proportion to the quaternion orientation error, \mathbf{q}_e , via control parameter $K_{rot,P}$. This error is computed via quaternion product and is expressed in scalar-vector component form. The vector component, $\mathbf{q}_{e,v}$, dictates the body rate direction and magnitude whereas the sign of the scalar component, $q_{e,w}$, avoids the unwinding phenomenon [314], caused by the dual representation of orientations in quaternion space.

$$\mathbf{q}_e = {}^W_B \mathbf{q}_{sp} \otimes {}^W_B \tilde{\mathbf{q}}^* = \begin{pmatrix} q_e \\ \mathbf{q}_{e,v} \end{pmatrix} \quad (3-3)$$

$$\boldsymbol{\omega}_{sp} = \text{sgn}(q_{w,err}) K_{rot,P} \cdot \mathbf{q}_{v,err} \quad (3-4)$$

The second attitude block then gives the desired body moments, ${}^B \boldsymbol{\tau}_{sp}$. These are proportional to the body rate error by constant $K_{rot,D}$, also compensating for the UAV inertia, \mathbf{J} , and the moment due to the net thrust and centre of mass offset from the origin of $\{B\}$, ${}^B \mathbf{r}_{off}$.

$${}^B \boldsymbol{\tau}_{sp} = K_{rot,D} (\boldsymbol{\omega}_{des} - \tilde{\boldsymbol{\omega}}) - {}^B \mathbf{r}_{off} \times {}^B \mathbf{F}_{sp} + \tilde{\boldsymbol{\omega}} \times \mathbf{J} \tilde{\boldsymbol{\omega}} \quad (3-5)$$

Finally, in the control allocation stage, desired arm angles, $\boldsymbol{\alpha}_{sp}$, and rotor speeds, \mathbf{n}_{sp} , minimizing total thrust are quickly calculated from the desired force and moment vector by Moore-Penrose pseudo-inverse of the static system matrix and trigonometric identity [185], [258], [262]. These command signals are then distributed to the arm actuation servos and propeller motor speed controllers.

3.2.2 Contact NDE Payload

The hardware that comprises the inspection payload consists of two main components: a dry-coupling wheel probe and a custom designed transceiver circuit board, both mounted aboard the UAV manipulator platform.

3.2.2.1 Dry-Coupling Wheel Probe

The wheel probe houses two piezoelectric ultrasonic elements with nominal centre frequency of 5 MHz [315]. These are embedded within an axial shaft, as depicted in the cross-sectional view of Figure 3-4. Surrounding the axle is a freely rotating, deformable, rubberised tyre providing minimal attenuation of the ultrasonic signal during propagation between the piezoelectric element and sample material.




Figure 3-4: A dual element wheel probe cross-sectional view and diagram of ultrasound wave propagation during measurement. The tyre outer diameter is approximately 25 mm, and its mass is 53 g including the connecting cables.

The most appropriate ultrasonic probe for a given inspection scenario varies with the defect feature of interest and target object. Among other options, a twin-element transducer such as this is considered well suited to remaining wall thickness measurements of corroded samples per current ISO [171] and ASTM [172] standards. Differing from the generic probe depicted in Figure 2-5, the V-shaped ultrasound propagation path increases sensitivity to echoes from the base of corrosive pits, useful for minimum wall thickness measurement. Further, dedicated transmit and receive elements prevent masking of the returning signal during the brief relaxation period of a piezoelectric transducer immediately following transmission, a common occurrence in thin samples. An internal acoustic barrier bisecting the tyre minimises direct crosstalk between elements and ensures the receiver can capture signals at the instant of transmission.

During inspection, application of compressive force causes deformation of the solid rubberised tyre and replicates the action of common liquid or gel couplants: displacing reflective or scattering acoustically discontinuous air pockets held at the tyre-sample interface, as in Figure 3-1. However, dry coupling mechanisms are sensitive to surface roughness and fine particulate dust as these features retain air pockets. They represent a complex effect on operability: influenced by material conformability, applied force, feature size and distribution [316], [317].

In commercial applications of this wheel probe, thickness measurements have been successful against in situ painted steel assets with typical mean arithmetic surface roughness (R_a) well above 10 μm . Brushing of the tyre surface with a thin layer of oil has also proven effective against dust particles [318]. Similar ability to measure through surface coatings to that of commercial thickness gages may be expected, typically up to 1 mm to 2.5 mm of paint [319], [320], provided the coating is non-scattering and well bonded to the substrate. Surface coatings failing to meet these criteria may require removal.

3.2.2.2 Ultrasonic Signal Transceiver

While embedded aboard the Voliro UAV, the piezoelectric elements of the dry-coupling wheel-probe are driven by a small form-factor transceiver. This hardware was custom designed for mobile robotic applications as part of prior research outside of this thesis [321]. Software forming the communications interface between the transceiver hardware and the embedded UAV PC has since been developed as part of this work to enable full integration with other real-time flight telemetry data, live data visualisation from a ground station, and on-the-fly reconfiguration of signal capture parameters. For reference, the functional components of the instrumentation are depicted in Figure 3-5.

The transceiver module contains two discrete channels: a transmission channel and a receiver channel, one for each piezoelectric element. Transmission uses a JFET transistor to briefly connect the transducer element to the DC-DC boosted supply voltage and produce a 180 V single pulse excitation at a repetition frequency of 100 Hz. Received signals pass through a discrete transimpedance amplifier and

variable gain amplifier (VGA) then are digitised to 8 bits by the analogue to digital converter (ADC) at a sampling rate of 100 MHz. Samples are buffered in FPGA memory before transfer to the UAV embedded computer (PC) via USB2.0. Available FPGA memory imposes maximum recorded signal length of 81.92 μ s per transmission, sufficient to capture thickness measurements up to 25 cm in aluminium. Recorded signals are passed through a Butterworth bandpass filter with 60 dB attenuation outside ± 2 MHz of the probe centre frequency. This mitigates noise induced by the high-power switching of UAV flight systems and any other electromagnetic interference.

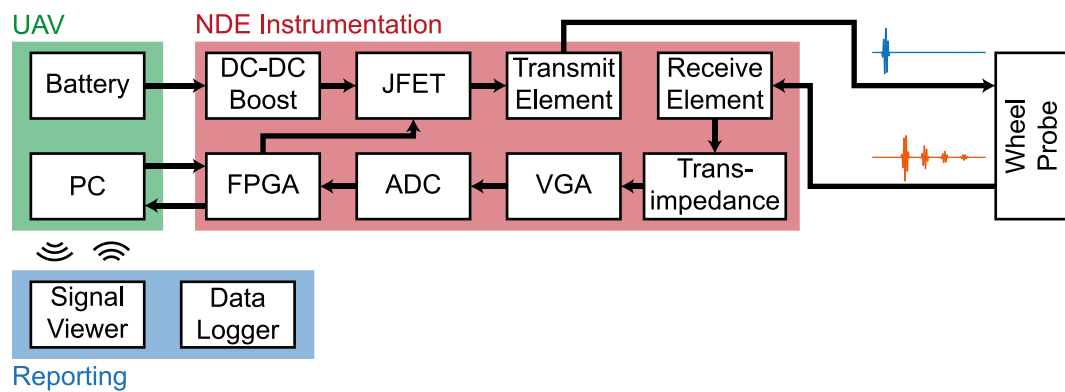


Figure 3-5: Functional block diagram of the embedded dual channel ultrasonic transceiver. Instrumentation hardware fits within a 136 mm x 61 mm x 33 mm footprint and has a mass of 78 g excluding connecting cables. Operational current draw from the onboard 12 V power supply is below 100 mA, posing minimal detriment to UAV flight duration.

Captured signals are timestamped and published in real-time to the ROS network hosted by the UAV embedded PC. This enables live display of A-scan signals via Wi-Fi link to a ground-based computer. Further, bi-directional communication through ROS is used to reconfigure instrument parameters such as gain, offset, and pulse repetition rate during flight, while also recording full ultrasound signal and UAV telemetry data on the ground control station. Numeric thickness inference and target relative measurement localization may thus be post-processed to generate a detailed inspection report, as in Section 3.3, below.

3.2.3 Acoustic Coupling Behaviour

Using the dry-coupled wheel probe, effective acoustic transduction requires that its interaction with the sample satisfy two conditions: sufficient compressive coupling force and well-aligned orientation relative to the surface normal, per Figure 3-6, below.

As an airborne floating-base, UAV deployment may introduce significant variation in these parameters. It is therefore necessary to examine the practical coupling behaviour of the probe to identify limits for feasible ultrasound transduction during interaction and to aid interpretation of the subsequent system-level evaluation. Thus, a brief empirical sensor characterisation was conducted.




Figure 3-6: Dry-coupling wheel probe coupling criteria include both compression and orientation. Compressive force acts along the probe centreline, while rotation about the two axes parallel to the target surface then define the probe pitch and yaw orientation.

This evaluation was designed to isolate and vary both the force and orientation parameters of the wheel probe during acoustic coupling interaction. To this end, the wheel probe was mounted to a KUKA KR 6 R900 sixx robotic manipulator arm, allowing precise pose control to within repeatability of ± 0.03 mm [322]. The robot was then used to instigate ultrasonic coupling to a rigidly affixed, bright-finished, steel plate of 10 mm thickness, as depicted in Figure 3-7. Next, an ideal coupling pose was identified by minutely altering the probe orientation to maximise ultrasound amplitude while maintaining a constant coupling force. The effects of angular displacement from this pose were then investigated via controlled incremental rotation about each axis.

Similarly, advancing or retracting the probe along its ideally aligned centreline would respectively raise and lower compression force, without affecting orientation.

Sensors recording force and orientation then support definition of the minimum coupling force and tolerance to misalignment. Forces and torques induced through the physical interaction are recorded at 1 kHz using an ATI Gamma Force-Torque sensor mounted between the probe and robot flange. Probe orientation is simultaneously reported at up to 250 Hz using the internal encoders of the manipulator arm. Post-processed analysis of the ultrasonic echo signals, recorded using the Tektronix DPO4054B oscilloscope, may then express any variation as a function of interaction parameters to quantitatively characterise the probe operating regions, as follows.




Figure 3-7: Experimental setup for identification of wheel probe dry-coupling requirements. Annotations denote key equipment using the coloured circles.

3.2.3.1 Force Requirements

Utilising this experimental setup, the robot may move the probe in the surface normal direction through repeatable cyclic loading from 0 N up to 100 N compressive force.

Throughout this motion, ultrasound readings are continuously recorded at a pulse repetition rate of 14 Hz using a benchtop transceiver and oscilloscope, per Figure 3-8, and are subsequently analysed to yield Figure 3-9, from which coupling force requirements may be identified.

Figure 3-8(a) shows five compressive loading cycles of the probe, each with up to 100 N acting through the force-torque sensor z-axis. Some minor forces and torques are also induced outside of this axis due to small misalignment between the primary ultrasound propagation direction and the force-torque sensor z-axis but these are inconsequential to this analysis.

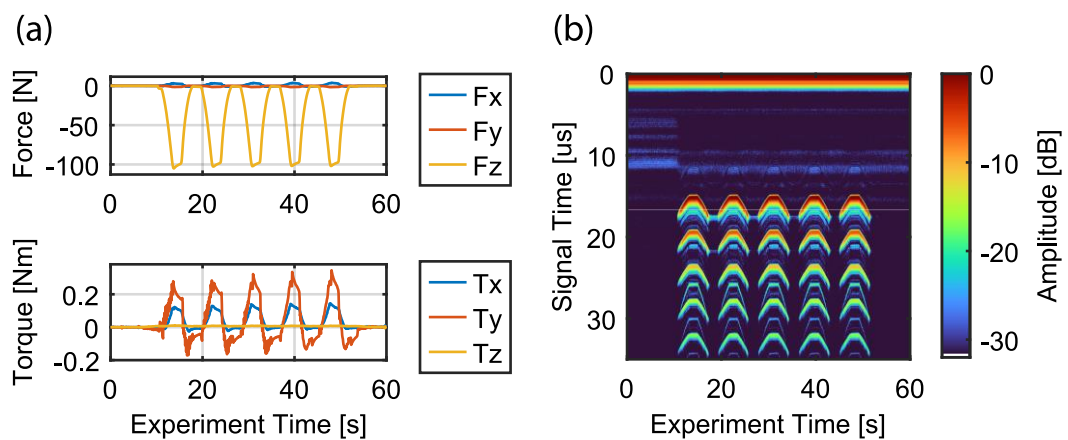


Figure 3-8: Wheel probe force-coupling behaviour assessment. (a) Force and torque acting through the wheel probe. (b) Hilbert-enveloped ultrasonic signals normalised relative to the global maximum echo signal amplitude.

The resulting ultrasound signals are recorded in Figure 3-8(b). There, each vertical column comprises the bounding envelope of a single A-scan waveform. This is computed via the Hilbert transform [323] to demodulate the echo response from the 5 MHz carrier frequency and aid visualisation. Subsequent pulses are arranged horizontally to show progression with time and probe movement. Repeatable ultrasonic coupling and de-coupling is clearly visible in correlation with the applied compressive force. Crosstalk between the transmit and receive elements is present at the instant of transmission but the time taken for the signal to propagate from the piezoelectric element to the sample and back ensures this does not affect measurement accuracy. Far smaller echoes are also visible as the signal propagates across the axle-tyre and tyre-sample interfaces, around 5 μ s and 12 μ s after transmission,

respectively, but again these pose no obstacle to signal interpretation. Notably, the arrival time of the first echo changes throughout the experiment, correlating with compressive deformation of the tyre material. Implications of this effect for thickness measurement are considered in Section 3.3.3, below. For coupling requirement elicitation, however, the force-amplitude response is of greater interest.

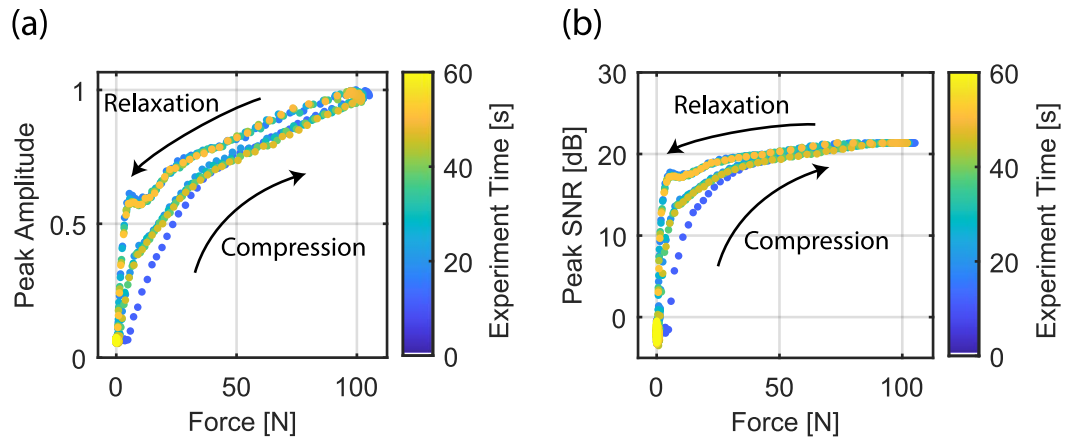


Figure 3-9: Ultrasound signal response to changing coupling forces. (a) Progression of the maximum amplitude of the recorded ultrasound signals. (b) Progression of the maximum SNR of the ultrasound signal. Annotations show the direction of the compressive hysteresis cycle.

This is examined through Figure 3-9(a), plotting the non-linear trend in peak amplitude of the first echo with compressive force. Note signal peak amplitude is normalised relative to the global maximum recorded, while data colour here indicates the sequence of their capture.

Therein, increasing compression is observed to yield rapid rise in signal amplitude up to around 10 N, consistent with the initial expungement of air from the tyre-sample boundary. Between this point and around 35 N coupling enters a transitory phase, thereafter showing a reduced but consistent rate of amplitude increase where behaviour is dictated by other, less-sensitive, effects. Hysteresis is also observed on relaxation as the rubberised tyre continues to grip the axle and sample materials, granting improved coupling versus the same force applied during the compression phase. This action may be beneficial where contact is maintained following a sharp initial collision with the target structure.

Amid these effects, it is clear that increased force offers improved coupling. However, the maximum examined force of 100 N is impractical for generation by an aerial manipulator, entailing significant energy cost and control challenges. An operating point balancing signal quality with applied force must therefore be identified.

This is supported by consideration of the response of the Signal to Noise Ratio (SNR) to applied force, as in Figure 3-9(b). Here SNR of the ultrasound signal is expressed in decibels as the ratio of the peak signal amplitude, A_s , to the peak noise amplitude, A_n , recorded in the time between transmission and the first returning echo.

$$SNR_{dB} = 20 \log_{10} \left(\frac{A_s}{A_n} \right) \quad (3-6)$$

Analysing Figure 3-9, compressive force of circa 20 N yields signal amplitude near the midpoint of the probe behaviour transition region. As a target operating point, this provides margin for error above the 10 N required for full initial coupling. Moreover, at 20 N, SNR is approximately 80 % of the maximum recorded, but requires only 20 % of the maximum force. The diminishing returns of the high-compression region are thereby avoided and a good compromise between signal quality and actuation requirements defined for the conditions under test.

Note, however, that increased surface roughness found on heavily corroded structures may trap larger air pockets at the tyre-sample boundary, reducing ultrasound transmissibility for a consistent coupling force and lowering the SNR. This effect can be mitigated by operating with increased compressive force or replacement of the dry-coupling tyre with a more conformable material, but can limit measurement feasibility if left unaddressed. Painted surfaces can further alter in situ coupling behaviour. If they become detached from the substrate and form a blister, paint coatings may block ultrasonic propagation and require removal to enable inspection of the area beneath. Nonetheless, provided paint is well bonded and conductive to ultrasound, through-coating thickness measurement remains possible using the inter-echo Mode 3 method described in Section 2.3.4 and used within the experimental assessment below.

Accordingly, the target force determined here is proposed to facilitate study in the absence of these effects, isolating the influence of the UAV platform, while also being

readily exerted by the Voliro UAV. Physical interaction at this level, akin to a firm single digit press by a human, may be supported by the vast majority of target structures, though prior structural assessment should be performed to avoid damage if inspecting lightweight structures, such as aerofoil skin or HVAC ducting.

3.2.3.2 Orientation Requirements

Acceptable orientations for wheel probe coupling are determined by the internal layout of its piezoelectric elements and the external geometry of the dry-coupling tyre and mounting structure. From first principles, optimal acoustic coupling may be expected where the probe yaw and pitch angles orient it such that the direction of ultrasonic propagation is best aligned to the surface-normal vector, as in Figure 3-6, granting a minimal transit path length through the material. The range of angles around this pose within which ultrasonic signal transmission and reception remain feasible will then define the tolerance of the probe to inadvertent misalignment and thus its ability support measurements during UAV deployment.

To identify these limits, the ultrasonic response to variation in orientation is evaluated using the experimental setup of Figure 3-7. As previously, the probe is initially aligned relative to surface to maximise echo signal amplitude while under 20 N compressive force, determining the ideal orientation. Angular changes relative to this pose are then introduced, rotating the probe about its tip in 0.5° increments. This is conducted separately in both directions about the probe pitch and yaw axes, recording the received A-scan signal while stationary at each point. Plotting the signal and its attributes as a function of the offset angle, as in Figure 3-10, then permits identification of suitable operating limits.

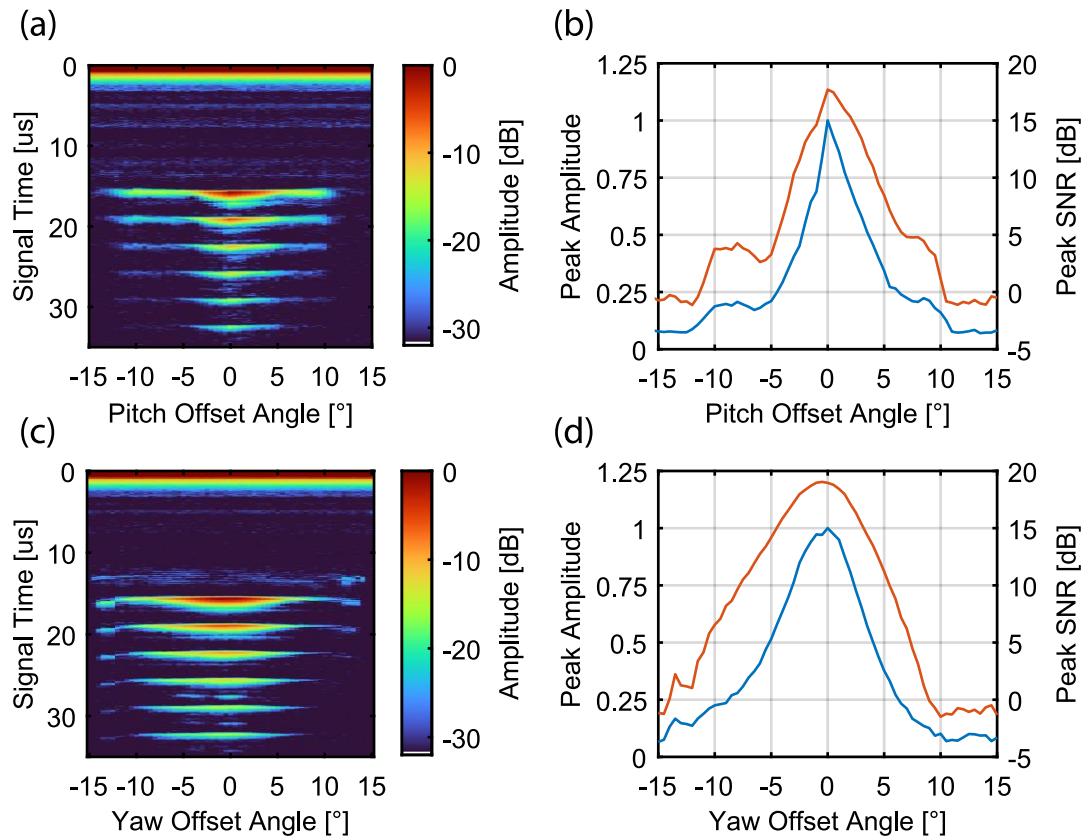


Figure 3-10: Effects of orientation on wheel probe acoustic coupling. The Hilbert-enveloped ultrasonic signals are given in (a) & (c), normalised relative to the global maximum echo amplitude in each trial. Variation of their peak amplitude (blue) and SNR (orange) are plotted in (b) & (d).

Considering the pitch response of Figure 3-10(a) & (b), it is noted that the behaviour is largely symmetric in the positive and negative angular direction, as may be expected given the symmetry of the probe in the horizontal plane, with a quick drop-off in signal strength across the region of $\pm 5^\circ$. This high sensitivity to pitch angle is attributed in part to the movement of the coupled tyre surface area relative to the internal path of the ultrasonic beam, leading to a decrease in the portion of the beam passing through the coupled area and thus a significant reduction in the energy that can be transmitted into or received from the sample. In a compounding effect, the changing angle of incidence will also lead to a V-shaped ultrasound propagation path in the pitch axis, displacing the high-energy centre of the returning beam relative to the receiver element such that only the low-energy extremities can be captured.

At angles above $\pm 5^\circ$, the pitch amplitude and SNR responses plateau briefly before approaching zero amplitude after $\pm 10^\circ$. From the echo arrival time per Figure 3-10(a)

and its consistency with the ideally aligned signal, this behaviour cannot be the result of a slower mode-converted shear wave. Instead, a potential explanation is found in reflections within the probe causing propagation pathways that present as sidelobes in its externally transmitted beam profile, as considered for a conventional circular piezoelectric probe in Appendix B.1. In small orientation ranges, the propagation paths of the sidelobes then intersect the receiver element, increasing the apparent signal amplitude relative to the diminished amplitude of the previous centreline path with only a minor propagation delay. Given the similarity of inter-echo timing per Figure 3-10(a), these effects may then meaningfully extend the functional angular range of the probe.

Defining a threshold of measurement feasibility by an SNR greater than or equal to +3 dB, the pitch axis coupling response of Figure 3-10(b) then suggests an orientation limit of approximately $\pm 10^\circ$ from the surface normal. Note, however, that this may be expected to vary with thickness of the target material as, for a consistent ultrasound angle of incidence, the surface-parallel displacement relative to the receiving element of an ultrasound beam returning via a V-shaped path is proportional to material thickness, influencing its recorded amplitude. The values given here therefore serve as a nominal indication of tolerance for a common use case and set expectations ahead of system-level evaluation in Section 3.4.

Considering the yaw coupling of Figure 3-10(c) & (d), a response with reduced sensitivity to orientation versus the pitch axis is apparent due to the distinct transmit and receive elements and their offset from the probe centre, not present in the pitch axis. This is visible in the flatter peak of the yaw SNR response given in Figure 3-10(d) versus that of pitch response in Figure 3-10(b), indicating reduced drop-off during small initial displacements from the ideal alignment. Quantitatively, at an offset of $\pm 5^\circ$, the yaw response retains SNR above +11 dB while the pitch response retains only +3.2 dB. In the yaw case, the geometry of the dual-element probe ensures that an angular offset increases the coupling of one element while reducing that of the other. Consequently, a consistently higher transmissibility is presented across the full signal path than in the pitch case, where angular offset reduces coupling of both elements simultaneously.

Note however that rotation about the yaw axis perpendicular to the wheel probe axle will bring its mounting structure into contact with the sample surface at offset angles over $\pm 12^\circ$. This leads to a discontinuity at the extremities of Figure 3-10(c) where it is necessary to apply a small corrective translation to the probe, backing it away from the surface to maintain desired coupling force and prevent it being crushed between the robot effector and the rigidly mounted sample. As the floating nature of an aerial platform will ensure that no forces sufficient to cause damage are applied during practical deployment, and this event occurs in the region where signal amplitude is already minimal, negligible influence is imparted upon the yaw coupling relationship.

Further insight into the probe yaw behaviour is then obtained from Figure 3-10(d). Therein it is clear that, for a consistent absolute yaw offset from the ideal coupling pose applied in either direction, overall signal propagation is asymmetric, being strongest when the receiver element is preferentially coupled. In this case, it is thought that the increase to coupled tyre area on the receiving side can better intercept the returning lower-energy beam and allow higher total signal energy capture than a smaller receiving area missed by the centre of a higher-energy beam due to the changing V-shaped propagation path.

In the presence of these effects, examination of the SNR response in Figure 3-10(d) using the same +3 dB threshold for measurement feasibility thereby indicates yaw offset limits of -11° and $+8^\circ$ in the sense indicated by Figure 3-6. Again, these values may increase for target structures with reduced thickness below this nominal point owing to effects of the reduced propagation path length.

3.2.3.3 Overview of Operational Coupling Requirements

Collecting the results of the empirical characterisation above, the full force and orientation coupling requirements may be defined as in Table 3-2. These parameters describe the desired physical interaction characteristics for ultrasonic measurement using the dry-coupling wheel probe and indicate tolerance limits for operation outside of this nominal interaction.

TABLE 3-2: SUMMARY OF WHEEL PROBE ACOUSTIC COUPLING CRITERIA

Parameter	Desired Operating Point	Tolerance
Compression Force *	20 N	-10 N, +80 N
Pitch Angle Offset	0°	±10°
Yaw Angle Offset	0°	-11°, +8°

* The coupling response to compressive force is tested up to 100 N. Behaviour outside of this region is not examined in light of typical UAV force generation capabilities.

As such, where a hard-faced probe has previously been shown unable to return an echo signal beyond a pitch offset angle of $\pm 3^\circ$ or a yaw offset of $\pm 6^\circ$ [251], the compliant dry-coupling probe presents markedly relaxed interaction requirements. Further support is then given to the case for its utilisation within airborne contact-based ultrasonic inspection in the presence of multiple intrinsic aerodynamic disturbances.

3.3 Experimental Assessment Methodology

With criteria for successful airborne ultrasonic coupling established, several trials were conducted within an indoor laboratory flight volume to quantify performance of the integrated system. Details of this facility and the assessment strategy employed are provided below to aid with understanding of results and support future assessments of other similar airborne NDE systems.

3.3.1 Flight Testing Facility

Experiments were conducted within a dedicated flight volume. This space included a flexible mounting frame: a repositionable planar wall upon which target inspection samples could be rigidly affixed. A similar structure allows for the positioning of test-pieces on the underside of an overhanging surface with adjustable inclination.

As the facility is indoors, the multicopter platform must operate in a GNSS denied environment, as is often the case in industrial settings. It may instead use localization provided by multiple Vicon motion capture cameras [311], giving full position and orientation measurement at 100 Hz.

Whereas the Voliro craft is similarly capable of environmental interaction when employing other localization technologies during outdoor operation, this setup is common to laboratory testing of aerial robotics. It provides accurate reference data for quantitative performance assessment and may be purposefully degraded to examine UAV controller function amid pose uncertainty.

The testing volume also includes a power tether, providing high voltage DC current to the UAV from a domestic supply via an onboard step-down converter. This is not necessary for flight but may replace standard LiPo batteries for practical convenience.

3.3.2 Inspection Samples

This section describes two samples used in the analysis of the inspection capability of the dry-coupling over-actuated remote thickness measurement system. An overview of their physical features supporting this assessment is first given, with the samples depicted in Figure 3-11 and their thickness dimension details listed in Table 3-3 and Table 3-4. The method used to determine a “ground truth” reference thickness for comparison with airborne ultrasonic measurements at any surface point then follows.

Thin plate and stepped bar geometries are chosen to be representative of defects arising from large scale corrosive material thickness loss in petrochemical storage tanks and similar inspection scenarios. Both are constructed of aluminium, highlighting capability for aerial contact interaction without relying upon magnetic adhesion.

Each sample is used to assess distinct aspects of system performance. The plate sample provides a large surface area over which the positional repeatability of interaction may be assessed. Additionally, the thinnest features enable performance characterization where the measured distance approaches the accepted practical limit of one ultrasonic wavelength in the material under test (less than 1.3 mm for a 5 MHz wave in aluminium) [171]. Concertedly, the bar sample is designed to assess the resolution of changing thickness during dynamic scanning of the wheel probe along a target surface. Geometry of the bar contains precision manufactured 1 mm stepped changes in thickness from 31.5 mm to 17.5 mm then further 0.1 mm steps down to a minimum thickness of 16.5 mm over a length of 500 mm at a step width of 20 mm.




Figure 3-11: Photographs and as-manufactured dimensions of the aluminium plate sample, (a) & (b), and stepped bar sample, (c) & (d). All dimensions are in millimetres. Thickness dimensions are the mean within each region, with the difference from this value illustrated via the colourmap. Note dimensions in (d) also indicate step thickness change between adjacent faces. The exterior surface finish of the plate is that of rolled aluminium with mild discoloration from natural ageing and no further processing. The bar sample faces are precision milled to ensure the uniform thickness of each step.

TABLE 3-3: PLATE SAMPLE THICKNESS AS MANUFACTURED

Region ID	Mean [mm]	± Max. Difference [mm]	Region ID	Mean [mm]	± Max. Difference [mm]
1	8.20	0.02	4	4.47	0.02
2	3.08	0.38	5	8.20	0.01
3	8.21	0.03			

Regions in Figure 3-11(a) are numbered top to bottom.

TABLE 3-4: BAR SAMPLE THICKNESS AS MANUFACTURED

Region ID	Mean [mm]	± Max. Difference [mm]	Region ID	Mean [mm]	± Max. Difference [mm]
1	31.46	0.03	14	18.60	0.07
2	30.56	0.02	15	17.60	0.07
3	29.54	0.01	16	17.51	0.07
4	28.56	0.01	17	17.40	0.08
5	27.57	0.02	18	17.29	0.08
6	26.58	0.02	19	17.19	0.09
7	25.59	0.04	20	17.08	0.09
8	24.60	0.05	21	16.98	0.09
9	23.62	0.04	22	16.88	0.09
10	22.65	0.06	23	16.78	0.09
11	21.62	0.07	24	16.68	0.10
12	20.63	0.06	25	16.59	0.12
13	19.60	0.06			

Regions in Figure 3-11(b) are numbered left to right.

Reference geometry of the samples is captured using calibrated micrometre callipers at regular intervals across the sample. Each sample is then reconstructed mathematically as a collection of planes, minimizing the total least squares distance from the measured points of to their fitted plane via Singular Value Decomposition (SVD) [324]. These objects have the benefit of encoding spatial variation in thickness across the sample, providing means to represent the location and depth of multiple fine-scale machined features. Small dimensional changes caused by manufacturing tolerances like non-parallel front and rear faces are also captured, enabling system assessment at a significantly greater level of detail than taking a single global average thickness value for numeric comparison.

The “ground truth” reference thickness at a given point may thus be determined by intersection of a line passing though that point with the plane which describes the sample back-wall, per the line-plane intersection equations for 3D geometry [325], as follows. This is illustrated in Figure 3-12.




Figure 3-12: Diagram showing extraction of reference object thickness at measurement locations using line-plane intersection. The grey plane is the object front-wall containing the measurement point, while the blue and orange planes represent back-walls. Note that a back-wall plane need not be parallel to the front wall, as depicted by the orange section.

The generic 3D line may be defined by a unit direction vector, $\hat{\mathbf{v}}$, and a known point on the line, \mathbf{r} . All coordinates lying on this line, including its point of plane interception, \mathbf{d} , are thereby expressible as the sum of \mathbf{r} and a multiple of $\hat{\mathbf{v}}$ by an appropriate scalar, t , such that

$$\mathbf{d} = \mathbf{r} + t\hat{\mathbf{v}} \quad (3-7)$$

The intercepted plane may likewise be defined by its surface-normal unit direction vector, $\hat{\mathbf{n}}$, and a known point on the plane, \mathbf{p} , such that \mathbf{d} and all other points on the plane satisfy the expression below.

$$(\mathbf{d} - \mathbf{p}) \cdot \hat{\mathbf{n}} = 0 \quad (3-8)$$

Combining these two equations, t , being the distance from \mathbf{r} to \mathbf{d} , may be found by

$$t = \frac{(\mathbf{p} - \mathbf{r}) \cdot \hat{\mathbf{n}}}{\hat{\mathbf{v}} \cdot \hat{\mathbf{n}}} \quad (3-9)$$

The coordinates of \mathbf{d} are then expressed by substitution of (3-9) into (3-7).

$$\mathbf{d} = \mathbf{r} + \left(\frac{(\mathbf{p} - \mathbf{r}) \cdot \hat{\mathbf{n}}}{\hat{\mathbf{v}} \cdot \hat{\mathbf{n}}} \right) \hat{\mathbf{v}} \quad (3-10)$$

Direct comparison between the ultrasonic measurements and manufactured geometry is thus enabled within the sample-relative coordinate frame $\{S\}$, as depicted in Figure 3-12. In this context, the front wall is represented by the datum $\{S\}$ x-y plane, containing the point of measurement, \mathbf{r} . The thickness line passes through \mathbf{r} , with direction parallel to the $\{S\}$ z-axis and perpendicular to the front-wall. It then intercepts the back-wall at \mathbf{d} . This remains true regardless of the back-wall orientation, as described by $\hat{\mathbf{n}}$. Reference thickness is then found via direct computation of t via (3-9).

Under this representation, finer features, such as shallow surface scrapes or other minor variations in surface texture, are not captured by the fitted planes. However, effects of any such features are negligible, owing to insensitivity of dual-element probes to surface roughness, advantageous when measuring corroded material [326]. Additionally, small scale surface warping or curvature is indistinguishable within the footprint of the transducer. These attributes justify simplification of the sample representation to construct the front face as a single geometric datum plane, with the back-wall planes retaining their relative offset and orientation. This approach thereby provides robust and accurate reference thickness data for direct numeric comparison with ultrasonic measurements.

3.3.3 Ultrasonic Thickness Measurement

As indicated by the introductory materials of Section 2.3, when assessing thickness ultrasonically, the geometry of the sample under test is inferred by directly measuring the time-of-flight (ToF), T , of the ultrasonic wave packet as it travels into the sample, reflects off the back-wall boundary and returns to the receiving transducer.

The material thickness, d , is then found by using this time, the speed of sound in the material for the propagating wave-mode, v , and an integer factor, n , correcting for the number of traversals of the material thickness, calculated as

$$d = \frac{vT}{n} \quad (3-11)$$

Note however, that this equation assumes that the wave transits the sample via a path directly across its thickness. Measurement error may then exist where the actual wave propagation differs from this assumed path. This is briefly examined in the context of a dual-element probe below, with greater detail provided via Appendix B.2. By convention a constant speed of sound is also assumed, as obtained from a table of nominal values [327] or determined via calibration against a sample of the same material and of known dimensions, per the relevant inspection standard [171], [172]. In this “pulse-echo” configuration, n again has the value 2 as T measures propagation of the ultrasonic wave-packet following reflection from a rear boundary then returning to the transducer on the front: a path through double the material thickness.

Thickness measurement accuracy thus depends on the confidence in the wave propagation speed, knowledge of its propagation path, and the ability to extract the time-of-flight from the recorded signal. However, effects specific to the dry-coupling dual-element wheel probe and its experimental usage must also be considered.

Speed of sound for each sample is determined experimentally using a manual calibration procedure against known geometry under optimal dry-coupled conditions. Results presented in Table 3-5 are within the standard range for rolled aluminium accounting for small metallurgical variations [327].

TABLE 3-5: SAMPLE ULTRASONIC WAVE VELOCITY CALIBRATION DATA

Sample	Number of Readings	Reference Thickness [mm]	Calculated Speed of Sound [m/s]	Standard Deviation [m/s]
Plate	2000	8.211	6417	14
Bar	2000	31.460	6405	4

In measuring the time-of-flight to infer thickness, note that thinning deformation of the dry-coupling transducer under pressure causes variation in the arrival time of the first echo from the tyre-sample boundary, depicted at t_1 in Figure 3-13. This highlights the behaviour of the ultrasonic signal when under compression, as previously observed in the experimental results of Figure 3-8.




Figure 3-13: A synthetic ultrasonic A-scan signal waveform showing the inter-echo time-of-flight, T_k . Increasing tyre compression shortens the first echo transit time, t_1 , and improves coupling, increasing signal amplitude. Inter-echo time is unaffected by compression and can be calculated when two or more back-wall echoes are available.

As illustrated, a stronger interaction force between the UAV and inspection sample compresses the rubberised tyre and shortens the return path length, also non-linearly increasing the signal amplitude. Without significant additional sensor hardware and UAV control complexity to ensure highly consistent tyre compression, or compensation for this changeable delay using real-time force readings and extensive prior calibration, any attempt to measure thickness using the Mode 1 or Mode 2 methods described in Section 2.3.4 will exhibit significant variance correlating to minor changes in UAV pose.

However, the inter-echo time-of-flight, T_k , between the k^{th} and $(k - 1)^{th}$ back-wall echoes, themselves arriving at times t_k and t_{k-1} respectively, is defined for $k \geq 2$ and independent of tyre compression, being set purely by the sample thickness.

$$T_k = t_k - t_{k-1} \quad (3-12)$$

Thickness is then more consistently determined via the Mode 3 thickness measurement described in Section 2.3.4. Here, this is performed by extracting the inter-echo ToF from the A-scan signal via autocorrelation: a method insensitive to variation in t_1 , as depicted in Figure 3-14. Tyre compression then influences only the signal amplitude.

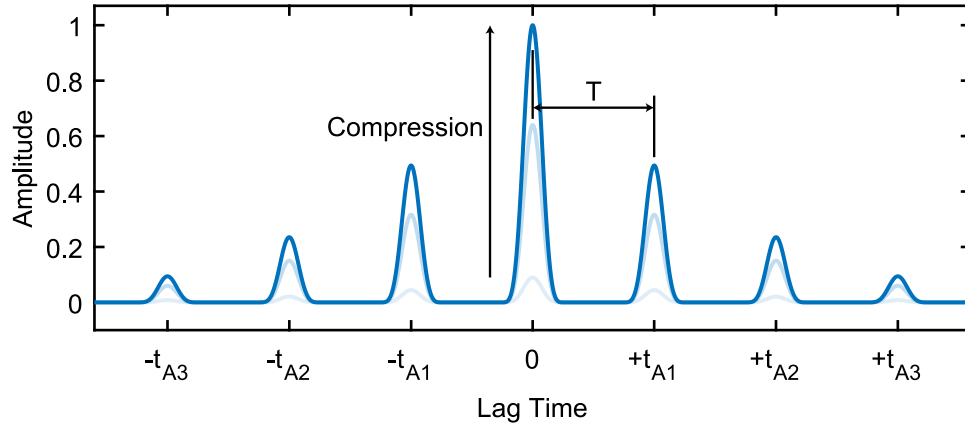


Figure 3-14: The auto-correlation of the synthetic A-scan signals depicted previously. Amplitude is globally normalised relative to the largest peak and increases with compressive coupling force. Note that despite variation in the arrival time of their first echo, the auto-correlated signal peaks occur at consistent lag times, denoted t_{Ai} and describing the relative offset between the two copies of the signal. This corresponds to alignment of the A-scan echo peaks and so is related to the inter-echo ToF, T .

Processing the A-scan signal as shown, application of the auto-correlation function returns a waveform symmetrical about zero lag. The inter-echo period may then be readily extracted via peak detection, taking the first positive lag time at which the amplitude reaches a local maximum, indicated by $+t_{A1}$. This peak corresponds to the time-of-flight between subsequent echo signals granting the best alignment of all recorded reverberations across the sample, weighted by their relative amplitude. Thickness may then be determined via (3-11) without the need for detailed manual review of each signal comprising the inspection dataset. Such a method also permits measurement of material beneath ultrasonically transmissive exterior coatings, e.g. paint, without manual gating, owing to its basis in inter-echo time-of-flight and the small amplitude of coating-substrate boundary reflections relative to the back-wall echoes.

Correlation based time-of-flight measurements such as this are well established in literature [328], [329] and are noted for their performance in cases of low SNR. In some regards they are considered optimal, owing to their use of all information held

in the signal [328]. However, a dual element probe, as considered here, will exhibit a small bias in the extracted time-of-flight due to the V-shaped propagation path illustrated in Figure 3-4. This effect arises due to transmit and receive element separation and is most consequential where this distance forms a large component of the total propagation path length, typically in samples thinner than 3 mm [172]. By this effect, inter-echo time-of-flight, depicted as T_2 , T_3 and T_4 in Figure 3-13, will vary slightly between sequential A-scan echoes. Accordingly, the time of the first peak with positive lag, $+t_{A1}$ in Figure 3-14, will not be exactly equivalent to the ToF for a pulse propagating via the shortest path directly across the sample width. Instead, $+t_{A1}$ is akin to an amplitude weighted average of across all recorded echoes due to the properties of autocorrelation, exhibiting small variation due to excitation pulse width and any signal noise.

To quantify the resultant measurement uncertainty within the context of the system assessment herein, Figure 3-15 depicts optimally coupled measurement of the thinnest region of the plate sample previously shown in Figure 3-11(a), where readings are most susceptible to V-path effects. Measurement is conducted by inter-echo ToF as in (3-12), and examined to establish worst-case bounds for the autocorrelated result.




Figure 3-15: The trend in measured thickness using the inter-echo time-of-flight, T_k . This is captured in the thinnest region of the plate sample, where the V-path effect is most significant, under optimal dry-coupling conditions. The trend formed by sequential inter-echo measurements is illustrated by the curve $f(k)$, fitting a generic asymptote curve via non-linear least squares error minimisation. The limit of this trend as $k \rightarrow \infty$ then corresponds to the asymptotic convergence of the V-path inter-echo thickness to the value measured directly across the sample width, as discussed in Appendix B.2.

Using the speed calibration data of Table 3-5 and converting propagation time to thickness via (3-11), inter-echo ToF ultrasonic thickness measurements of the thinnest sample region are thereby shown to be subject to a maximum under-sizing thickness uncertainty of up to 0.12 mm versus readings using the direct transit of the part thickness in the absence of V-path effects.

In examining the trend of Figure 3-15, note that the V-path taken to produce the multiple back-wall echoes received at t_1 , t_2 , t_3 and t_4 , will always be longer than the path which directly transits the component thickness. The additional time-of-flight in the V-path also proportionally diminishes as the number of reverberations increases and the horizontal element separation becomes a progressively lesser component of the total multi-echo path length. Perhaps counter intuitively, the inter-echo ToF used in Mode 3 thickness measurement, as described in Section 2.3.4, then initially under-sizes the component thickness owing to the relative decrease in additional path length between subsequent echoes. By the same effect, later inter-echo measurements then rise asymptotically towards the direct transit reading. A full discussion of this behaviour is provided in Appendix B.2.

In practice, however, these V-path effects are mitigated by the amplitude weighted averaging of multiple inter-echo times of flight imposed through autocorrelation, granting a reading closer to the true thickness. Other strategies such as empirical correction curve generation may further reduce error across the range of thicknesses under test [172]. Any residual error will then form a component of the thickness measurement uncertainty in the system-level assessment of Section 3.4.

3.3.4 Sample-Relative Measurement Localisation

To quantitatively profile system inspection performance, measurements taken while in contact with the sample must be identified and their results compared to reference geometry.

For the purposes of this assessment, null measurements taken outside contact, or when contact has failed to meet the acoustic coupling criteria, are identified by a peak amplitude below a threshold level set approximately 12 dB above the embedded

instrumentation root-mean-square (RMS) noise floor to ensure accurate sizing in successful measurements. This permits their direct identification based solely on the captured ultrasound signals.

Positions of the valid ultrasound readings relative to the sample are derived from the UAV pose estimate provided by the flight volume Vicon motion capture system when fused with the onboard Inertial Measurement Unit (IMU) data. Prior assessment of Vicon tracking performance for a similar setup has shown spatial pose estimation accuracy with a mean error of 1.48 mm and precision of under 10 μm [330]. In this case, IMU data fusion is conducted using the Extended Kalman Filter (EKF) functionality included in the Px4 flight controller and acts to improve sampling rate and reduce uncertainty in the pose estimate [331]. It may also be similarly applied when operating under other positioning technologies.

For the purposes of ultrasonic measurement localisation and assessment herein, this work uses the timestamps of both the ultrasound and UAV pose estimate messages, provided by a consistent computer clock, to determine the full 6 DoF pose of the craft in the world frame at the instant of thickness measurement by linear interpolation of its position and interpolation of its quaternion orientation via the Spherical Linear intERPolation (SLERP) method [332]. In contrast to purely linear quaternion interpolation, SLERP assumes a constant angular velocity between orientations. Movement during the 82 μs ultrasonic signal captured period is deemed negligible at the recorded flight speeds.

The location of each ultrasonically measured point is then expressed relative to the sample by projection from this UAV pose at the instant of signal capture. This is determined using line-plane intersection [325], as illustrated in Figure 3-16.




Figure 3-16: Projection of ultrasonic measurement locations onto the sample face. Line plane intersection is used to project the point of intersection from the instantaneous UAV pose, along the rigid arm where the probe is mounted, to a location on the front face of the fixed sample.

Note that this calculation is conducted in the world frame $\{W\}$ but references parameters from both the sample and UAV body frames, $\{S\}$ and $\{B\}$, respectively. These properties are indicated by leading subscript and superscript, such that ${}^W_B\mathbf{O}$ represents the origin of $\{B\}$ expressed in $\{W\}$, i.e. the instantaneous position of the UAV. Similarly, the unit z-axis of $\{S\}$ is expressed in $\{W\}$ as ${}^W_S\hat{\mathbf{z}}$, corresponding to the direction pointing into the front surface of the sample.

By computing the line along which measurements may be captured, this approach ensures projected locations are coplanar with the sample datum face and accounts for variable compression of the wheel probe tyre that slightly alters their location relative to the UAV. The measurement line is then defined by a unit vector parallel to the UAV probe arm, having direction ${}^W_B\hat{\mathbf{x}}$, and intersecting the $\{B\}$ origin, so passing through point ${}^W_B\mathbf{O}$, which is also the UAV centre of mass. Note that the axis of the end effector passes through the centre of mass as a feature of the Voliro design, minimizing the turning moment caused by interaction forces. The sample datum plane is defined by a unit vector opposing ${}^W_S\hat{\mathbf{z}}$ and contains the fixed point at the origin of $\{S\}$, ${}^W_S\mathbf{O}$.

Accordingly, the measurement location is then expressed in $\{W\}$ as ${}^W\mathbf{r}$, computed by adaptation of the solution to the line-plane intersection given in (3-10), as below.

$${}^W\mathbf{r} = {}^W_B\mathbf{O} + \left(\frac{({}^W_S\mathbf{O} - {}^W_B\mathbf{O}) \cdot {}^W_S\hat{\mathbf{z}}}{{}^W_B\hat{\mathbf{x}} \cdot {}^W_S\hat{\mathbf{z}}} \right) {}^W_B\hat{\mathbf{x}} \quad (3-13)$$

Registration of the sample is performed to empirically define its location within the world frame. The positions of multiple known points on the sample are measured using the Vicon system and used to produce a homogeneous transformation matrix, ${}^W_S\mathbf{T}$, via a least-squares rigid point cloud fitting algorithm [333]. The inverse of this transform then maps the world frame location of each ultrasound reading, ${}^W\mathbf{r}$, to the sample coordinate frame representation, ${}^S\mathbf{r}$, wherein the sample reference geometry may be queried for thickness at that location using the method described in Section 3.3.2.

$${}^S\mathbf{r} = {}^W_S\mathbf{T}^{-1} {}^W\mathbf{r} \quad (3-14)$$

Thickness error statistics are then obtained by comparison with the ultrasonically measured value.

3.4 Results

A number of experiments are conducted to demonstrate and assess the integrated ultrasonic inspection system¹. Whereas UAV flight telemetry and ultrasonic signals may be remotely viewed in real-time, the more detailed target relative localization and thickness reference comparison analyses are generated by post-processing of recorded data.

3.4.1 Vertical Wall Thickness Measurement

First, a simple interaction with the vertically mounted aluminium plate, as in Figure 3-17, is considered with the UAV manually piloted to engage at three points of different thickness. Surface contact is disengaged between static point measurements. The craft is then moved to the next point and the probe aligned to the surface before

¹ Video of these experiments is available via <https://youtu.be/uBNXE2eyTII>.

contact is re-entered for measurement. From take-off to landing, this process was completed in under 48 s.

Remaining in contact for a minimum period of 2 s at each point, the 100 Hz pulse repetition frequency allows capture of multiple thickness measurements. For illustrative purposes, a single ultrasonic echo signal or “A-scan” from each measurement location is presented in Figure 3-18. The characteristic echo signal is clearly visible, indicating that the system is able to successfully maintain probe orientation and compressive force during contact within the interaction tolerances expressed in Table 3-2 and validating its use as a platform for airborne ultrasonic NDE.

Considering all readings at each point, Figure 3-19 gives the mean measured numerical thicknesses, their standard deviations, and plots error versus reference geometry across the trial flight. Times at which the probe was in contact supporting measurement are highlighted using the timestamps of recorded A-scan signals. The error quantization bands are a consequence of the sampling period of the NDE instrumentation ADC and the propagation speed of the longitudinal ultrasonic wave. Given a sampling period of 10 ns and speed of approximately 6400 m/s, (3-11) dictates a minimum thickness resolution marginally above 0.032 mm, as is visible here.




Figure 3-17: The UAV positions and applies the probe to the plate sample face. This action satisfies the coupling pressure and orientation requirements to record thickness measurements. Rotor guards protecting against accidental surface collision are attached to the UAV body during testing. Power is supplied by the DC tether for convenience.

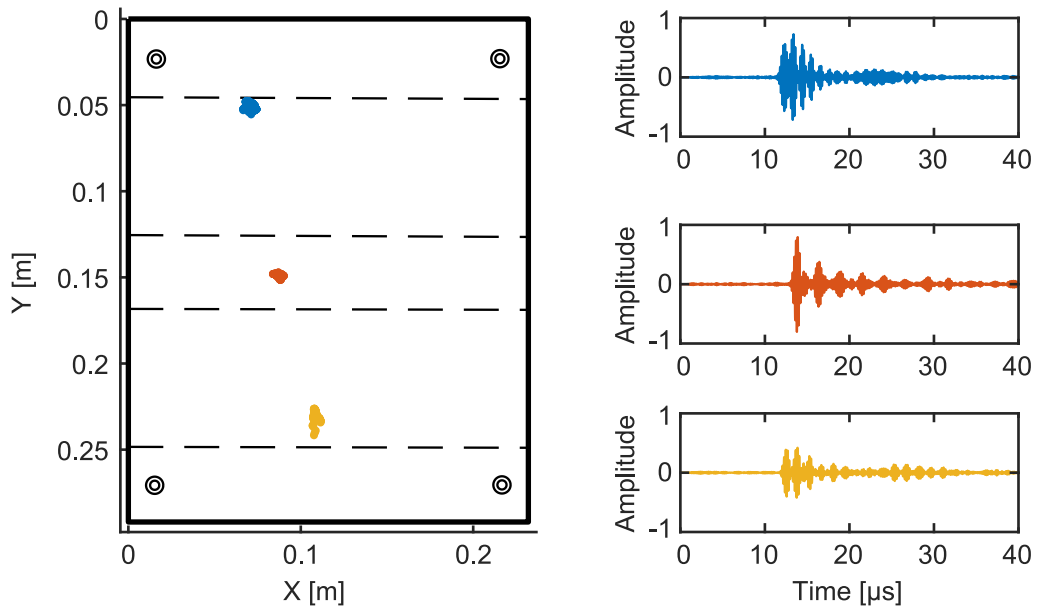


Figure 3-18: A map of the points on the sample outer face inspected in order of top to bottom during the test flight. Recorded ultrasonic echoes from each point to the correspond to the locations indicated opposite. Signals are normalised relative to the maximum amplitude recorded in the trial.




Figure 3-19: Error in the remote measurement versus reference geometry plotted across the time of the flight trial. Numerical annotations give mean thickness measurements and standard deviation. Highlights demarcate times of surface contact supporting measurement.

By analysis of the measurement data illustrated in Figure 3-19, a quantitative profile of system measurement accuracy in this scenario may be presented via the statistical data of Table 3-6. Therein, mean absolute error (MAE) describes system precision, indicating the distribution of successful measurements relative to the reference thickness. Mean Error (ME) indicates overall system accuracy, averaging data from

multiple readings to mitigate random effects and evaluate any systemic error. Typical lower and upper error boundaries of any single reading are indicated by the 5th and 95th percentile metrics, respectively.

TABLE 3-6: WALL THICKNESS MEASUREMENT ERROR STATISTICS

Number of Readings	Mean Absolute Error [mm]	5th Percentile Error [mm]	Mean Error [mm]	95th Percentile Error [mm]
766	0.0766	-0.1766	0.0112	0.1068

MAE below 0.1 mm indicates a level of system performance comparable to some commercially available hand-held ultrasonic thickness gauges with similar pricing to the instrumentation embedded aboard the UAV [319], [334], [335]. Additionally, system ME magnitude is less than the 0.032 mm minimum single measurement resolution for this instrumentation. Both ME and MAE are superior to thin sample measurements reported in literature using commercial, UAV-targeted, instrumentation, which showed consistent errors of 0.6 mm across multiple readings [253]. An asymmetrical error distribution is described by the 5th and 95th percentile errors, demonstrating a slight bias towards undersized thickness measurements, which may be attributed to residual V-path effects. However, the overall measurement accuracy shown in the positive mean error indicates that undersized measurements, though larger in magnitude, are less frequent. Despite the relative infancy of contact based aerial inspection, the observed error statistics then demonstrate measurement accuracy and repeatability sufficient to support practical utilisation.

3.4.2 Point Inspection Repeatability

In this experiment to profile system capability for repeated stable measurement interaction, the integrated UAV system is manually piloted through multiple interactions with the plate sample, targeting the marked target position. The locations of these measurements are presented in Figure 3-20.




Figure 3-20: Map of inspection points acquired through repeated interaction with upper right quadrant of target plate sample showing position uncertainty boundary. The targeted point on the plate is marked '+'. Point labels indicate the sequence of acquisition.

Position encoding of the measured thicknesses relative to the sample is projected from the UAV pose per Section 3.3.4 and so exhibits uncertainty related to both the translation and rotation estimates. This uncertainty will vary depending on the positioning technology used. In attempting to correct position error versus an estimate with excessive noise, the controller response may knock the UAV out of contact with the target surface. However, fusion with onboard inertial data mitigates global position noise effects and grants the stable relative pose estimate necessary for NDE interactions under a range of indoor and outdoor positioning technologies.

An indication of the system measurement localization accuracy under laboratory conditions is provided visually in Figure 3-20 and numerically in Table 3-7. These uncertainty radii quantify the worst-case distance from the mean location of each contact instance to the furthest single projected location captured in the same interaction.

TABLE 3-7: MEASUREMENT REPEATABILITY DATA

Point Index	Number of Samples	Distance to Mean Centre (X, Y) [mm]		Distance to Target [mm]	Position Uncertainty Radius [mm]	SNR [dB]	Thickness MAE [mm]
1	653	-32.01	-11.80	54.74	12.08	29.79	0.106
2	323	26.37	-1.46	95.54	12.29	39.39	0.111
3	385	-2.97	0.74	68.01	15.67	35.32	0.142
4	320	-35.31	-2.36	45.13	16.55	39.40	0.036
5	323	-8.41	11.22	58.55	7.40	39.53	0.079
6	393	-11.68	-5.53	64.17	7.63	39.96	0.155
7	308	16.55	4.18	84.35	9.81	36.63	0.073
8	260	12.05	-0.21	81.87	11.24	38.92	0.072
9	351	32.86	-1.39	101.58	10.11	31.58	0.148
10	242	-34.36	4.75	40.34	6.09	33.02	0.091
11	204	36.90	1.86	104.35	5.14	38.86	0.029

From the data of Table 3-7 it is clear that the UAV is able to make stable contact in a repeatable manner, satisfying the minimum probe coupling conditions in all interactions. These 11 measurements were completed in a period of under 143 s with minimal pilot intervention other than to approach and retreat from the sample surface. The system successfully captured data during each interaction; all maintained for at least 2 s.

As previously, SNR of an A-scan is estimated by the ratio of its overall maximum amplitude to the local maximum amplitude of the received signal before the first ultrasonic echo per (3-6). Results indicate all interactions grant mean SNR near 30 dB and above. Over all readings in this trial the system demonstrated thickness MAE of 0.095 mm, comparable to the previous experiment.

Considering precision of the transducer positioning, the mean location of each labelled contact point is within a maximum distance from the average across all points of 37 mm and 12 mm in the sample x and y axes, respectively. Over all trial contact point locations, the average distance to the target point is 73 mm, quantifying positioning accuracy. This accuracy, however, is expected to vary between repeated trials owing to the tolerance requested from the pilot, time taken to execute the motion, and other human factors, such as skill level and viewpoint.

3.4.3 Accessing the Underside of Overhanging Structures

Utilizing the ability of the Voliro aerial manipulator platform to actively re-orient its thrust relative to its body, a stable hover can be maintained in a number of pitch orientations. Further, it remains possible to exert force and perform environmental interaction in such poses, as in Figure 3-21. This enables application to numerous inspection scenarios concerning both vertical surfaces and the underside of overhanging structures.




Figure 3-21: The Voliro manipulator platform is able to enter and maintain stable contact with the underside of an overhanging surface. The inclination of the overhang is approximately 45° . Supports on the hidden side of the surface provide its rigidity, acting against force exerted by the UAV.

Following alignment of the UAV to the overhanging surface normal, its pose during repeated measurement acquisition is given in Figure 3-22. Therein, times of inspection are marked using the timestamps of captured A-scans with peak amplitude level higher than approximately 12 dB above the RMS noise floor and so indicative of successful deployment. Again, the probe criteria for acoustic coupling are fully satisfied within the limits stated in Table 3-2, meeting both orientation and interaction force requirements.




Figure 3-22: The UAV flightpath when inspecting the underside of an overhanging structure. Times when surface contact supported ultrasonic measurement are highlighted. Yaw, pitch and roll Euler angles relative to the ideal surface normal inspection pose are plotted under ZYX convention.




Figure 3-23: Error plot of thickness measurements from the underside of an overhanging surface versus reference geometry. Annotations give numerical thickness measurements and their standard deviation. Highlights demarcate times of contact measurements.

TABLE 3-8: TILTED MEASUREMENT ACCURACY STATISTICS

Number of Readings	Mean Absolute Error [mm]	5th Percentile Error [mm]	Mean Error [mm]	95th Percentile Error [mm]
1482	0.0924	-0.2249	-0.0536	0.1600

As before, a plot of the measured thickness error versus the reference sample geometry is provided in Figure 3-23 with statistical data describing the distribution of error magnitude given in Table 3-8.

The MAE of the ultrasound measurements versus reference geometry is comparable to previous results, indicating the performance is retained while operating in non-standard orientations on the underside of target structures. As before, the 5th and 95th percentile distribution points describe an asymmetrical distribution, favouring larger magnitude under-sizing errors. Overall measurement uncertainty has increased with the change in attitude but remains within usable limits.

This demonstrable and quantified capability to perform multi-orientation point thickness measurement with consistent accuracy and precision represents a significant development in the context of existing aerial ultrasonic NDE systems supporting measurement interaction in only one orientation [249], [273], or providing no quantitative discussion of measurement accuracy [258].

3.4.4 Dynamic Scanning Measurement Acquisition

A final trial is conducted as an assessment of system performance when capturing a stream of measurements, rolling along the sample surface. As such, extended area thickness maps may be constructed from a series of scanned lines. This modality also offers enhanced time efficiency and data density versus conducting serial discrete point measurements in a raster grid pattern. Here, assessment examines the capability to capture thickness features with a lateral spacing of 20 mm. Flight trajectory is set by the pilot while the flight controller maintains interaction force and orientation. Figure 3-24 illustrates this process through a series of timelapse images, showing completion of a scan back and forth along the surface of the stepped bar sample.




Figure 3-24: A top-down view of the UAV scanning across the stepped aluminium bar. Timelapse images were captured in 5 s intervals throughout the two passes across the bar sample. Image temporal progression from left-to-right then top-to-bottom shows the completed continuous contact interaction.

Figure 3-25 illustrates the measurements acquired via this dynamic scanning mode. This presents the ultrasonic data captured from motion in two contiguous passes across the bar sample surface, starting at the thin end before moving to the thicker sections and then returning back again. Figure 3-25(a) shows as a temporally encoded single-element B-scan, as introduced in Appendix B.3. Therein, sequential A-scan signals captured during the motion of the wheel probe form columns which are stacked horizontally to produce an image showing changes with time over the swept path. This differs from the more common array-based B-scan method which sequentially excites multiple ultrasonic elements to image the cross-sectional area beneath the sensor via beamforming and without physical motion. While such phased array NDE would offer finer resolution of small-scale defects like cracks and pitting, these results show the mechanical scanning approach to be sufficient for inspection processes at the coarser resolutions common to in situ asset thickness mapping. Further, thickness of a linear section is here recorded without carrying larger, more complex, phased-array instrumentation aboard the UAV, increasing maximum flight duration.

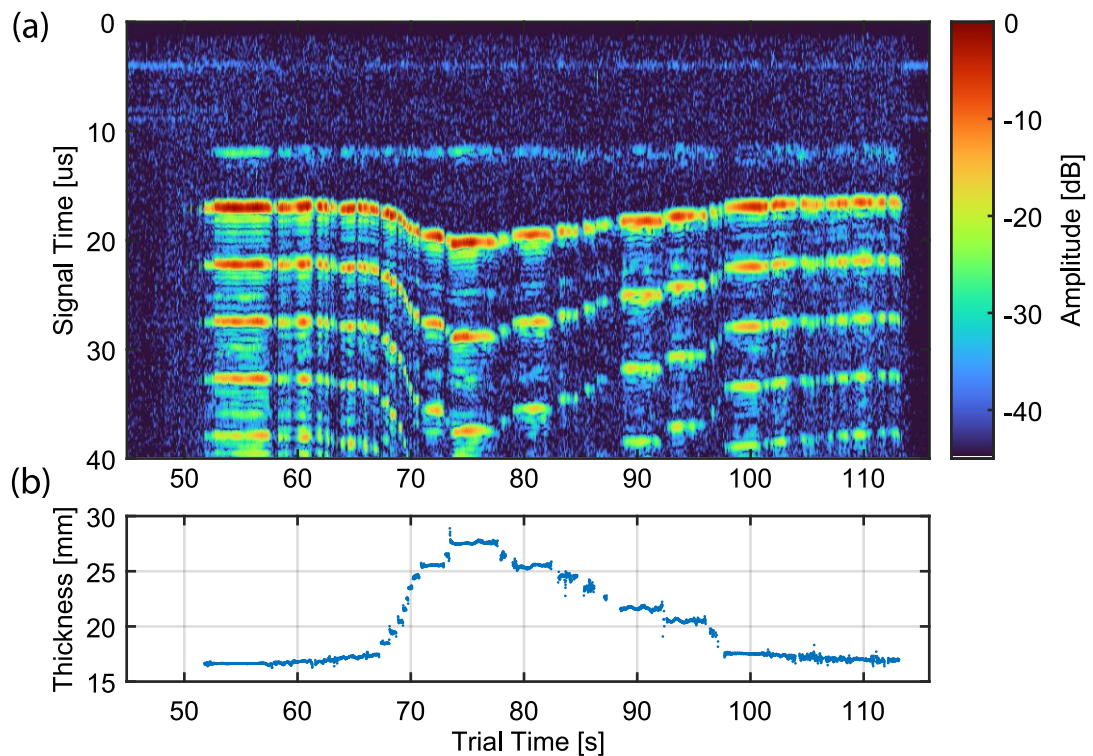


Figure 3-25: Ultrasonic data captured by rolling the wheel-probe along the bar x-axis and back again. (a) A temporally encoded ultrasonic B-scan normalised relative to maximum recorded signal amplitude. (b) A plot of thickness readings over time. Note all step features within the swept region are resolved.

Within Figure 3-25(a), the change in thickness of the sample may be directly observed in the vertical location of the first back-wall echo relative to the fainter front-wall echo appearing around $12 \mu\text{s}$ after transmission, as seen previously in Figure 3-8. This is more strongly confirmed by the varying ToF between higher amplitude subsequent back-wall echoes, which is used to extract numerical thickness via the method of Section 3.3.3. The variation in thickness over time is then plotted in Figure 3-25(b), indicating successful measurement of each the step feature in the swept path.

Further insight from the rolling inspection may be gathered by spatially encoding these data, as in Figure 3-26. A map of the rolling inspection path is first derived relative to the sample by projecting the measurement location of each ultrasound signal from the recorded UAV pose estimate per the method of Section 3.3.4, resulting in Figure 3-26(a).

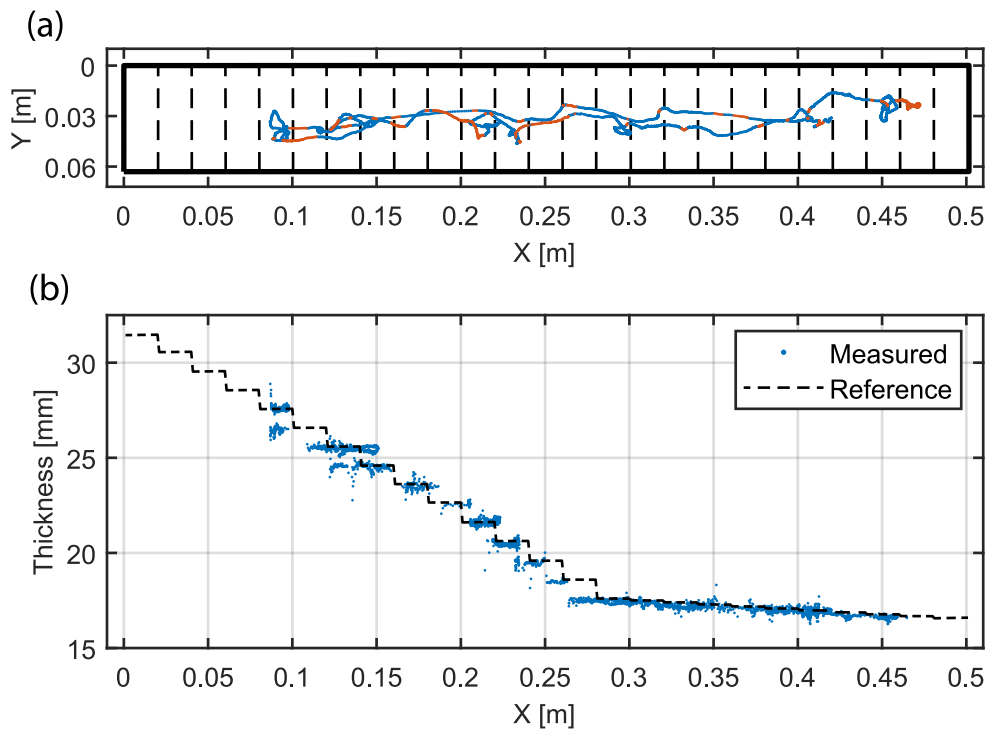


Figure 3-26: Aerial thickness measurement scans of the bar sample along its x-axis. (a) Points of inspection projected to the sample surface. Successful measurements are blue. Null or unsuccessful measurements are orange. (b) Position encoded bar thickness profile measurements. Each reading is a single blue point. The dashed line gives cross-sectional reference geometry. Note that these data exhibit spatial encoding uncertainty due to the effects discussed previously.

As such, the data of the two contiguous passes along the sample length, are shown to yield multiple successful measurements within the bounds of each thickness step. Remaining gaps where contiguous contact was maintained but fell short of the acoustic coupling requirements are smaller than the scale of reportable features for this trial. The largest single region without indicated coverage spanned a distance of 10.9 mm. It occurred near the leftmost point of Figure 3-26(a) as the scanning motion was halted and its direction changed before the return pass. The full scan path across the sample was completed in under 65 s, traveling a total distance of 0.71 m.

These measurement locations are then taken in combination with the data of Figure 3-25 to plot numerical thickness versus position in Figure 3-26(b), describing the dimensions of the sample cross-section. All thickness step features within the scanned region are again captured and visible in the depicted data, with 1 mm step thickness changes readily distinguishable. However, the smaller 0.1 mm steps are close to measurement accuracy limits established above so are challenging to discern.

As previously, error statistics are given in Table 3-9. These indicate some loss of accuracy during the dynamic process, as may be expected.

TABLE 3-9: SCANNING MEASUREMENT ACCURACY STATISTICS

Number of Readings	Mean Absolute Error [mm]	5th Percentile Error [mm]	Mean Error [mm]	95th Percentile Error [mm]
5197	0.2788	-1.0553	-0.1416	0.7389

These effects are attributed in part to measurement position uncertainty, as examined previously in Section 3.4.2. Visible as horizontal shift of some features in Figure 3-26(b), sample relative localization uncertainty of the magnitude given in Table 3-7 can see measurements wrongly assigned to neighbouring step features in the reference geometry. This is visible directly in the approach of the 5th and 95th percentile errors to the 1 mm step thickness change of the bar sample. However, within an industrial context, position uncertainty at this scale will only influence sizing and localization of smaller features. It will have minor impact on overall detectability or localization of larger scale flat pitting, or uniform and mesa-type corrosion features.

Through this contiguous scanning process, measurements were acquired with a success rate of 86.44 %. A recorded signal with peak amplitude falling below the 12 dB RMS noise-relative threshold for usable readings is considered a failed measurement. Here failure occurs where contact exists, as observable in the included video, but falls short of the minimum coupling force and orientation requirements for full acoustic coupling described in Table 3-2. This effect may also degrade SNR to marginally above the usability threshold and so cause the small number of outlier measurements visible in Figure 3-26(b).

The times at which coupling was successful are visible in the highlight overlay of Figure 3-27, plotting the world frame UAV flightpath. For optimal inspection performance the UAV should maintain the probe arm orientation along the sample surface normal with the direction in which the wheel-probe rolls parallel to the length of the sample.




Figure 3-27: The UAV flightpath while conducting a rolling scan measurement. Highlighted regions denote times when usable signals were acquired. Euler angle orientation is given relative to the ideal surface normal inspection pose.

Figure 3-27 also plots the angular error between this ideal pose and the craft orientation decomposed to ZYX convention Euler angles as before. Further, the estimated average speed of the probe during this motion, calculated using the derivative of its projected position and median filtered over a rolling period of 0.25 s to remove noise, was 8.5 mm/s but varied significantly. The probe both frequently stopped and exhibited sudden peaks in speed, up to a maximum of 115.5 mm/s, with standard deviation of 12.7 mm/s throughout the scan. Examination indicates that ineffectual measurements are common to sudden position changes or large angular offsets from surface normal, i.e. where pressure and orientation fall outside probe coupling limits. These effects are due to the nature of the tyre rolling resistance during dynamic motion and transition from the static case. Some insight into this complex behaviour may be gained via the free-body diagrams of Figure 3-28.

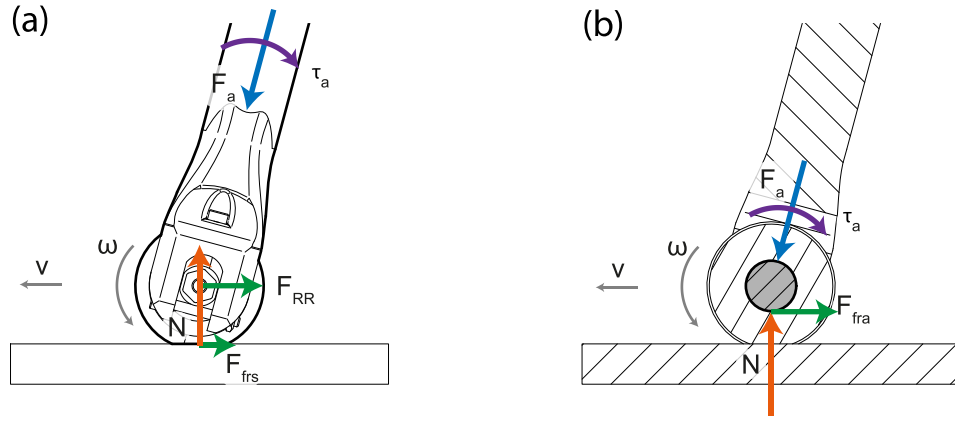


Figure 3-28: Freebody diagram of the wheel probe rolling scan interaction. (a) Considering the probe as a single unit. (b) Examining the effects acting upon the wheel axle in greater detail. Direction of travel and wheel rotation are indicated by vectors v and ω respectively.

Here the UAV applies force, F_a , and torque, τ_a , to enter and maintain a rolling dynamic equilibrium while satisfying the probe compression and orientation requirements determined in Section 3.2.3. Per Figure 3-28 (a), as in the static case, surface normal reaction force, N , is generated alongside static friction, F_{frs} , acting at the tyre-surface interface with magnitude up to a limit proportional to N per the friction coefficient, μ , i.e.

$$F_{frs} \leq \mu N \quad (3-15)$$

This friction force governs the rotational motion of the passive tyre about its axle, such that it rolls without slipping relative to the target surface, but has little effect on its translational motion. The rolling translation motion is more strongly influenced by rolling resistance force, F_{RR} . As this is a deformable wheel, F_{RR} may represent the combined action of many complex dynamic effects including hysteresis energy loss via continuous tyre compression and non-uniform pressure distributions across the contact area [336]. However, these effects are relatively consistent and pose minimal contribution to the observed discontinuities in motion.

Examination of forces at the probe axle via Figure 3-28 (b) better explains this behaviour. As before, compressive force acting through the tyre couples it to the piezoelectric elements in the axle, allowing ultrasonic transmission. However, in doing so the coupling force may also displace lubricant at the tyre-axle boundary, generating

a significant friction force, F_{fra} , that opposes relative motion between the tyre and axle. This effect increases in proportion to the applied coupling force per (3-15), introducing additional complexity. Furthermore, on commencement of motion, μ rapidly transitions from static friction to a weaker kinetic friction coefficient. Consequently, the wheel presents a reduced resistance to motion and accelerates rapidly upon starting to roll. Inertia of the UAV prevents its matching of this motion, instead arresting the wheel after a short distance. Repeated switching between these friction modes forms an irregular motion with low average speed along the target surface. The design of the probe prohibits introduction of bearings or other strategies to alleviate this friction effect without obstructing the ultrasound transmission pathway.

Nevertheless, the motion is shown as sufficient to support NDE measurement activity to the level illustrated in Figure 3-25 and Figure 3-26, successfully completing a rolling scan of the target cross-section. This result therefore serves to illustrate the potential for the scanning data capture mode once developed beyond this initial assessment.

3.5 Discussion

In consideration of the overall system, comparison is drawn to the state-of-the-art works listed in Table 3-1 to assess its relative capability for practical remote airborne ultrasonic NDE.

In this context, increased applicability to multiple in situ remote NDE applications is clearly demonstrated by the thrust vectoring, bi-axial tilting propeller, over-actuated tricopter architecture of the developed system. Whereas all works listed in Table 3-1 may conduct static, point thickness ultrasonic measurements in a single orientation, herein, NDE of multidirectional surfaces is performed through prolonged stable contact with both a vertical wall and overhanging structure, as in Section 3.4.1 and Section 3.4.3 respectively. Increased generality to the varying geometries found across high-value oil & gas, nuclear, and renewable energy generation structures is directly exhibited versus systems limited to single orientation interaction modes: contacting

vertically planar walls [249], [251], the underside of horizontally planar ceilings [250], or depositing a sensor package atop horizontal plates and pipes [273].

Full characterisation of the interaction repeatability and measurement positional uncertainty, as in Section 3.4.2, further highlights the utility of over-actuation to decouple attitude and translation control and enable force omnidirectionality. Here, as in the multidirectional thrust AeroX platform [258], wrench is generated in all directions about the body for interaction stabilisation from a static hover, reducing sensor motion incurred during force redirection in under-actuated quadrotors with unidirectional thrust, [250], [251], [273]. Requirements for mechanical stabilisation structures such as single DoF compliant effectors [249]–[251], or additional dedicated propellers which are otherwise unused in free flight [249] are then relaxed. There is also no reliance on material specific adhesion forces such as magnetism [273], enabling application to modern composite structures.

Further, over-actuation by dynamic vectoring of the propellers, as used here, removes both the counteracting internal wrench components found when using the net effect from multidirectional thrust superposition [258], and the resultant penalty to energy efficiency. A greater proportion of the total thrust is thereby available outside of the vertical direction, enabling pose omnidirectionality in stable flight with adjustable non-horizontal attitude. Where [258] makes omnidirectional contact with the target from a consistent horizontal pose using a 6 DoF actuated arm that rotates about the centre of mass, the system herein may directly align its body to the surface normal direction. Contact inspection is then reliably performed with a simple rigid effector, as shown throughout Section 3.4, permitting a UAV platform with significantly reduced mass, size and complexity and so supporting an increased ability for efficient remote inspection within crowded industrial airspace.

Direct quantitative comparison of ultrasonic measurement performance is limited by a lack of common procedures and metrics among existing publications. Of the works in Table 3-1, only [249], [250] and [251] quantify inspection measurement accuracy versus a known dimensioned target. Against a vertical aluminium plate with thickness ranging from 9 mm to 15 mm, [251] reports a mean absolute error of 0.12 mm, whereas [249] shows consistent error of -0.6 mm versus a 6.0 mm thick plate of

unspecified metal, attributed to sensor calibration. Performance validation in [250] reports an error of +0.46 mm, but the slower speed of sound in the 5.0 mm thick, horizontally mounted, acrylic sample limits comparability. Capability for both point thickness and scanning measurement is indicated in [258], but the authors present no quantitative ultrasonic data for comparison.

Addressing this, a full characterization of ultrasonic thickness measurement accuracy versus known dimensioned aluminium samples is presented herein using a well-defined procedure. Data obtained in repeated static point and rolling scan measurements of vertical aluminium samples using a dry-coupling ultrasonic wheel probe, show mean absolute error below 0.1 mm and 0.3 mm respectively. Whereas previous systems have required the manual application of ultrasonic couplant gel in an additional preparatory step [250], [251], or instead carry an onboard gel reservoir and dispensation mechanism [249], [258], this system presents a new mode of aerial ultrasonic inspection, improving process efficiency and aiding utilization where couplant gel may be contaminative. This inspection modality is also compatible with electrically non-conductive targets, increasing applicability over EMAT solutions [273]. Additionally, this work considers an application-focused quantification of both the repeat interaction positioning accuracy versus a target point and the propagation of combined uncertainty from UAV position and attitude estimates to the inferred, sample relative, measurement location. Such aspects are critical during extended timescale asset monitoring requiring repeat probe positioning at a known, consistent location. The results of Section 3.4 thus provide numerous practical benchmarks highlighting the utility of this combination of an advanced aerial robotic platform with a dry-coupling ultrasonic transducer for in situ NDE process automation.

3.6 Summary

Building on recently published works, operating principles for an aerial inspection platform integrating an over-actuated UAV with a 5 MHz dry-coupling ultrasonic wheel probe to perform contact based volumetric inspection are presented as a means to advance airborne NDE processes.

An application focused assessment of system capability shows that the UAV may readily perform point interaction with the test samples, meeting and maintaining the probe coupling and pressure requirements. Empirical trials demonstrate successful measurement acquisition in numerous orientations.

In point thickness measurements taken of vertical surfaces and on the underside of overhanging structures, mean absolute error below 0.1 mm is observed when compared to the sample reference dimensions. This indicates a level of performance comparable to some hand-held commercial thickness measurement equipment.

Lastly, utilization of the wheel probe is assessed for continuous thickness mapping along a linear path. This demonstrated the recording of the cross-sectional dimensions of a precision manufactured aluminium test piece. Stepped changes in thickness of 1.00 mm spaced at 20.01 mm increments along the sample length are clearly visible in the processed inspection data, despite small gaps in coverage of under 10.9 mm where the dynamics of the rolling motion cause a temporary loss of coupling.

Ultimately, the UAV aerial inspection system is found to be successful in its application of ultrasonic NDE and is presented with quantified, application contextualised, uncertainty bounds. This system therefore represents a means of improving the inspection process efficiency and operator safety in hazardous conditions common to industrial environments across the energy sector.

CHAPTER 4

A Hybrid Unmanned Aerial Vehicle – Crawler for Airborne Pipe Inspection

4.1 Introduction

As established in Section 1.1, corrosion of industrial infrastructure is a global problem, occurring through numerous mechanisms intrinsic to both asset function and operating environment. In many cases, it cannot be entirely prevented and must be actively managed to ensure continued safe and efficient operation of the vulnerable structure. As such, effective Operations and Maintenance (O&M) entails substantial NDE campaigns, regularly monitoring asset degradation relative to limits established by specific industrial standards [56], [66], [70].

Assessment of the most common structures across the energy generation sector, including pipework, storage tanks, and pressure vessels, comprises an integral part of this activity, preventing critical failures [300] but often presenting significant access challenges including work at height and in confined spaces due to surrounding obstructions. At facility scale, where insulated pipework can stretch for over 98 km in a mid-sized chemical processing plant or over 572 km in a mid-sized refinery [337], manual inspection costs of USD \$25 000 per 10 m elevated section quickly compound, with scaffolding comprising up to 80 % of the cost as estimated in 2019 [338].

To alleviate these recurring operating expenses, strong industrial desire is expressed for automation and robotic systems capable of matching or improving the structural insight of existing processes and further supporting predictive maintenance strategies.

4.1.1 Pipe Exterior Inspection Robots

Exterior surface robotic inspection systems are commonly adopted for in-service NDE processes and take many forms. These include both surface crawler robots and UAV platforms, with further subvariants created to address specific process requirements. An overview of existing approaches is given below, including examples from recent literature as reviewed in greater detail in Section 2.5, to aid assessment of their relative performance characteristics.

Comparing these systems, the platform stability offered by C-clamp style crawlers [267], [268] is unsurpassed. These leverage rigid mechanical clamping such that the target is near fully enclosed, entirely supporting the vehicle and ensuring its Centre of Mass (CoM) is located inside the pipe diameter. Multiple strong reactive contact forces allow these crawlers to utilise simple motion control algorithms and reliably deploy NDE sensors with high precision in immediate surface proximity. Accordingly, sensors have extremely short reach from the vehicle body, eliminating the negative effects of lever-arm disturbance amplification on NDE data. However, these full-enclosure crawlers cannot autonomously disengage and re-enter contact to pass circumferential obstacles like flanges, supports, or sharp bends. They may only pass small radial obstacles, such as pipe junctions or valves, by aligning the obstacle with the small gap in their clamping mechanism used to enable initial deployment. This has the undesirable effect of rendering certain areas of the target structure inaccessible without costly manual redeployment. Clamp geometry also restricts maximum enclosed asset diameter, limiting generality of a single platform.

A walking robot with two magnetic feet has been proposed as a means to step over these obstacles and increase diameter applicability [269], but such a strategy precludes application to insulated, composite or otherwise non-magnetic structures and entails more complex locomotion control. Increased CoM separation from the pipe centre may also minorly degrade stability.

Flight presents a more generic access solution, briefly exiting contact to bypass obstructions and allowing deployment without direct manual presence. Initial feasibility of such an approach has been shown by deposition of a discrete magnetic

sensor package atop a pipe using a standard unidirectional thrust quadrotor [273]. In this case, however, probe deployment success rate was limited to below 65 % due to near-surface aerodynamic disturbances [339], commonly restricting sensor placement accuracy and defect detection probability in underactuated UAVs [251].

As considered in Section 2.5.2, the increased dynamic capability of MDT [258], [260] and VT [259], [263] platforms permits direct counteraction of these effects to substantially improve probe positioning accuracy and stability. Capability for omnidirectional interaction also broadens application generality to include the full pipe circumference [258], [259], [263]. However, in such platforms the NDE probe extends far from the vehicle CoM to avoid rotor collisions with the target surface, entailing long-reach deployment via rigid arm [259], [263] or serial manipulator [258], [260] and amplifying pose variation at the effector tip. Lack of intrinsic support also necessitates strong interaction feedback pathways and complex control algorithms to facilitate contiguous contact scanning and repositioning, as examined in Chapter 3. Moreover, the additional actuation hardware can result in relatively large UAVs requiring a 1-2 m target clearance radius which could pose problems in crowded industrial facilities.

By contrast, smaller, simpler UAVs have shown stable short-reach contact using increased support from the target structure. In work by David & Zarrouk [270], a wall-crawling UAV is supported via friction from surface-normal thrust created by reversing its propellers, but high energy usage precludes large vehicles or low-friction surfaces. A mix of friction and thrust-based weight support may relax these issues, as shown by Myeong & Myung [271], but neither this nor [270] consider NDE or adhesion to curved surfaces.

In this context, recent UAV perching strategies have shown significant improvement, minimising energy use and increasing sensor stability over aerial manipulators. Commonly, the UAV suspends itself beneath the target structure [274], [275], or rests atop it [276]–[278] for extended duration deployment under passive support. Many use a gripper, either rigid [274], [275], [277] or conformable [276], to provide additional adhesion and further reject disturbances. Whereas some may exhibit locomotion using an inchworm mechanism [276] or driven wheels [277], [278] to

dynamically investigate suspected defects, others cannot move after first perching [274], [275], restricting potential NDE applications. Precise NDE transducer positioning has also been enabled via onboard robotic arm while the vehicle body remains in a position of passive support atop the pipe surface [277], [278]. However, arm mass and reach limit scalability of these systems to large diameter assets as they lack suitable adhesion mechanisms to support their weight in positions other than atop the pipe.

A novel, more extensible, paradigm is thus for the entire UAV to crawl about the asset surface. Preliminary works [279], [280] address this goal via control algorithms for non-flight-capable mock hardware. Zhao *et al* [279] use Non-Linear Model Predictive Control (NMPC) of wheel torque to stabilise a cart in the top $\pm 20^\circ$ of a pipe but omit any adhesion force necessary for circumnavigation. Nekoo *et al* [280] examine a benchtop test rig with variable-pitch propellers, regulating radial and angular position via the State-Dependent Riccati Equation (SDRE), but do not consider beneficial contact forces and so enforce a perpetual hovering condition susceptible to aerodynamic disturbance.

4.1.2 System Conceptualisation

A hybrid UAV-crawler vehicle is therefore identified as a novel and highly promising means to both combine the desirable aspects and resolve the shortcomings of the pipe-exterior NDE systems examined in Table 4-1.

TABLE 4-1: COMPARATIVE OVERVIEW OF PIPE-EXTERIOR INSPECTION ROBOTS

Ref.	[267]	[268]	[269]	[273]	[258]	[276]	[278]	Herein
Vehicle Category	Crawler	Crawler	Crawler	UAV	UAV	UAV	UAV	UAV-Crawler
Vehicle Sub-Category ^a	Clamped	Inch-worm	Walking	Uni	AM	Perch	Perch + RA	Hybrid
Contact Type ^b	Encircle	Encircle	Point	Point	Arm	Body	Body + Arm	Body
Access ^c	MP	MP	MP	Flight	Flight	Flight	Flight	Flight
Adhesion	Clamp	Clamp	Magnetic	Magnetic	Thrust	Soft Grip	Gravity	Thrust
Control Complexity	Low	Low	Mid	Low	High	Mid	High	Mid
Maximum Diameter [mm]	285	75	∞	∞	∞	300	250	∞
Pipe Coverage	Full	Full	Full	Top	Full	Top	Full	Full
Obstacle Bypass ^d	Radial	Radial	Radial + Circ.	Radial + Circ.	Radial + Circ.	Radial + Circ.	Radial + Circ.	Radial + Circ.
Reach ^e	Very Short	Very Short	Short	Short	Long	Short	Mid	Short
Probe Stability	Very High	Very High	High	Low	Mid	High	Mid	High

Key:

^a Defining platform architectural features i.e. *Uni*: Conventional unidirectional thrust layout UAV, *AM*: Aerial Manipulator platform, *RA*: Embedded Robotic Arm.

^b Nature of interaction with the target i.e. *Point*: Single point of contact. *Body*: Multiple points of contact distributed about the vehicle body. *Encircle*: Vehicle encircles the majority of the pipe circumference. *Arm*: Single point of contact at the end of an effector.

^c Means of initial vehicle deployment i.e. *MD*: Manual Placement on asset surface.

^d Types of obstacles that may be bypassed i.e. *Radial*: Obstacles extending radially from the surface such as pipe junctions and supports. *Circ.*: Obstacles blocking the full pipe circumference.

^e Relative distance from the vehicle point of support to the NDE probe.

Employing an airborne access method, such a vehicle may eliminate any costs due to initial manual deployment and, provided sufficient airspace clearance exists, completely bypass any surface obstacles preventing crawler access. By engaging in full-body pipe contact during inspection, a hybrid vehicle may greatly improve stability and disturbance rejection versus free-flying systems [258], [273], providing means to capture high-resolution surface images with minimal standoff or motion blur as a remote visual inspection platform. Crawling around the full circumference, it may also detect fine-scale damage, leaks, and exposed cracks without asset diameter

restrictions encountered by systems enclosing the target [267], [268], [276] or relying on specialised embedded manipulators to reach around the surface from a perching position [278]. Further, an ability to move relative to the surface without exiting contact would present significant utility to NDE applications, enabling high-accuracy probe placement and fine-scale repositioning to repeatedly access and monitor known defect locations in satisfaction of standard pipe inspection practice [66]. This could represent a significant advancement over existing free-flying systems with high sensitivity to aerodynamic disturbances [258], [273] while retaining their general flexibility to realise potential operating cost benefits for critical infrastructure across the global industrial sector.

4.1.3 Chapter Structure

The remainder of this chapter is concerned with the development of such a hybrid vehicle and is arranged as follows. Section 4.2 details the UAV-crawler system and models its environmental interaction. Section 4.3 considers optimal thrust support behaviour within the context of this model, while Section 4.4 details the control of vehicle flight, target interaction, and surface crawling. A description of the empirical methodology and supporting information used to assess these aspects follows in Section 4.5. Results of practical testing are then presented in Section 4.6 and discussed in Section 4.7, before the chapter summary of Section 4.8.

4.2 System Modelling & Prototype Vehicle

To develop a hybrid UAV-crawler vehicle, a rigid-body model for its thrust actuation and interaction with horizontal-axis industrial pipework or other cylindrical assets is established. This is introduced in the context of the constructed prototype but is generically extensible to vehicles of a similar archetype.

Herein, as previously, a vector or matrix is denoted by bold typeface. The signed magnitude scalar representation of a vector uses the same symbol in regular typeface. A leading superscript indicates the reference coordinate frame of a vector. Parameters with trailing numeral subscripts refer to the rotor of that index or i in the generic case. Other notation is defined as it is used.

Note that three spatial coordinate frames are used to establish the full rigid body model:

- $\{W\}$, the global world frame. This takes the standard North-East-Down sense [186] and is used to describe the asset and vehicle position relative to their environment.
- $\{B\}$, the body-relative frame. This employs a similar Front-Right-Down (FRD) coordinate system and is affixed to the vehicle with origin at its centre of mass (CoM) and is used to examine forces acting on the vehicle.
- $\{P\}$, the pipe-relative coordinate frame. This is defined with an origin in the pipe centre, an x-axis pointing along the pipe length, and an upwards pointing z-axis.

4.2.1 Prototype Vehicle

In consideration of the relative advantages of flight and surface crawling, the hybrid UAV-crawler vehicle is introduced, as in Figure 4-1.




Figure 4-1: A diagram of the prototype hybrid vehicle. Annotations show the body coordinate frame, $\{B\}$, and the unit vectors defining the positive thrust direction of each propeller. Red vectors spin counter-clockwise: blue spin clockwise. Bounding box dimensions are given in millimetres.

Initially landing atop a horizontal pipe, or similar cylindrical asset, this vehicle may then circumnavigate the surface in multi-point, short-reach, contact. Crawling under

thrust-based support, its design applicability extends to large, non-ferromagnetic assets and those where dirt or surface texture prevents vacuum adhesion.

Accordingly, its six MDT propellers grant omnidirectional net force in the $\{B\}$ y-z plane for both flight and crawling support. Each uses a fixed-pitch reversible HQprop 3D-5x3.5x3 propeller with a symmetric aerofoil profile. These have equal thrust efficacy in both spin directions, are driven by iFlight Xing 2207 motors, and powered by a 22.2 V (6S) 3000 mAh Li-Po battery. Adaptive commutation within the Aikon AK32 Electronic Speed Controller (ESC) firmware [340] enables run-time spin reversal without the additional hardware and complexity of variable pitch rotors [280], or the VT approach used in some aerial manipulator UAVs [259], [263].

Propeller layout, as in Table 4-2, is set following iterative design feasibility validation via the method detailed in Section 4.2.3, below, to ensure circumferential stability without actuator saturation. Thus, the main lift rotors, 1-4, are cambered to aid lateral thrust generation. Incline of the main lateral thrust rotors, 5 and 6, also reduces airflow obstruction by the vehicle body and ensures propeller tip clearance above the contact plane. Note also that the axes of these propellers are oriented such that none possess a component acting along the $\{B\}$ x-axis and so the array cannot generate thrust in that direction. This design choice removes counteracting thrust components, improving efficiency, and simplifies the system during initial examination of the contact-crawling approach. Mounted as described to a light, 3D-printed, airframe, the propellers may generate multidirectional thrust in the $\{B\}$ y-z plane, fully supporting the 1.953 kg total vehicle mass in arbitrary position about a horizontal axis pipe. Future objectives requiring omnidirectional thrust can then be met by reorienting propellers accordingly.

TABLE 4-2: UAV-CRAWLER HYBRID VEHICLE ROTOR LAYOUT

#	1	2	3	4	5	6
p_x [mm]	110	-110	110	-110	0	0
p_y [mm]	100	-100	-100	100	160	-160
p_z [mm]	-10	-10	-10	-10	-10	-10
ϕ [°]	-25	25	25	-25	75	-75

Key: p_x, p_y, p_z , Rotor position in $\{B\}$ axes. ϕ , Angle about $\{B\}$ x-axis of rotor axis from vertical.

Target contact occurs at six points, symmetrical in the $\{B\}$ y-z plane. At the front and rear, two FS5106R continuous rotation servo wheels with high-grip polyurethane elastomer tires provide surface locomotion under adaptive thrust support. Two pairs of rigid legs are also added, providing a further four points of contact to constrain vehicle yaw and increase stability. These are interchangeable: designed for a known approximate target diameter and trivially adapted to fit ASME standard pipes [341], or other targets above 160 mm diameter. The small vehicle form factor extends at most 180 mm radially beyond the asset surface, requiring far less clearance than hovering aerial manipulators [258]–[260] for multi-modal contact NDE.

Here, the vehicle is fitted with payload hardware for immediate proximity visual NDE. This includes a FLIR CM3-U3-50S5C, 5 MP, machine vision camera, and M0824-MPW2 lens, mounted at 55 mm surface standoff. These yield resolution of approximately $28 \mu\text{m}/\text{px}$ over a 56 mm by 69 mm image area. An embedded LED array allows image capture in complete darkness, obliquely lighting the target to enhance texture feature visibility. An onboard Raspberry Pi 4B computer supports real-time image streaming at over 15 Hz via Wi-Fi and internal high-resolution storage for later analysis and reporting. Its Robotic Operating System (ROS) integration with the Pixhawk4 flight control unit permits remote UAV telemetry and camera parameter changes for flexible NDE deployment.

4.2.2 Rigid-Body Interaction Model

A mathematical model describing this vehicle and its interaction with a target pipe or cylindrical asset is developed as depicted in Figure 4-2. Note that, per mechanical constraint by its legs, the vehicle angular position about the surface of the pipe is fully described in $\{P\}$ by clockface angle, $\theta \in [-180^\circ, +180^\circ]$. This is zero when atop the pipe and positive with clockwise rotation as pictured in Figure 4-2(a), conveniently matching UAV roll angle.

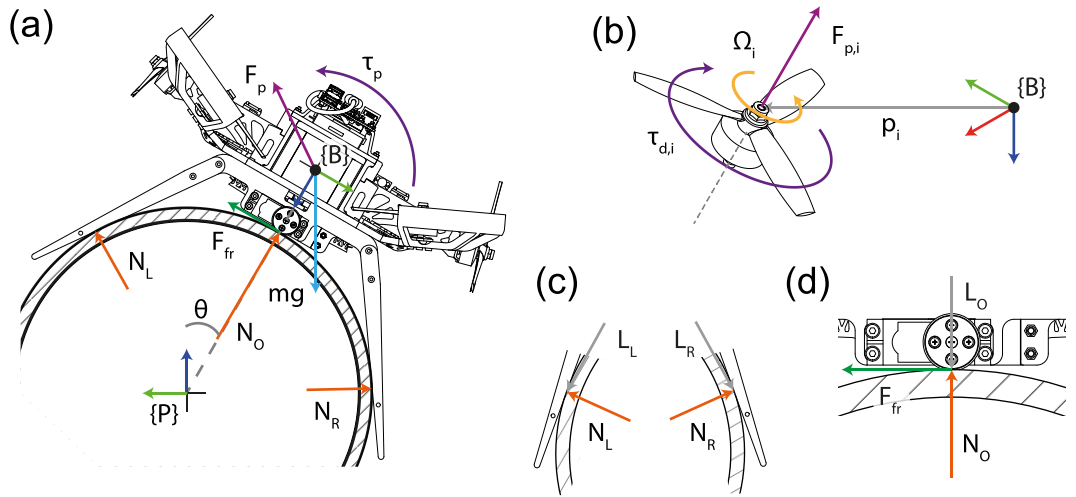


Figure 4-2: Free-body diagram showing forces and torques acting on the hybrid UAV-crawler vehicle. (a) Overview. (b) Propellers. (c) Legs. (d) Central wheels. Grey vectors indicate force application points relative to $\{B\}$.

Accordingly, the wrench components acting on the vehicle are examined in $\{B\}$ to develop a vehicle relative interaction model that will later serve as the foundation for its onboard control. Exploiting system symmetry in the $\{B\}$ y-z plane, this interaction can be reduced to a 2D model in the case of the horizontal-axis cylindrical targets examined here, as depicted in Figure 4-2(a). Contact forces are then approximately equal between the vehicle front and rear, and so are combined in single parameters.

Thus, the wheels equally split the normal contact and Coulomb friction force vectors, N_O and F_{fr} , respectively. Both F_{fr} and N_O act at L_O , directed in the $\{B\}$ y-axis and negative z-axis respectively as depicted in Figure 4-2(d). The magnitude of the friction vector at the wheel is then limited by the static friction coefficient, μ , as below. Note that this condition remains applicable for inclined pipes, in which case friction acts in the $\{B\}$ x-y plane to support the component of the vehicle weight directed along the pipe and so provides a tolerance to small deviations from horizontal target orientation.

$$|F_{fr}| \leq \mu N_O \quad (4-1)$$

As in Figure 4-2(c), the remaining normal contact forces of the left and right leg pairs, N_L and N_R , act at points L_L and L_R with direction unit vectors \hat{u}_L and \hat{u}_R , respectively. Friction effects at these contact points are neglected for simplicity owing to the low-

friction interaction between the plastic legs and target surface material relative to that of the rubberised tyres.

Within the model for horizontal pipe interaction, the vehicle weight vector, $m\mathbf{g}(\theta)$, is expressed as a function of θ with constant magnitude, redirecting the product of the acceleration due to gravity, g , and vehicle mass, m , according to clockface position.

$$m\mathbf{g}(\theta) = mg[0 \quad \sin \theta \quad \cos \theta]^T \quad (4-2)$$

Expanding the model framework for a UDT craft established in Section 2.4, collective propeller behaviour is then set by their net force, \mathbf{F}_p , and torque, $\boldsymbol{\tau}_p$, vectors, acting through the centre of mass and causing rotation about it, respectively. As shown in Figure 4-2(b), each propeller exerts static thrust, $\mathbf{F}_{p,i}$, and drag torque, $\boldsymbol{\tau}_{d,i}$, varying as a function of its rotation speed, Ω_i , and lumped thrust, C_T , or drag, C_D , coefficients,

$$\mathbf{F}_{p,i} = C_T \Omega_i^2 \cdot \hat{\mathbf{a}}_i, \quad \boldsymbol{\tau}_{d,i} = s_i C_D \Omega_i^2 \cdot \hat{\mathbf{a}}_i \quad (4-3)$$

These coefficients are determined by the propeller aerofoil and vary with air density and rotor radius, as described previously. Here, thrust acts on the airframe at point \mathbf{p}_i , with direction parallel to the rotor axis unit vector, $\hat{\mathbf{a}}_i$. Drag torque acts about axis $\hat{\mathbf{a}}_i$, in the spin direction opposing Ω_i , denoted $s_i \in \{-1, +1\}$ per the right-hand rule. Thus $\mathbf{F}_{p,i}$ and $\boldsymbol{\tau}_{d,i}$ may be related via a scalar factor, $K_M = C_D/C_T$, such that

$$\boldsymbol{\tau}_{d,i} = s_i K_M \mathbf{F}_{p,i} \quad (4-4)$$

The total torque exerted by each propeller, $\boldsymbol{\tau}_{p,i}$, is then the summation of $\boldsymbol{\tau}_{d,i}$ and the thrust moment, computed via the vector cross-product $\mathbf{p}_i \times \mathbf{F}_{p,i}$, with the result expressible as a function of thrust scalar, $F_{p,i}$, and a fixed direction vector sum, $\boldsymbol{\tau}_{\Sigma,i}$.

$$\boldsymbol{\tau}_{p,i} = F_{p,i} \boldsymbol{\tau}_{\Sigma,i} \quad (4-5)$$

$$\boldsymbol{\tau}_{\Sigma,i} = (\mathbf{p}_i \times \hat{\mathbf{a}}_i) + s_i K_M \hat{\mathbf{a}}_i \quad (4-6)$$

Consequently, net propeller wrench $[\mathbf{F}_p, \boldsymbol{\tau}_p]^T$ is the product of the collective effectiveness matrix, \mathbf{A}_p , and the column vector of all propeller force scalars, \mathbf{F}_q .

$$\begin{bmatrix} \mathbf{F}_p \\ \boldsymbol{\tau}_p \end{bmatrix} = \begin{bmatrix} \sum \mathbf{F}_{p,i} \\ \sum \boldsymbol{\tau}_{p,i} \end{bmatrix} = \mathbf{A}_p \mathbf{F}_q \quad (4-7)$$

Where n is the number of rotors, and \mathbf{A}_p is formed by the vertical concatenation of the force and torque effectiveness matrices, \mathbf{A}_F and \mathbf{A}_τ , respectively.

$$\mathbf{A}_p = \begin{bmatrix} \mathbf{A}_F \\ \mathbf{A}_\tau \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{a}}_1 & \hat{\mathbf{a}}_2 & \dots & \hat{\mathbf{a}}_n \\ \boldsymbol{\tau}_{\Sigma,1} & \boldsymbol{\tau}_{\Sigma,2} & \dots & \boldsymbol{\tau}_{\Sigma,n} \end{bmatrix} \quad (4-8)$$

$$\mathbf{F}_q = [F_{p,1} \quad F_{p,2} \quad \dots \quad F_{p,n}]^T \quad (4-9)$$

Note that this expression of collective propeller wrench presents key differences from that given previously in (2-23). In this case, formation of \mathbf{A}_F using the normalised propeller direction vectors, $\hat{\mathbf{a}}_i$, allows \mathbf{A}_p to express both 3D torque and 3D force behaviour so as to model fully and over-actuated UAVs, as is necessary for the MDT design of the hybrid vehicle. Further, expression as a function of individual propeller thrusts enables the use of non-linear models describing bi-directional actuation and explicit incorporation of asymmetric thrust limits, as examined in Section 4.4.3.

The collective propeller wrench may then be included in consideration of the net wrench across all interaction forces and torques to complete the interaction system model. Accordingly, summation gives a single net force acting through the centre of mass and a single net torque acting about it, as below.

$$\sum {}^B \mathbf{F} = \mathbf{N}_O + \mathbf{F}_{fr} + \mathbf{N}_L + \mathbf{N}_R + m\mathbf{g}(\theta) + \mathbf{F}_p \quad (4-10)$$

$$\sum {}^B \boldsymbol{\tau} = (\mathbf{L}_O \times \mathbf{N}_O) + (\mathbf{L}_O \times \mathbf{F}_{fr}) + (\mathbf{L}_L \times \mathbf{N}_L) + (\mathbf{L}_R \times \mathbf{N}_R) + \boldsymbol{\tau}_p \quad (4-11)$$

Where system dynamics are of interest, this may be extended per a Newton-Euler formulism [342], expressing in $\{B\}$ the motion resulting from the effect of all forces and torques acting on the rigid body, as discussed in Section 2.4.3.

$$\begin{bmatrix} \sum {}^B \mathbf{F} \\ \sum {}^B \boldsymbol{\tau} \end{bmatrix} = \begin{bmatrix} m\mathbf{I} & 0 \\ 0 & \mathbf{J} \end{bmatrix} \begin{bmatrix} {}^B \dot{\mathbf{v}} \\ {}^B \dot{\boldsymbol{\omega}} \end{bmatrix} + \begin{bmatrix} {}^B \boldsymbol{\omega} \times m {}^B \mathbf{v} \\ {}^B \boldsymbol{\omega} \times \mathbf{J} {}^B \boldsymbol{\omega} \end{bmatrix} \quad (4-12)$$

Where $\mathbf{I} \in \mathbb{R}^{3 \times 3}$ represents the identity matrix, \mathbf{J} is the vehicle inertia tensor expressed in $\{B\}$, ${}^B \mathbf{v}$ and ${}^B \boldsymbol{\omega}$ are the translational and angular velocity, and their derivatives, ${}^B \dot{\mathbf{v}}$ and ${}^B \dot{\boldsymbol{\omega}}$, denote the respective accelerations.

For the purposes of vehicle support about the target circumference, the interaction may be considered quasi-static with net force-torque vectors thus equal to the zero vector.

$$\begin{bmatrix} \sum {}^B \mathbf{F} \\ \sum {}^B \boldsymbol{\tau} \end{bmatrix} = \mathbf{0} \quad (4-13)$$

As the slow motion common to NDE activity requires minimal centripetal force relative to the weight support force, this condition may also be applied in an approximation to simplify vehicle surface crawling behaviour during inspection.

4.2.3 Vehicle Design Feasibility

To construct a hybrid vehicle capable of circumferential contact inspection across the range of pipes, tanks and pressure vessels used in the energy sector, a condition to qualify generic design feasibility must be established. This will primarily assess the ability of the craft to support its weight around a surface of non-specific diameter.

4.2.3.1 Minimum Supporting Wrench Requirements

As such, an expression for the minimum supporting wrench required to maintain arbitrary circumferential pose is first developed using the interaction model above. Recognizing that the vehicle legs are subject to change with asset diameters, and that they provide contact forces beneficial to support, a stringent condition for generic

feasibility omits their effects and thereby guarantees feasibility across all diameters. From (4-10), the minimum supporting force vector required at a given clockface angle is then found by minimizing the magnitude of the residual force vector it must support, i.e.

$$\underset{N_o, F_{fr}}{\operatorname{argmin}} \quad \|\mathbf{N}_o + \mathbf{F}_{fr} + m\mathbf{g}(\theta)\| \quad (4-14)$$

s.t

$$N_o \geq N_{min}$$

$$-\mu N_o \leq F_{fr} \leq +\mu N_o$$

In addition to the limits of static friction expressed in (4-1), this also imposes a non-zero minimum permissible wheel contact force at any point around the target circumference, N_{min} . Whereas the minimum supporting force may be further reduced by the trivial condition of $N_{min} = 0$, under such a criterion the vehicle is in essence hovering in free-flight immediately adjacent to the target surface and therefore very sensitive to disturbances. Here, N_{min} is thus set at 10 N, ensuring that the vehicle exerts strong radial contact force, aiding position stability and robustness to modelling uncertainty, while also representing the compressive coupling requirements of some NDE transducers. A safety factor is also applied to the static friction coefficient between the tyres and the pipe, μ , setting conservative limits for its use to stabilise the interaction. While a typical value of 0.20 may be appropriate between steel and plasticised materials similar to the tyre [343], here the limit is thus set with μ of 0.05. At N_{min} of 10 N, this will allow the vehicle to use static friction forces for support up to 0.5 N. A reserve of 1.5 N versus the theoretical value, around 7.8 % of vehicle weight, then provides tolerance for model uncertainty and force generation error in the $\{B\}$ y-axis, also granting robustness against the effects of small pipe inclines, in situ lubricants like particulate dust, and moisture on circumferential position stability.

Substitution of the N_o and F_{fr} solving (4-14) into (4-10) then yields the minimum supporting force vector as a function of clockface angle, $\mathbf{F}_s(\theta)$, while their substitution in (4-11) gives the accompanying torque vector required for static equilibrium, $\boldsymbol{\tau}_s(\theta)$.

Together, these define the required minimum support wrench, $\mathbf{W}_s(\theta)$.

$$\mathbf{W}_s(\theta) = \begin{bmatrix} \mathbf{F}_s(\theta) \\ \boldsymbol{\tau}_s(\theta) \end{bmatrix} \quad (4-15)$$

Generic feasibility is thus achieved if, at all clockface angles, the required support wrench lies within the set of feasible net propeller wrench, Γ_p . Formally,

$$\mathbf{W}_s(\theta) \in \Gamma_p \quad \forall \theta \quad (4-16)$$

The design objective is therefore to define an array of propellers in terms of their position, orientation, and individual thrust generation capability such that this condition is satisfied by the candidate vehicle. However, this is a complex and highly multidimensional problem with no unique solution, also having several interlinked practical considerations regarding vehicle mass, airframe structure, and power usage. An iterative design approach is therefore taken, proposing a candidate vehicle before evaluating its feasibility and revising accordingly until the MDT array design is satisfactory and can be fully implemented as a prototype vehicle.

Accordingly, means to determine Γ_p and evaluate (4-16) is required. This is conducted in two steps: a coarse maximum thrust envelope comparison, and a quantification of the radial thrust feasibility margin.

4.2.3.2 Maximum Thrust Envelope Comparison

This first design validation stage examines the maximum thrust force the vehicle may exert in each direction about the $\{B\}$ y-z plane. If this is greater than the minimum feasible support force in all cases, then the design is potentially valid and progresses to the next stage. As it requires only the propeller orientations in $\{B\}$ and their thrust limits to assess a design, this stage is well suited to initial evaluation of candidate vehicles where the full propeller layout is not yet defined. When plotted visually, results may also support development of intuition regarding the effects of propeller orientation on feasible vehicle behaviour amid other practical constraints such as airflow and rotor clearance. As such, candidate designs may be quickly evaluated with minimal intermediate design revision work.

To support this evaluation, the maximum feasible thrust, $\mathbf{F}_{p,max}$, in a given direction, $\hat{\mathbf{u}}$, for an arrangement of propellers is determined via a linear optimization problem maximising the scalar length of the sum of individual motor thrust vectors.

$$\mathbf{F}_{p,max} = \left(\max_{F_n} \left\| \sum_{i=1}^n \mathbf{F}_{p,i} \right\| \right) \cdot \hat{\mathbf{u}} \quad (4-17)$$

$$\begin{aligned} \text{s.t.} \quad & \sum \mathbf{F}_{p,i} = \lambda \hat{\mathbf{u}} \\ & \lambda \geq 0 \\ & \mathbf{F}_{lp} \leq \mathbf{F}_q \leq \mathbf{F}_{up} \end{aligned}$$

Here, $\mathbf{F}_{p,max}$ must be the product of a positive scalar, $\lambda \in \mathbb{R}$, and $\hat{\mathbf{u}}$ to satisfy the direction constraint. The force of each propeller is also subject to lower and upper limits, represented by vectors \mathbf{F}_{lp} and \mathbf{F}_{up} , respectively. These limits depend on the specific motor, propeller aerofoil, and supply voltage employed and are not well estimated from commonly available manufacturer specifications. They are instead measured experimentally using a Tyto Robotics Series 1580 dynamometer [344], showing symmetrical bi-directional thrust generation up to ± 8.5 N in static thrust testing.

Using these thrust limits and the propeller layout of Table 4-2, the maximum thrust envelope for the prototype vehicle is evaluated via (4-17) and displayed in Figure 4-3. This indicates that the proposed vehicle may generate thrust force of up to ± 30 N and ± 35 N in the $\{B\}$ y and z axes, respectively, well above the vehicle weight in both cases. Notably, the maximum vertical thrust in the negative $\{B\}$ z-axis approaches double the vehicle weight, also suggesting strong flight performance with suitable overhead for aerial manoeuvrability. As the maximum thrust envelope encloses the line describing the minimum required stabilising thrust, the vehicle is shown to fully satisfy this first feasibility check.




Figure 4-3: Maximum propeller thrust of the candidate vehicle in $\{B\}$ relative to minimum required supporting force. The vehicle weight is also indicated for scale. No consideration is given for torque effects at this stage.

This, however, is insufficient to guarantee generic design feasibility across all asset diameters as torque effects have not yet been considered. Calculation of the maximum thrust force envelope per (4-17) imposes no conditions on net torque accompanying the force, and as such may not fully satisfy static equilibrium or describe a feasible interaction. Adding a further condition of zero net torque, or any other consistent value, likewise does not satisfy support requirements in all orientations about the surface. Further consideration is therefore required to ensure the vehicle can satisfy the supporting wrench criteria beyond the pure force term.

4.2.3.3 Radial Thrust Feasibility Margin

The second design validation stage provides a numeric metric of surplus wrench capacity beyond the minimum required supporting wrench, indicating the ability of the candidate vehicle to support airframe alterations, additional payload, and its robustness to modelling uncertainty. Full consideration is given to both force and torque requirements, allowing a passing design to guarantee generic feasibility.

As such, this stage evaluates at each clockface angle the maximum generable centripetal force in the $\{B\}$ z-axis, $F_c(\theta)$, while satisfying (4-16) in the other axes, using the vehicle actuator force effectiveness matrix derived in (4-8). Note that if no

solution exists, then the vehicle cannot produce any net wrench satisfying $\mathbf{W}_s(\theta)$ and has failed the feasibility check.

$$F_c(\theta) = \max_{\mathbf{F}_q} \mathbf{A}_{F,z} \mathbf{F}_q \quad (4-18)$$

$$\text{s.t} \quad \begin{bmatrix} \mathbf{A}_{F,\bar{z}} \\ \mathbf{A}_\tau \end{bmatrix} \mathbf{F}_q = \mathbf{W}_{s,\bar{F}_z}(\theta)$$

$$\mathbf{F}_{lp} \leq \mathbf{F}_q \leq \mathbf{F}_{up}$$

Here, $\mathbf{A}_{F,z}$ is the row vector of \mathbf{A}_F yielding net z-axis force. $\mathbf{A}_{F,\bar{z}}$ is then \mathbf{A}_F omitting this row vector, and \mathbf{W}_{s,\bar{F}_z} is \mathbf{W}_s omitting the matching z-axis force row. As previously, the column vectors \mathbf{F}_{lp} and \mathbf{F}_{up} define the lower and upper thrust limits of each motor, respectively.

The design margin for generic circumferential interaction is then the minimum difference between $F_c(\theta)$ and $F_s(\theta)$ in the $\{B\}$ z-axis, denoted F_Δ as below.

$$F_\Delta = \min \left(F_c(\theta) - F_{s,z}(\theta) \right) \quad (4-19)$$

In this interaction, note that any value of $F_c(\theta)$ is inherently stabilised by the reactive nature of N_o , with increased radial propeller force opposed by an equal increase in surface contact force with no resulting disturbance to vehicle position. Therefore, any combination of propeller effects granting $F_c(\theta)$ above the support threshold is guaranteed to represent a feasible interaction solution. By satisfying (4-16) in this manner, it is explicitly shown that the candidate vehicle may meet both the minimum force and torque requirements of static equilibrium at all points around the asset circumference regardless of diameter, guaranteeing full generic feasibility.

Evaluating (4-18) for the proposed vehicle with the propeller array described in Table 4-2 yields the results depicted in Figure 4-4. As shown, the design is generically feasible, with positive F_Δ of 3.20 N, and so may hold arbitrary position at any point around a cylindrical asset of any diameter.

Furthermore, the surplus design capacity indicated by F_Δ allows relative comparison between alternate designs, with a larger value representing increased support capacity.

However, note that the interdependent nature of the interaction system means that F_{Δ} cannot be directly converted to additional payload mass due its effects on $W_s(\theta)$. A balance must also be struck between F_{Δ} and the additional mass and power requirements of any unused actuation capability. The design metric therefore indicates a small capacity for further payload or system changes, sufficient to validate the vehicle design ahead of prototype construction.




Figure 4-4: Design feasibility margin plot for the candidate vehicle design. F_{Δ} is 3.20 N, occurring at θ of $\pm 164^\circ$, indicated in red.

4.3 Optimal Thrust Support Setpoint

With the vehicle design validated, a practical framework to minimise thrust-support power usage while ensuring interaction stability is developed, implemented as a piecewise function for real-time feed-forward control of vehicles sharing the hybrid UAV-crawler archetype. This examines the opportunities posed by the specific system beyond the generic interaction criteria described above.

4.3.1 Wrench Energy Optimization

At each clockface angle about the horizontal pipe, the goal is to determine the set of argument parameters describing the interaction forces and net propeller wrench that minimises total motor power usage, P . This cost function is approximated by the sum of each squared motor thrust, minimising control effort in the least-squares sense as is

common in previous works concerning net wrench allocation between multiple UAV propellers [185], [219], [345].

$$P(\mathbf{F}_q) = \sum F_{p,i}^2 \quad (4-20)$$

Practical constraints beyond the quasi-static equilibrium of (4-10) and (4-11) are also imposed. Column vectors \mathbf{F}_{lp} and \mathbf{F}_{up} again define the lower and upper propeller thrust limits, respectively. Similarly, the non-zero normal force condition and static friction limits used in (4-14) again include a safety factor to mitigate model uncertainty and ensure a robust interaction. Lastly, the legs and wheel contact points enclose a cylinder marginally larger than the target diameter. Thus N_L and N_R cannot be negative or simultaneously nonzero without airframe deformation. The minimal energy support solution is then found by the expression below.

$$\underset{\mathbf{F}_q, N_O, N_L, N_R, F_{fr}}{\operatorname{argmin}} P(\mathbf{F}_q) \quad (4-21)$$

s.t.

$$\begin{aligned} \mathbf{F}_{lp} &\leq \mathbf{F}_q \leq \mathbf{F}_{up} \\ N_O &\geq N_{min} \\ -\mu N_O &\leq F_{fr} \leq +\mu N_O \\ N_L, N_R &\geq 0 \\ N_L N_R &= 0 \\ \sum \mathbf{F} &= 0 \\ \sum \boldsymbol{\tau} &= 0 \end{aligned}$$

Using this with the prototype vehicle parameters described above yields the desired interaction controller behaviour, as in Figure 4-5. This is robustly solved using 50 randomised trials for 1024 clockface angles evenly distributed around the circumference using the MATLAB optimization toolbox, taking approximately 28 min when using a 4 parallel cores of an i7-7700HQ laptop CPU at 3.5 GHz clock speed. At the average solve rate of 7.6 Hz per core, the high dimensionality of the problem therefore renders on-line optimisation of (4-21) intractable within an embedded flight controller, precluding its direct use within an MPC framework. An alternate approach using off-line optimisation is therefore considered.

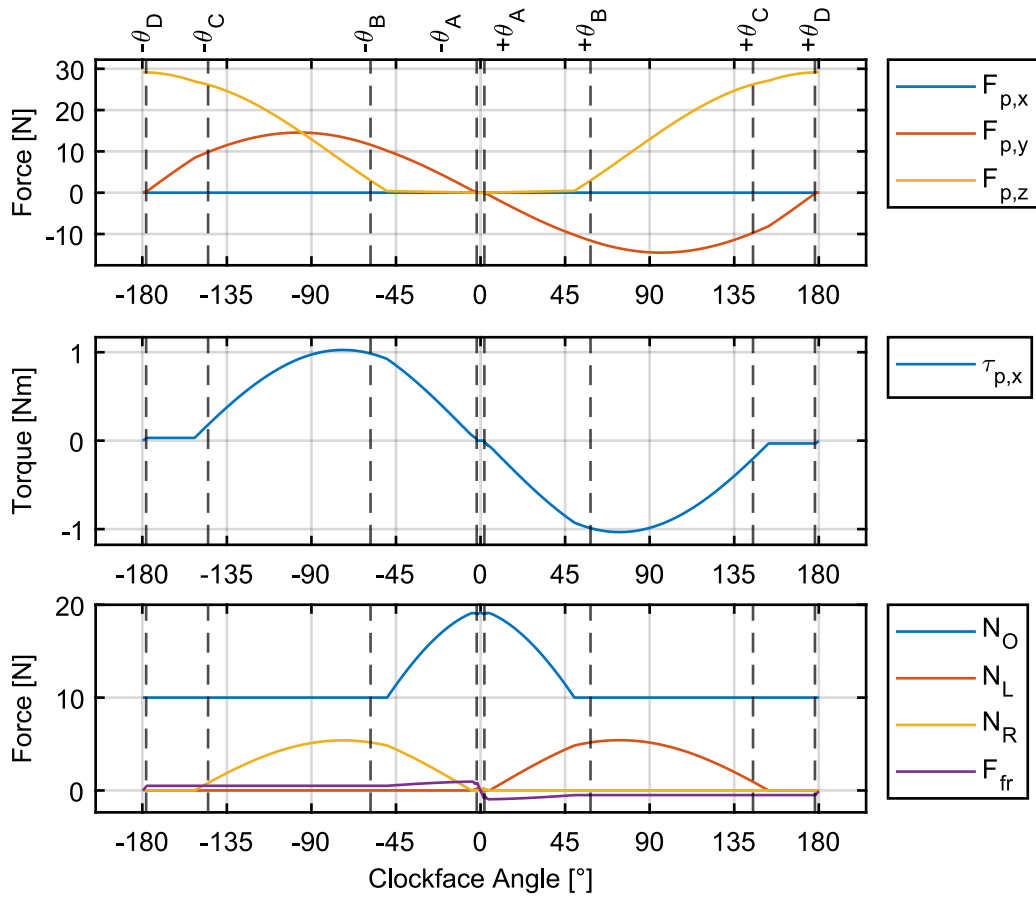


Figure 4-5: Energy optimal interaction force-torque setpoint and contact forces computed around the target surface. Key clockface angles are denoted for further consideration, below.

4.3.2 Optimal Supporting Leg Torque

Here energy-optimal behaviour is examined at several key points around the pipe, matching the annotations in Figure 4-5. Recognizing the significant role of leg torque, a simple function approximating its optimal response is derived to permit a solution to (4-21) to be found in real-time at the embedded controller loop rate of 200 Hz. This is sufficiently generic to also be applicable to other target assets and in control of other vehicles of the same archetype. Given the symmetry of the vehicle in the $\{B\}$ x-z plane, only the region $\theta \in [0, +180^\circ]$ is explicitly considered. Behaviour for $\theta < 0^\circ$ is then the negative of that in the examined region.

Here, scalar components of a vector aligned with a given axis are denoted using trailing subscripts, e.g., the component of $\hat{\mathbf{u}}_L$ along the x-axis is $\hat{u}_{L,x}$. Left and right leg torque,

$\boldsymbol{\tau}_L$ and $\boldsymbol{\tau}_R$, may also be expressed as the product of the normal contact force scalar and a constant vector moment arm, \mathbf{m}_L and \mathbf{m}_R .

$$\boldsymbol{\tau}_L = N_L(\mathbf{L}_L \times \hat{\mathbf{u}}_L) = N_L \mathbf{m}_L \quad (4-22)$$

$$\boldsymbol{\tau}_R = N_R(\mathbf{L}_R \times \hat{\mathbf{u}}_R) = N_R \mathbf{m}_R \quad (4-23)$$

4.3.2.1 Top & Bottom, $\theta = 0$ & $\theta = \pm 180^\circ$

At the top and bottom points of the circumference, vehicle weight acts purely in the $\{B\}$ z-axis. Thus, contact forces outside this axis are not beneficial and zero torque is applied.

4.3.2.2 Point A, $\theta = \theta_A$

Here, the component of $m\mathbf{g}$ in the $\{B\}$ y-axis reaches the limit of passive support by F_{fr} and the propellers begin to act. Prior to θ_A , unbalanced torque is entirely supported by the downhill leg. Hence, the system is subject to the conditions

$$F_{fr} = \mu N_O, \quad F_p = \tau_p = 0, \quad N_L = 0 \quad (4-24)$$

Applying quasi-static equilibrium by (4-10) and (4-11), these conditions define θ_A and the optimal leg torque, τ_A , at that point

$$\tan(\theta_A) = \left(\frac{\frac{m_{R,x}}{L_O} + \hat{u}_{R,y}}{\frac{m_{R,x}}{\mu L_O} + \hat{u}_{R,z}} \right) \quad (4-25)$$

$$\tau_A = - \left(\frac{mg \sin(\theta_A)}{\frac{m_R}{L_O} + \hat{u}_{R,y}} \right) m_R \quad (4-26)$$

Per a small angle assumption for θ_A to ease computation, torque exerted by the legs on the UAV in the region $\theta \in [0, \theta_A]$ is approximately linear, passing through the origin with gradient m_A defined by

$$m_A = \frac{\tau_A}{\theta_A} \quad (4-27)$$

4.3.2.3 Point B, $\theta = \theta_B$

Here, θ_B is defined as the clockface angle where the component of $m\mathbf{g}$ in the $\{B\}$ z-axis is reduced to the limit N_{min} . It follows then that at this point

$$N_O = N_{min}, \quad F_{fr} = \mu N_{min} \quad (4-28)$$

However, these criteria prove insufficient to define \mathbf{F}_p and $\boldsymbol{\tau}_p$ by (4-10) and (4-11). The following additional constraints are therefore imposed:

$$F_{p,y} = k_f N_{L,y}, \quad F_{p,z} = -N_{L,z} \quad (4-29)$$

Here k_f sets the ratio at which support for $m\mathbf{g}$ in the $\{B\}$ y-axis is divided between \mathbf{F}_p and the uphill leg contact force, generated via $\boldsymbol{\tau}_p$. The value of k_f thereby depends on the vehicle-specific relative efficiency of net propeller force and torque generation. Here, a value of $k_f = 2.75$ is found to give good agreement with the energy-optimal interaction solution.

With $m\mathbf{g}$ in the $\{B\}$ z-axis passively supported by N_O , (4-29) poses a further logical constraint when minimizing propeller power use, exerting the minimum thrust component to counteract $N_{L,z}$ and grant equilibrium. In fact, the optimal $F_{p,z}$ marginally exceeds this condition, increasing \mathbf{F}_{fr} and reducing θ_B . However, this effect is non-trivially expressed alongside the y-axis constraint of (4-29) and is highly sensitive to system parameter changes, so is omitted for generality.

By (4-10) and (4-11), the constraints of (4-29) then yield the definition for θ_B and the leg torque exerted there, τ_B

$$\cos(\theta_B) = \frac{N_{min}}{mg} \quad (4-30)$$

$$\tau_B = m_{L,x} \left(\frac{\mu N_{min} - mg \sin(\theta_B)}{(k_f + 1) \hat{u}_{L,y}} \right) \quad (4-31)$$

Optimal leg torque in the region $\theta \in (\theta_A, \theta_B]$ is then approximated by a straight line with gradient m_B and offset c_B given by

$$m_B = \frac{\tau_B - \tau_A}{\theta_B - \theta_A}, \quad c_B = \tau_B - m_B \theta_A \quad (4-32)$$

4.3.2.4 Point C, $\theta = \theta_C$

Here angle θ_C is defined as the point where the uphill leg contact force ceases opposing the vehicle weight and so becomes detrimental to support, imposing the conditions

$$\hat{\mathbf{u}}_L \cdot m\mathbf{g} = 0, \quad \tau_C = 0 \quad (4-33)$$

By (4-2) and (4-33), θ_C may be expressed in terms of the axial components of $\hat{\mathbf{u}}_L$ as

$$\tan(\theta_C) = \left(-\frac{\hat{u}_{L,x}}{\hat{u}_{L,y}} \right) \quad (4-34)$$

This differs slightly from the optimal behaviour as, for a small region past θ_C , the specific rotor configuration used here may more efficiently generate additional leg torque and support the resulting downward normal force than avoid creating the torque at all. Again, this phenomenon cannot be generically captured in net propeller behaviour and must be omitted. Given the basis of this function on clockface angle and noting that the optimal solution for τ_{LR} in the region $\theta \in (\theta_B, \theta_C]$ resembles a sine function, an angular scalar factor, f_C , may be defined to give zero torque at θ_C .

$$f_C = \frac{\pi}{\theta_C} \quad (4-35)$$

Combining this with τ_B , the sine amplitude, A_C , is then

$$A_C = \frac{\tau_B}{\sin(f_c \theta_B)} \quad (4-36)$$

Note that a similar response may also be obtained using an appropriate quadratic function in the case where the embedded controller cannot easily compute sine functions.

4.3.2.5 Point D, $\theta = \theta_D$

This final point is added in consideration of the practical system. To prevent leg roll chatter due to mechanical tolerances, a small but non-zero minimum torque, τ_{min} , is applied in the region $\theta \in (\theta_B, 180^\circ]$, improving overall stability at a small cost to efficiency. Arbitrary angle θ_D is then set between θ_C and π as the point starting linear transition to zero torque at $\theta = \pm 180^\circ$, preventing discontinuous step-changes when crossing the bottom point. For sake of convenience, the following equality is applied.

$$\theta_D = \pi - \theta_A, \quad \tau_D = \tau_{min} \quad (4-37)$$

The linear transition region then has gradient, m_D

$$m_D = -\frac{\tau_{min}}{\theta_A} \quad (4-38)$$

4.3.2.6 Optimal Leg Torque Approximation Function

Using the above piecewise equations, and exploiting vehicle symmetry, the function approximating optimal supporting leg torque about the $\{B\}$ x-axis, $\tau_{LR}(\theta)$, is then,

$$\tau_{LR}(\theta) = \begin{cases} m_A \theta, & \theta \in [0, \theta_A] \\ m_B \theta + c_B, & \theta \in (\theta_A, \theta_B] \\ \max(A_C \sin(f_c \theta), \tau_{min}), & \theta \in (\theta_B, \theta_D] \\ m_D (\theta - \pi), & \theta \in (\theta_D, \pi] \\ -\tau_{LR}(-\theta), & \theta \in [-\pi, 0) \end{cases} \quad (4-39)$$

This function is plotted alongside the optimal behaviour for the prototype vehicle in Figure 4-6. Good general agreement is shown, with maximum absolute difference under 0.171 Nm. This is due to the position of θ_B and θ_C amid the complex relationship between MDT net force, torque, and energy usage, as indicated above. Influence on interaction is quantified within the analysis of the larger control structure, below.




Figure 4-6: Comparison of energy optimal leg torque setpoint (solid line) and the piecewise function setpoint (dashed line). Annotated clockface angles match the points of interest from the analysis above.

4.4 System Control

Using the function for energy-optimal supporting leg torque, the full vehicle control structure is implemented in a custom version of Px4 [346], as illustrated in the signal flow diagram of Figure 4-7 and discussed in detail hereafter. Within this control architecture, discrete flight and interaction control pathways generate independent wrench setpoints for the propellers. A mode-select parameter, set either by the pilot or in software, determines which wrench signal is allocated to the propeller motor throttle signals and temporarily disables the unused path. In free flight, this permits vehicle attitude stabilisation and tracking of desired trajectories. During contact, the propellers generate the wrench required to adaptively stabilise the interaction while separate differential steering control of two driven wheels effects motion about the asset surface.

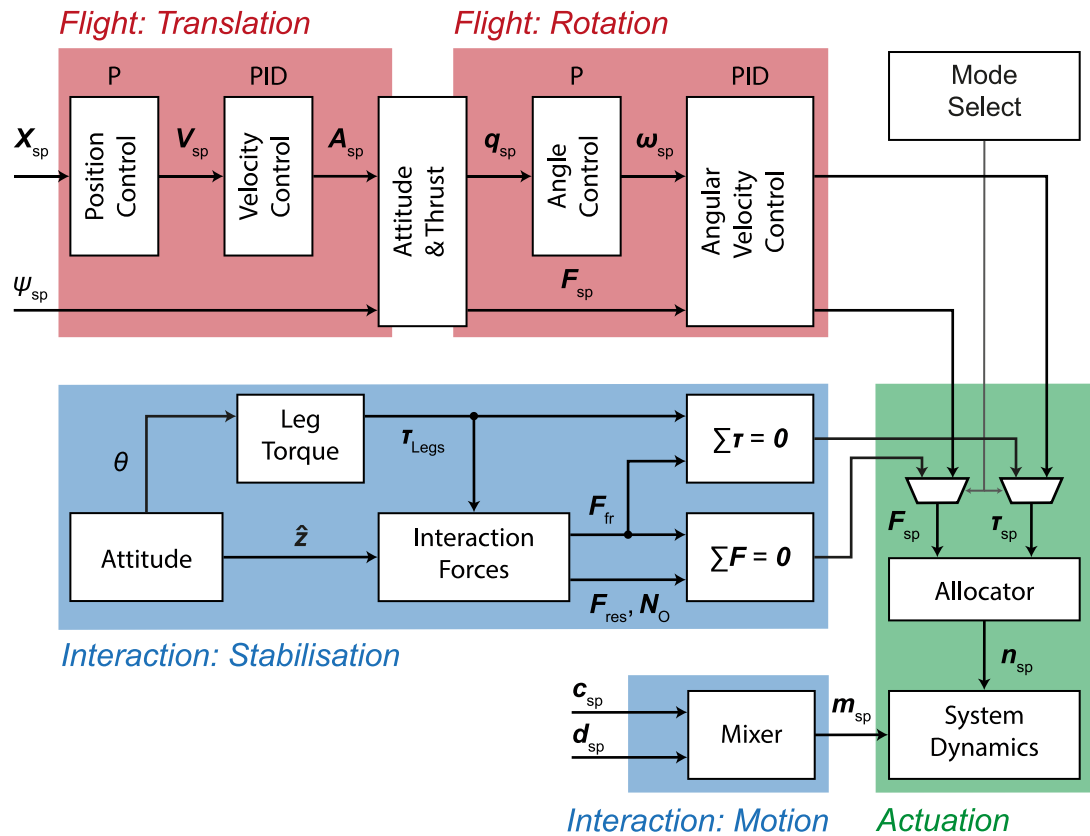


Figure 4-7: UAV-crawler hybrid controller block diagram. Separate signal flow paths are used for flight and interaction control functions. Symbols are defined in their related control section.

In implementing this control structure, some standard components of the Px4 firmware have been reused or modified. The existing Position, Velocity, Angle, and Angular Velocity control blocks are functionally independent from the vehicle design and so require only tuning of internal parameters to account for vehicle mass and inertia. However, the non-linear function of the Attitude & Thrust block is revised to enable multidirectional thrust generation while in flight. The mode selection state machine is also substantially modified to ensure the new interaction control mode retains compatibility with the various safety cases, fail-safes, and telemetry logging structures used in the standard system [347]. Modifications to the Allocator block enable it to function with the array of bi-directional thrust propellers used by the hybrid vehicle, supplying their body-relative pose and thrust generation limits while also applying a suitable internal model to convert desired per-motor thrust to a normalised throttle signal. The novel interaction control structure is newly implemented alongside the differential steering signal mixer to complete the control system depicted in Figure

4-7. The operating principles of all of these component modules are examined below to provide a full understanding of system function and highlight beneficial aspects.

Here, the relative orientation of coordinate frames may be described via rotation matrices. For example, ${}^W_B\mathbf{R}$ is the vehicle attitude: the orientation of $\{B\}$ relative to world inertial frame $\{W\}$. A quaternion denoting the same is ${}^W_B\mathbf{q}$. The unit basis vectors of one frame, $\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}$, share this notation when referenced in another frame. For example, ${}^B_W\hat{\mathbf{z}}$ is the $\{W\}$ z-axis expressed in $\{B\}$, giving the direction of gravity relative to the vehicle.

4.4.1 Flight Control

In free-flight, vehicle position and attitude control use the nested proportional, integral, derivative (PID) structure of Px4, providing vehicle stabilization and trajectory setpoint tracking with a high-level architecture similar to that depicted in Figure 2-12. A summary of key points and modifications made to the open-source base firmware for this unique vehicle is therefore given.

In this section, multiple distinct gain vectors are denoted as $\mathbf{K} \in \mathbb{R}^3$. Subscripts give the relevant controller and PID term. Estimated values of a given parameter are also indicated by a tilde symbol, e.g. the vehicle position estimate, $\tilde{\mathbf{X}}$.

4.4.1.1 Translation Control

Translation control is computed in the inertial world frame $\{W\}$ by the outermost loop. This generates a velocity setpoint vector, \mathbf{V}_{sp} , by proportional gain of the error in the current desired position, \mathbf{X}_{sp} , versus its estimate, $\tilde{\mathbf{X}}$.

$$\mathbf{V}_{sp} = \mathbf{K}_{pos,P}(\mathbf{X}_{sp} - \tilde{\mathbf{X}}) \quad (4-40)$$

Desired acceleration setpoint, \mathbf{A}_{sp} , is then given by the velocity error, \mathbf{V}_e , in the current estimated velocity, $\tilde{\mathbf{V}}$, versus the setpoint per a PID algorithm. Basing the derivative term on the estimated acceleration, $\tilde{\mathbf{A}}$, reduces the impact of noise in the numerical derivative of the estimated velocity and sudden impulse responses due to step changes in the velocity setpoint.

The velocity error and PID function are then expressed as follows.

$$\mathbf{V}_e = \mathbf{V}_{sp} - \tilde{\mathbf{V}} \quad (4-41)$$

$$\mathbf{A}_{sp} = \mathbf{K}_{vel,P}\mathbf{V}_e + \int \mathbf{K}_{vel,I}\mathbf{V}_e dt - \mathbf{K}_{vel,D}\tilde{\mathbf{A}} \quad (4-42)$$

4.4.1.2 Attitude and Thrust Setpoints

Here modifications to the base Px4 controller enable translation in the $\{B\}$ y-axis without roll. This is highly beneficial to precision landing in the context of pipe and cylindrical asset inspection due to the many complex disturbance mechanisms acting to shed the vehicle from its target location.

When landing atop a pipe, the “ground effect” caused by interaction of the propeller downdraft with the environment differs significantly from that experienced while landing on a planar surface [339], [348]. In the pipe case, the colliding airflow alters both lift and lateral force components, varying in strength with surface proximity, and acting to displace the vehicle from the landing position directly above the pipe axis. This is exacerbated by vehicle roll which, in attempting to correct lateral position, redirects the main thrust airflow relative to the pipe and alters the aerodynamic disturbance in a highly non-linear manner. Further, such roll may cause unintended interaction of the UAV legs with the pipe. In the absence of an appropriate sensor feedback pathway, these forces are treated as an unmodelled disturbance leading to additional potential for position controller instability.

Direct thrust generation in the $\{B\}$ y-axis enables the main lift thrust to retain a more consistent airflow direction, while another lateral airflow effects translation without significant pipe interaction. Without the need to reorient the main thrust via roll, undesired leg contact during landing is markedly reduced. Dynamic responsiveness of the vehicle position control is also thereby improved, replacing limits posed by the vehicle moment of inertia about the $\{B\}$ x-axis with those of the significantly faster propeller motor throttle response.

To this end, a control structure capable of lateral translation without roll is developed as a more performant means to reject disturbances and regulate vehicle position above the pipe axis.

As such, vehicle attitude and thrust setpoints are computed via an intermediate coordinate frame $\{A\}$ aligned to the desired yaw heading, ψ_{sp} , by rotation of $\{W\}$ about its z-axis via ${}^W_A\mathbf{R}$.

$${}^W_A\mathbf{R} = \begin{bmatrix} \cos(\psi_{sp}) & -\sin(\psi_{sp}) & 0 \\ \sin(\psi_{sp}) & \cos(\psi_{sp}) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (4-43)$$

First the thrust setpoint in $\{A\}$, ${}^A\mathbf{F}_{sp}$, is found by adding acceleration opposing gravity, \mathbf{g} , to \mathbf{A}_{sp} , giving acceleration about a stable hover, before scaling by the vehicle mass and rotating the result into $\{A\}$

$${}^A\mathbf{F}_{sp} = {}^W_A\mathbf{R}^{-1} \left(m({}^W\mathbf{A}_{sp} - {}^W\mathbf{g}) \right) \quad (4-44)$$

The ${}^A\mathbf{F}_{sp}$ vector is then split into components directly effectible in the $\{A\}$ y-axis, ${}^A\mathbf{F}_{sp,dir}$, and those made collectively by pitching the vehicle about this axis, ${}^A\mathbf{F}_{sp,att}$.

$${}^A\mathbf{F}_{sp,dir} = [0 \quad {}^A F_{sp,y} \quad 0]^T \quad (4-45)$$

$${}^A\mathbf{F}_{sp,att} = [{}^A F_{sp,x} \quad 0 \quad {}^A F_{sp,z}]^T \quad (4-46)$$

Desired vehicle attitude sets ${}^A_B\hat{\mathbf{z}}$ as the unit vector opposing $\mathbf{F}_{sp,att}$, per the FRD coordinate convention. Without roll, $\{A\}$ and $\{B\}$ y-axis alignment then yields ${}^A_B\hat{\mathbf{x}}$ by vector cross product. Collection of orthonormal vectors thereafter sets ${}^A_B\mathbf{R}$ and thus the desired attitude rotation matrix ${}^W_B\mathbf{R}$.

$${}^A_B\hat{\mathbf{x}} = {}^A_B\hat{\mathbf{y}} \times {}^A_B\hat{\mathbf{z}} \quad (4-47)$$

$${}^A_B\mathbf{R} = [{}^A_B\hat{\mathbf{x}} \quad {}^A_B\hat{\mathbf{y}} \quad {}^A_B\hat{\mathbf{z}}], \quad {}^W_B\mathbf{R} = {}^W_A\mathbf{R} {}^A_B\mathbf{R} \quad (4-48)$$

This is subsequently expressed as a quaternion, \mathbf{q}_{sp} , using Shepperd's algorithm [349] for use as the desired orientation within the rotation controller.

The thrust setpoint in the desired $\{B\}$ frame is then found via its common y-axis direction with $\{A\}$ and the alignment of the $\{B\}$ z-axis to the desired attitude thrust direction.

$${}^B\mathbf{F}_{sp} = {}^A\mathbf{F}_{sp,dir} - \left\| {}^A\mathbf{F}_{sp,att} \right\| \hat{\mathbf{z}} \quad (4-49)$$

Note that this assumes good alignment between the current vehicle roll angle and the desired orientation per action of the attitude controller. Any residual effects due to small misalignment are handled as disturbances by the position control loop.

4.4.1.3 Rotation Control

Rotation control is computed as in the standard Px4 firmware using the angular rate setpoint generation algorithm proposed by Brescianini *et al* [350]. The approach is similar to that used in Section 3.2.1 but uses a different expression for the desired torque, so is revisited here to fully describe the control structure of this hybrid vehicle.

First, the difference between the current estimated attitude quaternion, $\tilde{\mathbf{q}}$, and the desired attitude setpoint, \mathbf{q}_{sp} , is evaluated. Their difference quaternion is computed with the quaternion product operator and the inverse of $\tilde{\mathbf{q}}$ as follows.

$$\mathbf{q}_e = \mathbf{q}_{sp} \otimes \tilde{\mathbf{q}}^{-1} = \begin{bmatrix} q_w \\ \mathbf{q}_v \end{bmatrix} \quad (4-50)$$

The desired angular rate vector is then generated in proportion to the vector component of the quaternion attitude error, \mathbf{q}_v . Note that the ‘‘unwinding problem’’, where the vehicle unnecessarily moves through a full rotation due to the dual representation of attitudes in quaternion space [314], is resolved using the sign of the quaternion error scalar, q_w .

$$\boldsymbol{\omega}_{sp} = 2 \operatorname{sgn}(q_w) \mathbf{K}_{att,P} \mathbf{q}_v \quad (4-51)$$

Finally, desired torque in $\{B\}$, $\boldsymbol{\tau}_{sp}$, is then given by a PID expression of the rotation rate error, $\boldsymbol{\omega}_e$, versus the estimated rotation rate, $\tilde{\boldsymbol{\omega}}$, proportional to the vehicle

moment of inertia matrix, \mathbf{J} , with a feed-forward term accounting for the rotation of the body frame. Note that the hybrid vehicle is not required to exert strong interaction forces during flight, so the desired interaction force term may be omitted versus the formulation used by the Voliro platform in (3-5). The absence of obstructions to rotation introduced by in-flight interactions also allows for inclusion of an integral term, improving pose tracking performance without risk of saturation instability.

$$\boldsymbol{\omega}_e = \boldsymbol{\omega}_{sp} - \tilde{\boldsymbol{\omega}} \quad (4-52)$$

$$\boldsymbol{\tau}_{sp} = \mathbf{J}(\mathbf{K}_{rot,P}\boldsymbol{\omega}_e - \mathbf{K}_{rot,D}\tilde{\boldsymbol{\alpha}}) + \mathbf{K}_{rot,I} \int \boldsymbol{\omega}_e dt + \tilde{\boldsymbol{\omega}} \times \mathbf{J}\tilde{\boldsymbol{\omega}} \quad (4-53)$$

As with the translation control, the differential term is based on the estimated angular acceleration, $\tilde{\boldsymbol{\alpha}}$, minimising disturbances due to numerical differentiation and sudden setpoint changes. This completes the functionality of the flight controller and provides vehicle stabilization and trajectory setpoint tracking while in free flight outside of contact with the inspection target.

4.4.2 Interaction Control

Forming the second component of Figure 4-7, a feed-forward algorithm generating the thrust-support propeller wrench setpoint in real-time is proposed utilising an approximation of the minimum energy solution to (4-21) formed via (4-39). This allows portable deployment with low processor and sensor requirements, suitable for inclusion in various control structures and is presented in Algorithm 1, below.

Initially, Algorithm 1 uses the current body attitude estimate quaternion, ${}^W_B\tilde{\mathbf{q}}$, to determine the world “down” direction in $\{B\}$, i.e. ${}^B_W\hat{\mathbf{z}}$. Inversion of ${}^W_B\tilde{\mathbf{q}}$ gives the orientation of the world frame relative to the UAV body. Expression as a rotation matrix, ${}^B_W\mathbf{R}$, then allows ${}^B_W\hat{\mathbf{z}}$ to be read from the 3rd column [192]. Clockface angle, θ , is subsequently found by trigonometry.

Supporting leg torque, τ_{legS} , is then evaluated per (4-39). This then gives the normal contact vectors \mathbf{N}_L and \mathbf{N}_R . As in the optimal setpoint formulation of (4-21), their

physical nature means both must be non-negative but only one may be positive, allowing their determination by the sign of τ_{Legs} and UAV geometry.

With the vehicle vector weight in $\{B\}$, $m\mathbf{g}$, given by scaling ${}^B_W\hat{\mathbf{z}}$, the residual force, \mathbf{F}_{res} , is found by summation. To minimise the effort by the propellers, the wheel normal contact force, \mathbf{N}_O , and friction, \mathbf{F}_{fr} , are then set to support as much of \mathbf{F}_{res} as possible, per conditions $N_O \geq N_{min}$ and $|F_{fr}| \leq \mu N_O$.

Finally, desired net propeller wrench, $[\mathbf{F}_p \quad \boldsymbol{\tau}_p]^T$, is calculated by the condition for quasi-static equilibrium.

Algorithm 1: Interaction Force-Torque Setpoint

Input: Current orientation estimate quaternion, ${}^W_B\tilde{\mathbf{q}}$

Output: Desired Propeller Wrench Setpoint, $[\mathbf{F}_p, \boldsymbol{\tau}_p]$

1: Get Attitude

$$\begin{aligned} {}^B_W\mathbf{R} &\leftarrow \text{quat2rotm}({}^W_B\tilde{\mathbf{q}}^{-1}) \\ {}^B_W\hat{\mathbf{z}} &\leftarrow {}^B_W\mathbf{R}_{i,3} \\ \theta &\leftarrow \text{atan2}({}^B_W\hat{\mathbf{z}}_y, {}^B_W\hat{\mathbf{z}}_z) \end{aligned}$$

2: Get Leg Torque

$$\tau_{Legs} \leftarrow \tau_{LR}(\theta)$$

3: Get Interaction Forces

$$\begin{aligned} s &\leftarrow \text{sgn}(\tau_{legs}) \\ \mathbf{N}_L &\leftarrow \max(0, s) \frac{\tau_{legs}}{\|\mathbf{L}_L\|} \hat{\mathbf{u}}_L \\ \mathbf{N}_R &\leftarrow \min(0, s) \frac{\tau_{legs}}{\|\mathbf{L}_R\|} \hat{\mathbf{u}}_R \\ m\mathbf{g} &\leftarrow m\mathbf{g} {}^B_W\hat{\mathbf{z}} \\ \mathbf{F}_{res} &\leftarrow m\mathbf{g} + \mathbf{N}_L + \mathbf{N}_R \\ \mathbf{N}_O &\leftarrow \min(-N_{min}, -F_{res,z}) \hat{\mathbf{z}} \\ \mathbf{F}_{fr} &\leftarrow -\text{sgn}(F_{res,y}) \min(\mu \|\mathbf{N}_O\|, |F_{res,y}|) \hat{\mathbf{y}} \end{aligned}$$

4: Apply Quasi-Static Equilibrium

$$\begin{aligned} \mathbf{F}_p &\leftarrow -(\mathbf{N}_O + \mathbf{F}_{Fr} + \mathbf{F}_{res}) \\ \boldsymbol{\tau}_p &\leftarrow -(\mathbf{l}_O \times \mathbf{F}_{fr} + \tau_{legs} \hat{\mathbf{x}}) \end{aligned}$$

A comparison of the optimal wrench setpoint results and component forces to those obtained by Algorithm 1 is given in Figure 4-8. These show very close agreement, with difference largely due to the positioning of θ_B and θ_C in (4-39). Maximum absolute difference in \mathbf{F}_p is less than 0.758 N and 0.530 N in the $\{B\}$ y and z axes, respectively, and in $\boldsymbol{\tau}_p$ is less than 0.171 Nm. Force differences then fall within 4 % of system weight. Torque difference is within 17 % of peak applied torque but is of limited concern with only a minor impact on the total energy usage. Algorithm 1 is thereby validated in comparison to the energy-optimal setpoint established in Section 4.3.1, yielding feasible real-time thrust support within an acceptable margin for error.

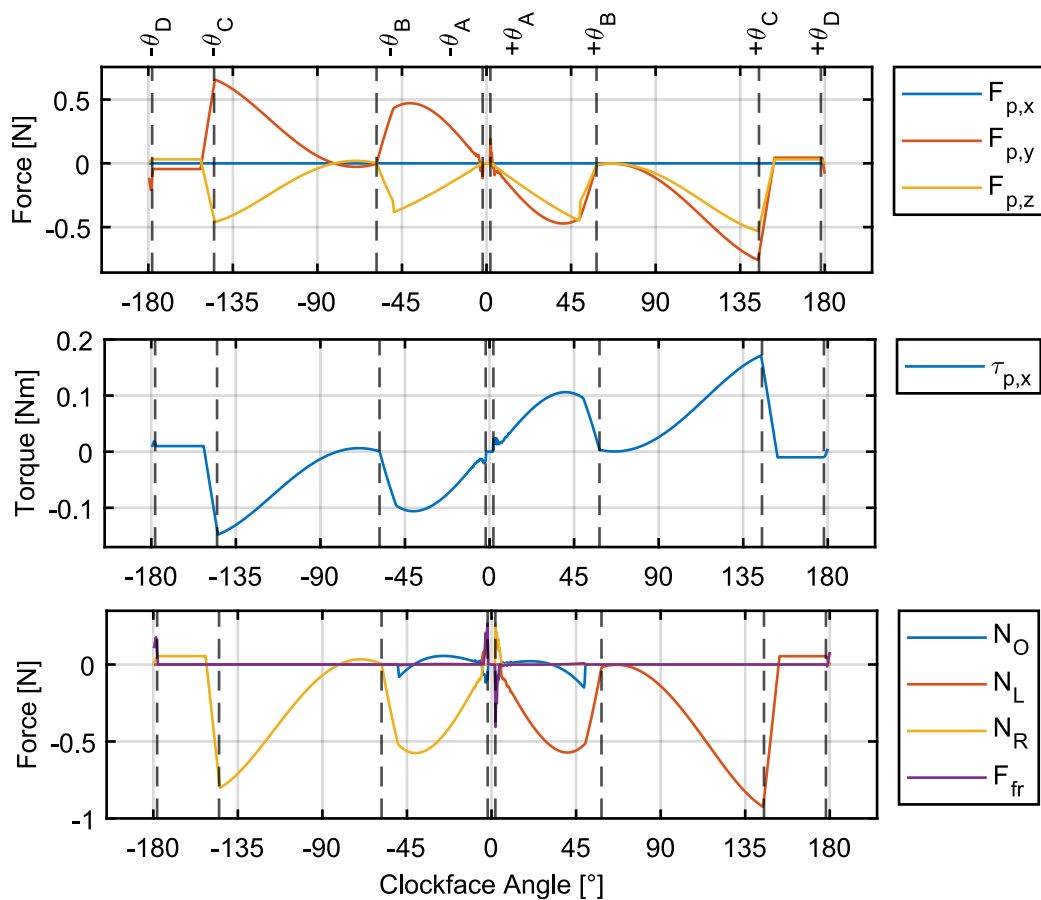


Figure 4-8: A plot of error between the optimal interaction propeller force-torque setpoint and the approximation generated using Algorithm 1. Errors in the interaction forces are also shown.

4.4.3 Wrench Allocation

Completing the propeller control architecture of Figure 4-7, the allocator maps the wrench setpoint from either the flight or interaction modules to an appropriate motor throttle signal.

Here, the Moore-Penrose pseudo-inverse of the system model given in (4-7) is used similarly to its deployment in previous MDT and VT UAV systems [258], [263]. This yields a least-squares solution of individual propeller forces [218], with net behaviour matching the desired wrench in the absence of saturation.

$$\mathbf{F}_q = \begin{bmatrix} \mathbf{A}_F \\ \mathbf{A}_\tau \end{bmatrix}^\dagger \begin{bmatrix} \mathbf{F}_p \\ \boldsymbol{\tau}_p \end{bmatrix} \quad (4-54)$$

Each normalised bi-directional motor throttle setpoint, $n_{sp,i} \in [-1, +1]$, giving propeller force $F_{p,i}$ is then found using an inverse motor model. As motor speed set by the ESCs correlates approximately linearly to the throttle signal, (4-3) is rearranged and expressed with a normalised thrust coefficient, C_n , as below.

$$n_{sp,i} = \text{sgn}(F_{p,i}) \sqrt{\frac{|F_{p,i}|}{C_n}} \quad (4-55)$$

Thus, the signal flow from both the flight and interaction control modules to the individual propeller motor throttle signals is fully described, defining how the vehicle may perform stable flight and maintain position on the pipe surface.

4.4.4 Surface Crawling

The final component of the control system depicted in Figure 4-7 governs the crawling motion of the hybrid vehicle during contact with the target asset. It provides throttle signals to two active servo motors, one mounted at the front of the vehicle and another at the back. Together these drive the vehicle around the asset circumference, enabling access to the desired inspection location.

To fully inspect the pipe exterior surface, the vehicle must also be capable of motion along the pipe length. For large displacements, this is best achieved by re-entering

flight and landing at the desired point. However, the small corrections required to examine a specific feature of interest or return to an exact defect position to track its growth over time are more easily achieved by moving along the surface in contiguous contact. This may be conducted by employing the front and rear servo motors in a differential steering approach, similar to that common in robotic crawlers [191]. A kinematic overview of the application of this behaviour within the hybrid UAV-crawler vehicle is provided below and depicted in Figure 4-9.




Figure 4-9: Diagram of differential steering crawler kinematic used by the hybrid vehicle. Annotations show key vehicle dimension and their relationship to the instantaneous velocities in $\{B\}$.

Acting collectively, the two vehicle servo motors may generate motion in a single translation and rotation axis relative to the vehicle body frame. The translational component, v_y , acts in the $\{B\}$ y-axis, while the rotational component, ω_z , acts to reorient the vehicle through yaw about the $\{B\}$ z-axis. The result is motion about an Instantaneous Centre of Rotation (ICR) which has variable position along the shared axis of the two servo wheels.

By the relationship between linear and angular velocity, it may be shown that this distance from the origin of $\{B\}$ to the ICR, denoted w_R , can be expressed as a function of the translational velocity of the front wheel, v_1 , and the rear wheel, v_2 , and their distances from the origin, w_1 and w_2 , respectively.

$$w_R = \frac{v_1 w_2 + v_2 w_1}{v_2 - v_1} \quad (4-56)$$

Using this, the translational and rotational velocities can then be similarly expressed.

$$v_y = \frac{v_1 w_2 + v_2 w_1}{w_1 + w_2} \quad (4-57)$$

$$\omega_z = \frac{v_2 - v_1}{w_1 + w_2} \quad (4-58)$$

Under the differential steering configuration, the velocity of each servo wheel is then set by two normalised command setpoint signals, combined within the mixer block of Figure 4-7. The first, $c_{sp} \in [-1, +1]$, is a common command setting desired speed about the pipe circumference and moving the vehicle in the $\{B\}$ y-axis. The second, $d_{sp} \in [-1, +1]$, is a differential command governing desired angular velocity about the $\{B\}$ z-axis. Accordingly, the vector of normalised signals used to drive the vehicle servo motors, $\mathbf{m}_{sp} \in [-1, +1]$, is set as follows, with each row driving the servo of the same index.

$$\mathbf{m}_{sp} = \begin{bmatrix} c_{sp} - d_{sp} \\ c_{sp} + d_{sp} \end{bmatrix} \quad (4-59)$$

Kinematically, these are related to the translational velocity of each wheel by the servo proportional constant k_s . Assumed approximately equal in each servo, this accounts for the radius of the tyres and the scaling factor between the normalised servo input signal and its rotational speed.

$$\begin{aligned} v_1 &= k_s (c_{sp} - d_{sp}) \\ v_2 &= k_s (c_{sp} + d_{sp}) \end{aligned} \quad (4-60)$$

Using these command signals within (4-57) and (4-58), it may be shown that the 2D crawling kinematics of the vehicle are described as below. Note that here the simplifying assumption is made that w_1 and w_2 are approximately equal and both have value w .

$$v_y = k_s c_{sp}, \quad \omega_z = \frac{k_s}{w} d_{sp} \quad (4-61)$$

Thus, the resultant vehicle motion components are linearly proportional to the normalised command signals c_{sp} and d_{sp} , and crawling may be easily controlled by either a pilot setting these values, or an appropriately developed feedback scheme. Surface relative motion in 2D is thereby enabled as in the case of wheeled mobile robots on a planar surface.

In the context of the pipe geometry, the vehicle may traverse its surface by manipulation of v_y and ω_z . By controlling ω_z , the yaw angle of the vehicle body frame $\{B\}$ relative to the pipe frame $\{P\}$ may be altered within a small range of approximately $\pm 5^\circ$. This is enabled by the enclosure of a diameter slightly larger than the pipe by the vehicle legs, permitting adjustment of vehicle orientation marginally beyond perfect axial alignment until diagonally opposite legs contact the surface and constrain further rotation. Such a small orientation offset is well within design tolerance of the thrust-support algorithm described in Section 4.4.2 and has minimal influence on stability. As such, redirection of v_y using ω_z may create motion components outside the purely circumferential $\{P\}$ y-z plane and so be used to control motion along the $\{P\}$ x-axis.

NDE activities involving traversal of the pipe length are therefore possible using a helical path with a pitch length set by the vehicle yaw angle, supporting surface area inspection. Additionally, this provides capability for fine-scale position correction in 2D about the pipe surface while in contiguous contact, permitting repeat visitation and monitoring of specific target locations.

4.5 Experimental Assessment Methodology

This section details the methods used to assess the utility of the hybrid UAV-crawler vehicle for practical NDE. Accordingly, consideration is given to characterisation of both spatial positioning and imaging quality criteria. Data processing methods used to prepare photogrammetric models and unrolled orthoimages of the pipe surface analysis are also examined ahead of their results in Section 4.6.

4.5.1 Motion Tracking

Experimental work characterising aspects of system performance relating to physical surface crawling and flight access is proposed to validate the interaction control algorithm and novel hybrid vehicle as an NDE deployment method. This is performed within an indoor laboratory flight volume comparable to that described in Section 3.3.1, allowing testing conditions to be tightly controlled for isolation and assessment of relevant performance aspects.

Within this facility offboard position feedback is provided using a Vicon tracker system with twelve T160 cameras [351]. Similar systems are commonly used to assess motion control of ground-based and aerial robotic systems, here providing an external means of measuring vehicle position and orientation at a frequency of 100 Hz, as in the assessment of Section 3.4. It functions in place of a GNSS system within the indoor laboratory and, if desired, may be purposefully degraded to mimic GNSS accuracy characteristics. During these experiments the Vicon tracker provides pose feedback to the flight controller enabling the vehicle to autonomously hold position and reducing pilot skill requirements. Note, however, that during thrust-supported crawling the vehicle requires only the attitude estimate from the internal IMU to maintain a stable interaction and the Vicon data is used purely for evaluation purposes.

Vicon pose measurements of both the vehicle and target asset are thus recorded throughout each trial in support of post-processed motion analyses. While the Vicon tracker global pose measurement accuracy can show Mean Absolute Error (MAE) on the order of 4 mm [330], typical pose measurement precision for static readings are specified at 0.007 mm RMS deviation, increasing to 0.142 mm for a moving object

[352]. The system is therefore well suited to precisely measure body stability relative to a static position about the test asset circumference, readily detecting slip, and profile both flight and crawling motion, where accuracy is of acceptable scale relative to vehicle motion. Note that high frequency vibrations common to multirotor vehicles are not considered here owing to limitations of the Vicon tracker sampling rate and their limited impact on the contact interaction.

4.5.2 Imaging Performance

Remote visual NDE requires the capture of high-quality images capable of accurately resolving fine surface detail to detect and track incipient flaws and satisfy inspection standards. This spatial resolution is influenced by a range of compound effects including pixel sizing and layout within the image sensor, lens distortions, variation in lighting conditions, focal blur, and motion blur [234]. Of these, aspects relating to the camera sensor and lens may be well controlled by appropriate hardware specification. Inclusion of an LED array within the vehicle similarly provides means to set lighting parameters, while the nature of the surface crawling contact ensures a well-known and consistent focal point relative to the vehicle. Thus, within the context of thrust-supported deployment, camera motion during image capture is identified as the most influential of these aspects.

Experimentation therefore examines the impact to visual NDE data quality incurred due to the vehicle motion profile when holding static position, correlating physical motion characteristics with image spatial resolution to identify causal effects. This is intended to provide an operational performance specification that may be used in evidence of suitability of the hybrid vehicle for practical visual inspection processes.

To this end, the quality of images captured in static poses about the full sample circumference is assessed and numerically quantified via United States Air Force (USAF) MIL-STD-150A procedure [353], as recommended by HOIS industrial guidance for UAV visual inspection system qualification [354]. This testing process is built around target artefacts such as the one depicted in Figure 4-10.

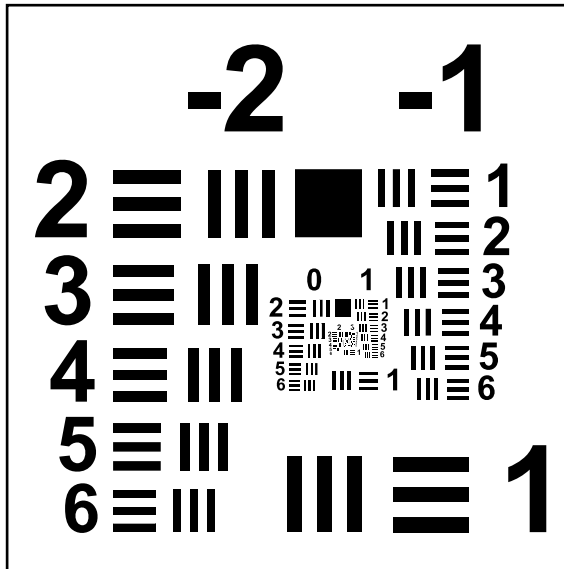


Figure 4-10: Image resolution artefact used under USAF MIL-STD-150A spatial resolution assessment procedure. This image contains line pairs from group -2 down to group 7, extending well beyond the limits of commercial printer resolution. If reproduced as shown on A4 paper, the image will retain 1:1 scale with the standard artefact, with a bounding box of side length of 75 mm.

This artefact contains a number of line-pair elements, each here denoted by an element number, n_e , and consisting of three pairs of black and white lines of specific width, set in both horizontal and vertical orientation. These are further arranged into groups, denoted n_g , each containing six elements. A feature may then be referred to using its group and element number in the format n_g-n_e , e.g. group 2 element 3 is denoted 2-3. The width of each line pair is then specified such that the image spatial resolution needed to observe the element, here denoted R_{lp} , may be determined by the following equation, expressing the result in line-pairs per millimetre, i.e. lp/mm.

$$R_{lp} = 2^{\left(n_g + \frac{(n_e - 1)}{6}\right)} \quad (4-62)$$

The spatial resolution of an image showing the target artefact may then be quantified by manual identification of the smallest element for which all included line segments are visually distinct. For example, an image where element 2-3 is visible will have spatial resolution of at least 5.657 lp/mm, corresponding to lines with width of 88.39 μm .

An assessment based on the USAF MIL-STD-150A procedure therefore provides direct indication of the minimum resolvable feature size for practical NDE amid the complex interaction of the imaging system parameters and any degressive influences.

4.5.3 Geometry & Surface Texture Inspection

As discussed in detail in Appendix A, photogrammetric reconstruction of a target scene provides means to accurately measure asset geometry and track any changes over time, also providing spatial reference for other NDE data. Versus the approximate knowledge of the nominal pipe diameter that may exist following initial construction, a reconstructed mesh then provides substantial additional information. Its analysis can enable measurement of complex spatial features, detecting and assessing any pipe ovality, deformation, sag or bending, and allowing quantitative tracking of defects like corrosive blistering that may develop over its service life. The process however requires high quality source images, captured without significant distortion, and in a pattern ensuring full surface coverage with overlap between images sufficient to allow their matching and alignment.

The feasibility of this process using images captured by the hybrid vehicle in continuous helical motion around a pipe sample is therefore examined using commercially available photogrammetry software, Agisoft Metashape [355], commonly used in UAV-based NDE processes [234]. This testing provides practical evidence of functionality in the context of a highly self-similar surface texture and the motion profile of a surface crawling vehicle under thrust-based adhesion. If successful, results will demonstrate the ability of the vehicle to provide complete coverage inspection of cylindrical assets, yielding photorealistic models. Including surfaces not visible from the ground, the reconstructions would then be well suited for use in Building Information Modelling (BIM) processes or as digital twins tracking structural health over asset lifespan.

Furthermore, assessment of the hybrid vehicle utility for remote visual inspection is considered via a photogrammetric process created to produce a single, high-resolution, orthophoto image of the contiguous exterior surface of the inspected area. This has the potential to ease the data examination process, requiring little modification to current

NDE practices using 2D images, while also providing intrinsic localisation of features relative to more distinct reference points, and removing context confusion due to the self-similarity common in industrial structures. Production of this orthoimage is then approached via the unwrapping of the reconstructed mesh, retaining its benefits as a method for complex geometric assessments while ensuring that the features depicted in the resulting image are not distorted or warped by any deviation from the ideal cylindrical case.

Post-processing of the reconstructed mesh is then required to produce a cohesive unwrapped surface texture image useful for NDE. This is necessary as the mesh texture exported from the SfM process is commonly fractured and discontinuous, as shown in Figure A-8. As such, the linkage between x-y-z spatial coordinates and u-v texture image coordinates, termed the “UV-map”, does not retain similar relative spacing between spatial points and their locations in the texture image. To view the full surface of the pipe in a single cohesive image, this must be corrected. The method used to do so here is detailed as follows, supporting assessment of the hybrid vehicle for surface texture inspection.

4.5.3.1 Coordinate Frame Alignment

Photogrammetric reconstruction of the 3D mesh is performed without awareness of the camera poses in $\{W\}$, and so the resulting textured mesh arbitrarily forms and isolates the object with undefined orientation in its reference frame. This reconstructed mesh frame is denoted $\{M\}$, as depicted in Figure 4-11(a). Expression of the reconstruction in a known, well-aligned, coordinate frame enables geometry manipulation and is therefore necessary to extract relevant NDE surface image.

Accordingly, a cylinder-aligned frame, $\{C\}$, is determined using the vertices of the reconstructed pipe-section mesh, as are expressed in $\{M\}$ and so denoted ${}^M\mathbf{v}_i$. These vertices are fit to a cylindrical primitive by importing the reconstructed mesh object into MATLAB and using a non-linear, least-squares algorithm provided within the PROTO toolbox [356], [357], to minimise separation between the vertex points and the candidate cylinder.




Figure 4-11: Fitting of the reconstructed pipe mesh to a cylinder-aligned coordinate frame, $\{C\}$. (a) The initial photogrammetric output mesh is arbitrarily aligned in its reference frame, $\{M\}$. (b) Diagram of the method defining $\{C\}$ relative to the surface mesh. Note the cylinder mesh faces are shaded top to bottom to aid visualisation.

Following optimisation, a best-fit cylinder is defined in $\{M\}$ by a unit vector describing the orientation of its central axis, \hat{A} , an arbitrary point on this axis, ${}^M\mathbf{v}_A$, and the least-squares radius, \bar{r} . These parameters, however, are insufficient to define $\{C\}$ in isolation. An origin point with a known location relative to the pipe mesh and full constraint of its orientation are still required. These are determined as depicted in Figure 4-11(b).

To this end, the origin of $\{C\}$ expressed in $\{M\}$, denoted ${}^M\mathbf{c}_O$, is set at the midrange point along the cylinder axis, determined by considering the signed distance along \hat{A} of each vertex relative to ${}^M\mathbf{v}_A$, denoted d_i , as below.

$$d_i = \text{proj}_{\hat{A}}(\overline{{}^M\mathbf{v}_i - {}^M\mathbf{v}_A}) \quad (4-63)$$

The midrange signed distance from ${}^M\mathbf{v}_A$, termed λ , is then found.

$$\lambda = \min(d_i) + \frac{\max(d_i) - \min(d_i)}{2} \quad (4-64)$$

Consequently, the line of the cylinder axis lastly yields the definition of ${}^M\mathbf{c}_O$ at the mesh midrange point per the equation below.

$${}^M\mathbf{c}_O = {}^M\mathbf{v}_A + \lambda\hat{A} \quad (4-65)$$

Orientation of $\{C\}$ relative to $\{M\}$ is next set by constructing the rotation matrix ${}^M_C\mathbf{R}$ using appropriate basis vectors.

$${}^M_C\mathbf{R} = [{}^M_C\hat{\mathbf{x}} \quad {}^M_C\hat{\mathbf{y}} \quad {}^M_C\hat{\mathbf{z}}] \quad (4-66)$$

The first of these, ${}^M_C\hat{\mathbf{x}}$, expressing the $\{C\}$ x-axis in $\{M\}$ is trivially defined by the best-fit cylinder axis.

$${}^M_C\hat{\mathbf{x}} = \hat{\mathbf{A}} \quad (4-67)$$

The second basis vector defining the z-axis, ${}^M_C\hat{\mathbf{z}}$, is set such that the $\{C\}$ z-axis opposes the direction of gravity in the reconstruction frame, $\{M\}$. Doing so, the clockface angle of a point on the pipe surface given in $\{C\}$ will match the angle of the same point in the spatial pipe-relative frame, $\{P\}$, defined above, facilitating intuitive comparison. The $\{W\}$ z-axis in $\{M\}$, ${}^M_W\hat{\mathbf{z}}$, defining the direction of gravity is found as follows

$${}^M_C\hat{\mathbf{z}} = -{}^M_W\hat{\mathbf{z}} = -{}^M_I\mathbf{R} \quad {}^I_B\mathbf{R} \quad {}^B_W\mathbf{R} \quad {}^W\hat{\mathbf{z}} \quad (4-68)$$

This uses a number of relationships between coordinate frames to link rotational transforms. Starting from a North-East-Down $\{W\}$ z-axis, ${}^B_W\mathbf{R}$ expresses its orientation relative to the vehicle body frame, $\{B\}$, estimated at the instant of image capture using the internal IMU data as in Algorithm 1. By design the camera is rigidly mounted aboard the vehicle such that its image plane normal direction aligns with the $\{B\}$ z-axis. This mechanically defines orientation of the vehicle body relative to the imaging coordinate frame $\{I\}$, setting ${}^I_B\mathbf{R}$. The camera extrinsic estimate computed during SfM image alignment then defines the orientation of $\{I\}$ relative to the mesh frame $\{M\}$, yielding ${}^M_I\mathbf{R}$. As both ${}^B_W\mathbf{R}$ and ${}^M_I\mathbf{R}$ vary with time, their values are taken simultaneously at the instant of capture for the first image used in the photogrammetric reconstruction. The SfM optimisation stage and the IMU data fusion used in the vehicle pose estimate ensure these provide a robust transform with minimal spurious error, granting a reliable estimate of the direction of gravity.

The third axis, ${}^M_C\hat{\mathbf{y}}$, is then found by the vector cross product, completing the basis vectors to define ${}^M_C\mathbf{R}$, such that the top of the reconstructed pipe aligns with its physical counterpart.

$${}^M_C\hat{\mathbf{y}} = {}^M_C\hat{\mathbf{x}} \times {}^M_C\hat{\mathbf{z}} \quad (4-69)$$

Pose of the cylinder-aligned frame $\{C\}$ is then fully described relative to the reconstruction frame $\{M\}$ by the transform matrix, ${}^M_C\mathbf{T}$, composed as follows

$${}^M_C\mathbf{T} = \begin{bmatrix} {}^M_C\mathbf{R} & {}^M_C\mathbf{O} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \quad (4-70)$$

Locations of reconstructed mesh vertices may then be expressed in $\{C\}$ using the transformation matrix ${}^C_M\mathbf{T}$, i.e. the inverse of ${}^M_C\mathbf{T}$, such that

$${}^C\mathbf{v}_i = {}^C_M\mathbf{T} {}^M\mathbf{v}_i \quad (4-71)$$

This method thereby defines a practical coordinate system relative to the reconstructed surface mesh of the pipe to facilitate further processing of NDE data captured in a single inspection. If desired, extension to align multiple reconstructions from different inspections may be facilitated by expressing each $\{C\}$ frame relative to a globally consistent absolute coordinate system. This may be derived from the estimated pose of the vehicle during data capture, or established from NDE images of the same location in each dataset as in the camera alignment stage of SfM photogrammetry, discussed in Appendix A

4.5.3.2 Unrolling the Surface Texture

Working within the aligned frame, $\{C\}$, unrolling of the reconstructed mesh texture while preserving the spatial adjacency of points on its surface becomes practical. By further non-linear transformation, similar to conversion between cartesian and cylindrical coordinates [358], the mesh vertex positions may now be given in an unrolled imaging coordinate frame, $\{U\}$, as depicted in Figure 4-12. The reconstructed pipe section can thereby be represented as quasi-planar object and subsequently rendered to yield a radial orthoimage of its surface.




Figure 4-12: Unrolling of the reconstructed pipe mesh texture. (a) Spatial attributes of the pipe in the cylinder-aligned frame $\{C\}$. (b) Mapping of these properties to the unrolled image frame $\{U\}$. Note unwrapped mesh faces preserve their top to bottom shading, visualising the consistency of face UV-mappings through this transform process, but now form a planar surface.

Transformation of the object vertices from $\{C\}$ to $\{U\}$ is derived based on components of their 3D position. The clockface angle of each point in $\{C\}$, θ_i , is first determined via trigonometry. Note that an offset of 90° is applied to position the zero point at the top of the pipe per the clockface angle definition used previously.

$$\theta_i = \text{atan2}({}^c v_{i,z}, {}^c v_{i,y}) - 90^\circ \quad (4-72)$$

This then yields the signed arc length, s_i , from the top of the pipe section, around its best-fit radius, to the vertex point.

$$s_i = \bar{r}\theta_i \quad (4-73)$$

The radius of the specific vertex from the best-fit axis, r_i , is its Euclidean distance from the origin of $\{C\}$ in the y - z plane.

$$r_i = \sqrt{({}^c v_{i,y})^2 + ({}^c v_{i,z})^2} \quad (4-74)$$

Using (4-72), (4-73) and (4-74), the vertex location in the unrolled frame, ${}^U \mathbf{v}_i$, is then summarily expressed by non-linear transformation from $\{C\}$, as follows.

$${}^U \mathbf{v}_i = \begin{bmatrix} {}^c v_{i,x} \\ s_i \\ r_i \end{bmatrix} \quad (4-75)$$

Note that x-axis position of each point remains consistent across $\{C\}$ and $\{U\}$. Additionally, the vertex point radius in $\{C\}$ is mapped to the $\{U\}$ z-axis and so appears out of plane in Figure 4-12(b). This ensures all geometry information is retained and that radial features or deformations remain detectable in $\{U\}$.

Versus the standard cylindrical coordinate expression using θ_i , formation of the $\{U\}$ y-axis using arc length preserves feature aspect ratio in the unwrapped geometry. By calculating s_i with the best-fit radius \bar{r} , a further source of distortion is removed. If the point radius r_i were used, the arc length of radially higher features would increase disproportionately versus nearby radially-lower features, causing their horizontal shift in $\{U\}$ and distorting the texture image. Expression as in (4-73) ensures that each point in the image appears as if viewed radially from directly above it and so allows accurate defect sizing from the resulting orthoimage without camera perspective effects.

A final spatial processing step, also modifying the object UV-map, is taken to account for triangular mesh faces split across the seam at θ_i of $\pm 180^\circ$. Such faces are first detected where any θ_i differs from that of another vertex in the face by more than $\pm 180^\circ$. To prevent the face being stretched across the horizontal span of the image, a new vertex is introduced. This vertex replaces the one of the three comprising the face that lies on the minority side of the divide. It is calculated by offsetting θ_i in a copy of the minority vertex by $\pm 360^\circ$ to reposition it near the majority vertices. Addition of the new vertex ensures adjacent faces sharing the minority vertex are not disturbed. As shown in Figure 4-12(b), this new vertex thereby extends the surface arc length marginally outside the previous $\pm \pi \bar{r}$ bounds, also moving the unrolling seam to lie along the edges of the faces on this line.

With the unrolling process complete, the reconstructed mesh may be imported into Blender [359] and rendered with an orthographic perspective from a viewpoint in the $\{U\}$ z-axis. This fills in the texture for each mesh face, referencing its location in the fractured photogrammetric source texture via the UV-map and applying the image patch at the corrected pose in the unrolled geometry. The quasi-planar surface is resampled and interpolated to populate the pixel grid of the resultant texture image, yielding a well-structured, contiguous surface depiction similar to Figure 4-12(b).

Using this process, the vehicle may capture visual NDE data to provide a single image surface view of the entire pipe-section of interest. The effectiveness of this method for inspection is considered in practical experiment as below.

4.6 Results

Quantitative validation of the vehicle design and interaction control is conducted within an indoor laboratory flight volume as described above. Results from multiple experiments² examine key system performance aspects concerning vehicle interaction stability, its motion profile when crawling around the pipe circumference, its ability to access target locations from flight, and, ultimately, its ability to conduct immediate proximity visual NDE.

4.6.1 Interaction Stability

Firstly, assessment considers ability of the vehicle ability to maintain sequential positions at 45° increments around the circumference of a 12.75-inch outer diameter, schedule 80, steel pipe sample, as depicted in Figure 4-13. This mimics common NDE points under recommended pipe inspection practice [66]. In a continuous test under adaptive thrust support, the pilot sets wheel servo speed to move the vehicle between these points and halt as desired.

To quantify interaction stability, Vicon pose data is examined when the vehicle is commanded to be stationary: that is, having a normalised speed control signal under $\pm 2.5\%$, per a small dead-zone. The vehicle remains in this state for a minimum of 2.5 s to allow for steady-state position data capture. Motion relative to the mean pose during these times is then characterised as in Table 4-3.

² Video of these experiments is available via https://youtu.be/Z9I9yj6_NXM.

TABLE 4-3: STABILITY OF THE VEHICLE AROUND THE PIPE

N	θ [°]	$\bar{\theta}$ [°]	$\overline{\Delta X}$ [mm]	$\overline{\Delta Y}$ [mm]	$\overline{\Delta Z}$ [mm]	\bar{L} [mm]	L_{95} [mm]
500	0	+3.06	0.03	0.04	0.06	0.09	0.13
347	+45	+48.04	0.02	0.27	0.28	0.39	1.27
345	+90	+91.14	0.05	0.07	0.91	0.92	1.67
367	+135	+134.86	0.14	0.61	0.52	0.83	2.36
384	180	+178.93	0.04	0.12	0.20	0.27	1.03
344	-135	-135.55	0.13	0.77	0.67	1.03	2.10
270	-90	-91.93	0.06	0.07	0.08	0.14	0.23
341	-45	-37.79	0.10	0.04	0.06	0.14	0.23
500	0	+1.67	0.04	0.23	0.05	0.25	0.62

Key: N , Number of samples. θ , target clockface angle. $\bar{\theta}$, Mean clockface angle. $\overline{\Delta X}$, $\overline{\Delta Y}$, $\overline{\Delta Z}$, Mean absolute position deviation in $\{P\}$ x,y,z axes. \bar{L} , mean position deviation vector length. L_{95} , 95th percentile deviation length.

Position deviation is primarily tangential, showing successful adhesion. Differences in behaviour noted at similar clockface angles on opposing sides are attributed to motion transients following a stop command. When moving upwards, gravity aids deceleration, reducing stopping distance and vice versa. After the transient, minor position drift arises due to craft roll-axis vibration, reducing wheel contact force and permitting brief instants of tire slip. Worst case deviation occurs at $\pm 135^\circ$, correlating with the reduced leg torque near θ_C , as in Figure 4-6. Higher roll stabilizing torque may mitigate this but increases energy use. Algorithm 1 is thereby validated, showing successful stabilization of the vehicle about the pipe.




Figure 4-13: Sequential image series showing the vehicle holding position at each 45° station about the pipe. Temporal progression runs from left to right, top to bottom. Annotations denote the targeted clockface angle about the pipe long axis.

4.6.2 Imaging Performance

Here spatial imaging resolution is assessed in the context of full circumferential pipe inspection.

A total of eight USAF test artefacts, as depicted in Figure 4-10, are affixed at 45° increments around the pipe starting from a clockface angle of 0°. Repeating the motion path used in Section 4.6.1, the vehicle is piloted between target locations under automatic adaptive thrust support, pausing at each for a minimum of 2.5 s as previously. Images of these artefacts and the pipe surface are captured continuously at

a rate of over 16 Hz. Analysis of the control signals is then used to identify images taken while the vehicle was static, having a speed command signal less than $\pm 2.5\%$.

Images typical of those recorded at each position about the circumference are given in Figure 4-14. These are cropped and enlarged, without any further processing, to show the finest scale features at the centre of the USAF artefact in Figure 4-15, allowing clear identification of the smallest visible line group. Here, each clearly shows the 2-2 group, with borderline visibility of the 2-3 group, corresponding to spatial resolution between 4.49 lp/mm and 5.04 lp/mm per equation (4-62).

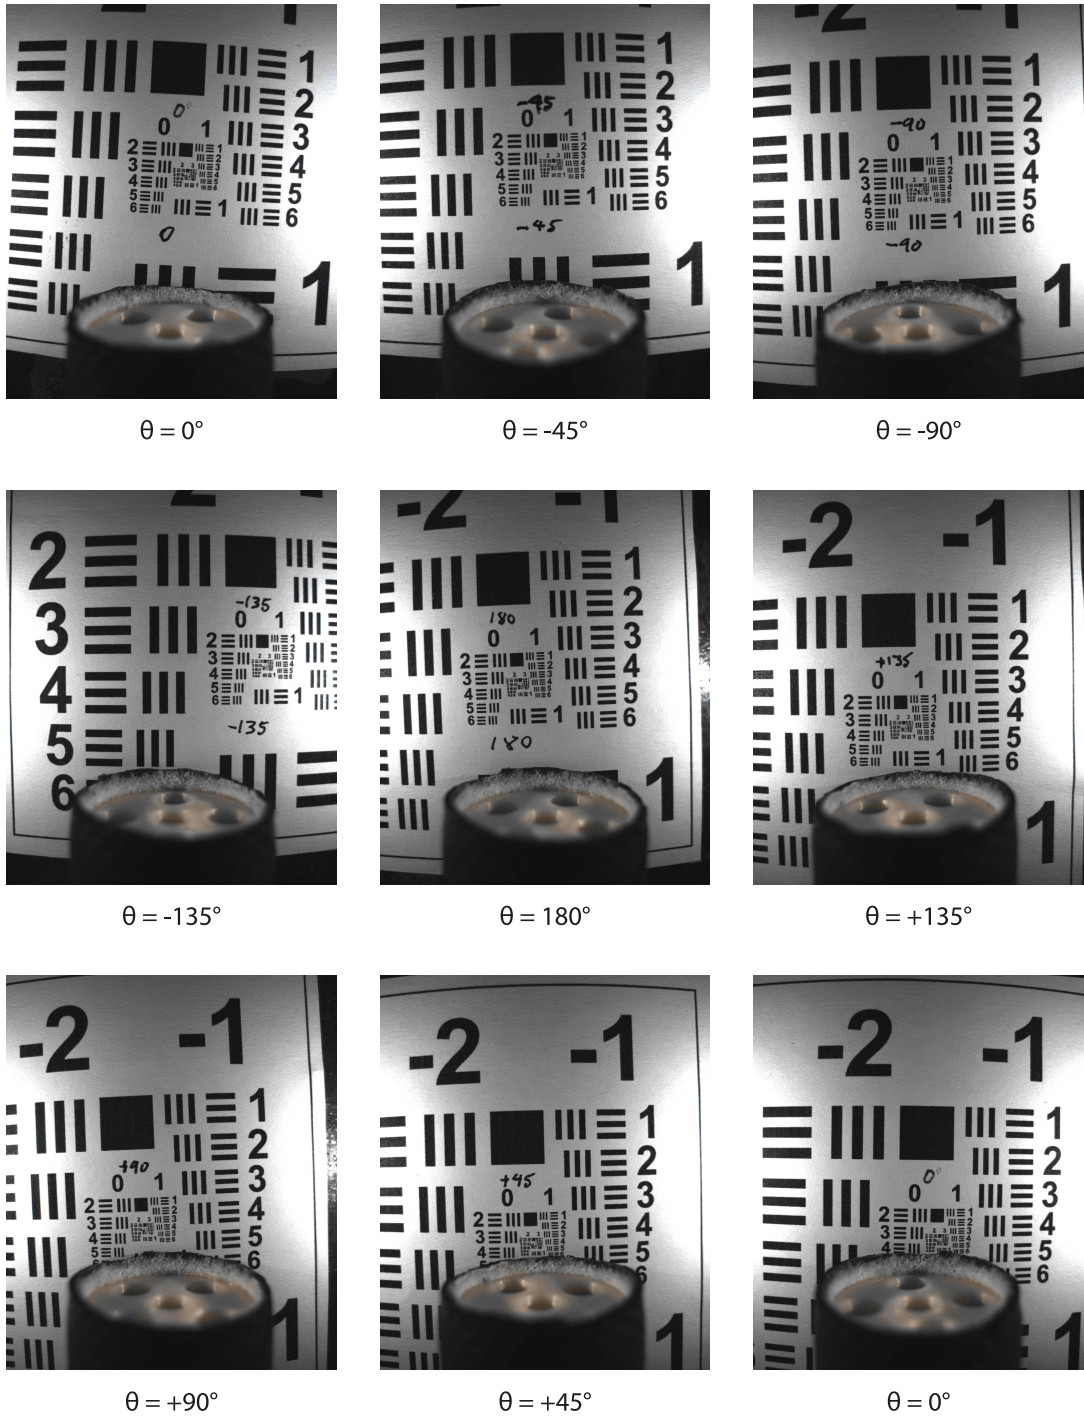


Figure 4-14: Unprocessed full-frame images of the USAF artefact captured by the vehicle during static pose holding at 45° intervals around the pipe circumference. Annotations denote the position where each artefact is centred by their clockface angle. For scale, the largest visible black square in these images has side length of 9 mm.

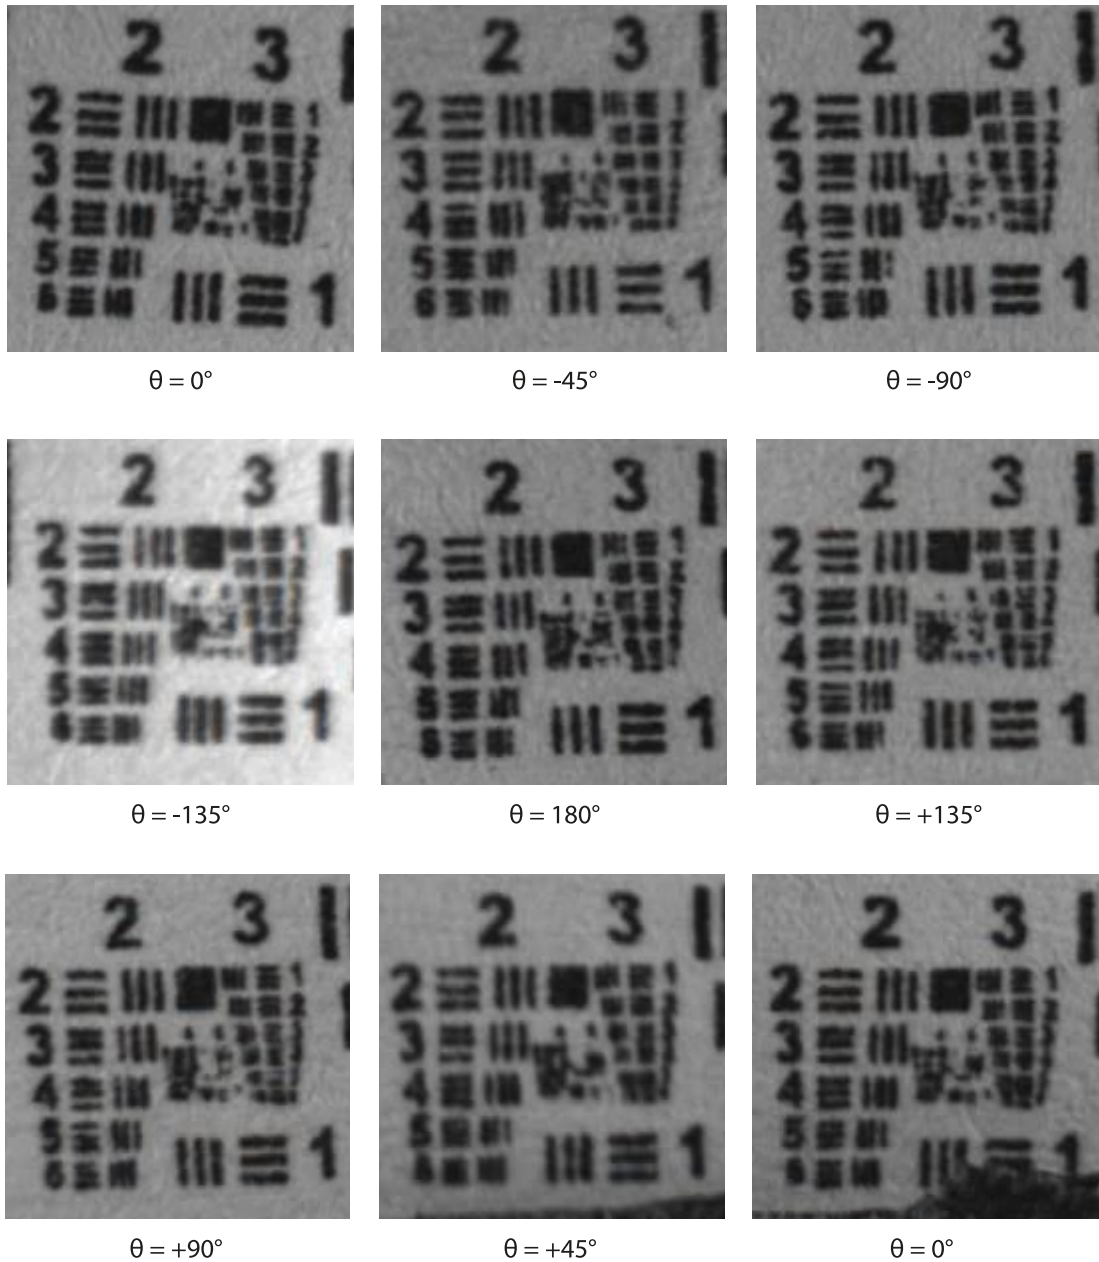


Figure 4-15: Images of the central region of the USAF artefact cropped and enlarged from those of Figure 4-14 without re-sampling or any further processing steps to enhance resolution. Annotations give the clockface angle of each artefact. These images may be used to identify the smallest line-pair feature visible in each position. For scale, the largest visible black square in these images has side length of 0.56 mm.

Further quantitative evaluation is conducted by application of MIL-STD-150A analysis across all 514 images captured while the vehicle was stationary, including those shown in Figure 4-14 and Figure 4-15. Accordingly, imaging performance obtained at each station is numerically evaluated via the statistical metrics of Table

4-4. These describe the distribution and variation of image resolutions, establishing numeric confidence bounds for practical deployment of the NDE method. Visualisation of these data as a time series, as in Figure 4-16, provides additional insight into the factors influencing this performance, as examined below.




Figure 4-16: Resolution of images captured while in static position at 45° increments around the pipe. Each blue dot is one image. Orange highlights indicate the times when the vehicle was commanded to be static. Annotations show the clockface angular position of each USAF artefact.

TABLE 4-4: IMAGE RESOLUTION WHEN STATIONARY AROUND THE PIPE

N	θ [°]	$\bar{\theta}$ [°]	\bar{R}_{lp} [lp/mm]	σ_{lp} [lp/mm]	R_{95} [lp/mm]
92	0	-0.65	4.550	0.172	5.040
43	-45	-45.34	3.724	0.908	5.040
45	-90	-93.89	4.563	0.189	5.040
52	-135	-146.75	4.116	0.374	4.490
53	180	176.55	4.475	0.201	4.957
42	+135	132.48	4.531	0.478	5.040
52	+90	89.53	4.564	0.352	5.040
57	+45	45.90	3.645	0.667	4.490
78	0	-1.79	4.554	0.374	5.040

Key: N , Number of samples. θ , Artefact clockface angle position. $\bar{\theta}$, Mean vehicle clockface angle position while static. \bar{R}_{lp} , Mean resolution across images in each location. σ_{lp} , Standard deviation of image resolution in each location. R_{95} , 95th percentile maximum resolution of images in each location.

Collectively, the data of Table 4-4 show a mean resolution of 4.330 lp/mm with a standard deviation of 0.560 lp/mm. However, the nature of visual NDE is such that if multiple redundant images exist only the highest quality image of a given consistent

target need be considered. A more practical figure of merit is found in the 95th percentile maximum spatial resolution which reduces the negative influence of unused redundant images while still accounting for some empirical variation. Accordingly, an average 95th percentile resolution of 4.915 lp/mm is shown here across all trial image data.

Considering variation around the pipe, it is apparent that the images captured at 0° clockface angle yield highest, most consistent, performance. This occurs at the top of the pipe, where the vehicle motors are not active and it rests under purely passive support, eliminating any vibration effects due to propeller motion. Performance there thus defines an upper bound for the spatial resolution attainable by the imaging equipment within this assessment, free of any degressive effects introduced by motion of the thrust-supported platform.

Such vehicle effects are most prominent at the $\pm 45^\circ$ positions, where the highest resolution variance within a single group is observed. In spite of the high stability retained by the vehicle body per the results of Table 4-3, the images captured there exhibit cyclically variable distortion akin to motion blur, as indicated by their temporal sequence in Figure 4-16. This is therefore attributed to vibration in the rigid camera mount, induced as the increasing propeller support force brings their rotation speeds near a structural resonance frequency. Beyond this point, further motor speed increase moves their rotation frequency past resonance and significantly reduces the effect, reducing variance in image resolution at other points under the pipe.

Another influence of the mobile platform is observed at the -135° point where Figure 4-16 highlights a difference in behaviour versus the same point on the opposite side of the pipe. In this case, the pilot positioned the vehicle approximately 10° off centre from the USAF artefact. While the target centre remains visible in the image of Figure 4-14, curvature of the pipe surface sees it positioned further from the image focal point. This contributes additional focal blur above the degressive effects of the motion profile at $\pm 135^\circ$ shown in Table 4-3, leading to a small reduction in average spatial resolution.

A final effect of the mobile platform is also made apparent in Figure 4-16 where, on visiting some stations, the first images captured after arrival exhibit a marked reduction

in resolution versus the location average. This is attributed to motion transients in the instants following the reduction of the vehicle speed command signal below the 2.5 % threshold for static position holding. The vehicle takes a small but distinct time to decelerate, during which time motion blur of the finest scale features is exaggerated. Following this transient, the image performance is far more consistent and in better agreement with the baseline established atop the pipe.

Investigation into system visual NDE performance has thereby shown reliable performance for point screening of the pipe surface through the capture of multiple redundant images, with effects impacting their spatial resolution identified and quantified to support informed practical utilisation.

4.6.3 Landing Repeatability

Here assessment profiles the ability of the system to enter contact with the target asset from the air, considering landing accuracy for precision NDE delivery to a specific location.

Landing may be conducted manually, utilising the ability of the UAV to exert lateral thrust without body roll to finely control its position above the pipe axis, but for the purposes of this assessment the flight controller uses Vicon pose feedback to eliminate human repeatability factors. Using markers affixed to its surface, the real-time position of the pipe is then provided by the Vicon tracker, functioning as described in Section 4.5.1. While other works have demonstrated that similar pipe-relative positioning may be provided via onboard monocular vision [360], depth cameras [361], or lidar [362], using Vicon feedback minimises the influence of uncertainty in the target pose estimation and permits focussed assessment of the underlying performance characteristics of the novel hybrid vehicle. Note also that, as in previous experiments, no offboard pose feedback data is used once the hybrid vehicle has entered contact with the pipe and switched to interaction control mode.

To conduct the automated landing manoeuvre assessed here, the UAV descends from an approach pose 0.75 m above the pipe central axis at up to 0.1 m/s, targeting a central location marginally below the asset surface under pose and velocity regulation by the

flight controller described in Section 4.4.1. On contact, interaction control mode is enabled and the vehicle rests statically atop the pipe for no less than 2 s. Take-off reverses this, returning to the approach pose. The cycle is repeated in two sets of five landings to aid result confidence.

Landing was successful in all trial instances. Figure 4-17 plots the target relative landing locations in the $\{P\}$ x-y plane. Table 4-5 gives their numerical distribution.




Figure 4-17: Landed vehicle positions versus target location in $\{P\}$, which forms origin of this plot. Numerical indices denote the temporal sequence of landings.

TABLE 4-5: LANDING LOCATION DEVIATION FROM TARGET

Axis	Min [mm]	Mean [mm]	Max [mm]	σ [mm]	L_{95} [mm]
X	3.12	11.99	39.47	12.53	39.47
Y	-6.65	-1.39	2.70	2.93	6.65
Direct	3.29	12.68	39.52	12.20	39.52

Key: *Direct*, the Euclidean distance from the target to the point of landing. σ , Position standard deviation across landing trials. L_{95} , 95th percentile deviation length.

Here, the MDT rotors permit direct thrust-based yaw and $\{B\}$ y-axis motion control, yielding strong regulation of craft heading and position transverse to the pipe axis amid turbulent near-surface aerodynamics. Consistent roll angle removes the discontinuous lateral ground effect components otherwise caused by downdraft reorientation relative to the pipe, improving landed pose consistency in the $\{P\}$ y-axis. Larger deviation in the x-axis is instead governed by pitch axis inertia, limiting responsiveness but not affecting landing success.

Surface crawling provides means for further fine corrections after initial contact, allowing the vehicle to travel large distances to access the target area via flight before refining its position while in continuous contact crawling to repeatably access a precise target location. Considering landing accuracy and the camera image area, performance is readily sufficient for direct delivery of NDE to features at the scale of interest in surface corrosion inspection [300].

4.6.4 Fly-Crawl-Fly Inspection

To assess overall system feasibility for contact and immediate proximity NDE, a full fly-crawl-fly process is examined, with flight trajectory plotted in Figure 4-18.

Here, the vehicle is piloted to a position near the target pipe before autonomously landing atop it, as before. Control then transitions to interaction mode. The pilot sets a constant scanning speed, and the vehicle completes a full circumnavigation under adaptive thrust support before re-entering flight and returning to its start position on the ground.

NDE images are captured continuously during this motion and streamed to the operator, permitting flexible investigation. As in the example of Figure 4-19, the weld cap and pipe surface features are clearly visible without significant motion blur, indicating successful rolling scan imaging.

Scan motion consistency is further examined to assess dynamic performance of Algorithm 1 and utility for general NDE applications. Figure 4-20 plots the vehicle scan speed under constant command setpoint and without applied differential steering. Speed is derived from the Vicon pose data by numerical differentiation of each axis, median filtering out high-frequency numerical noise with a 0.25 s moving window, and returning the Euclidean vector norm. As such, a small sinusoidal disturbance is observed about the 0.075 m/s mean speed. Maximum speed of 0.106 m/s occurs near θ of -136° , travelling downwards, whereas minimum speed of 0.061 m/s occurs at $+107^\circ$, traveling upwards, suggesting gravitational influence. However, speed standard deviation under 0.012 m/s, in concert with the steady roll angle change of Figure 4-18, shows sufficient consistency for contact-based scanning NDE.

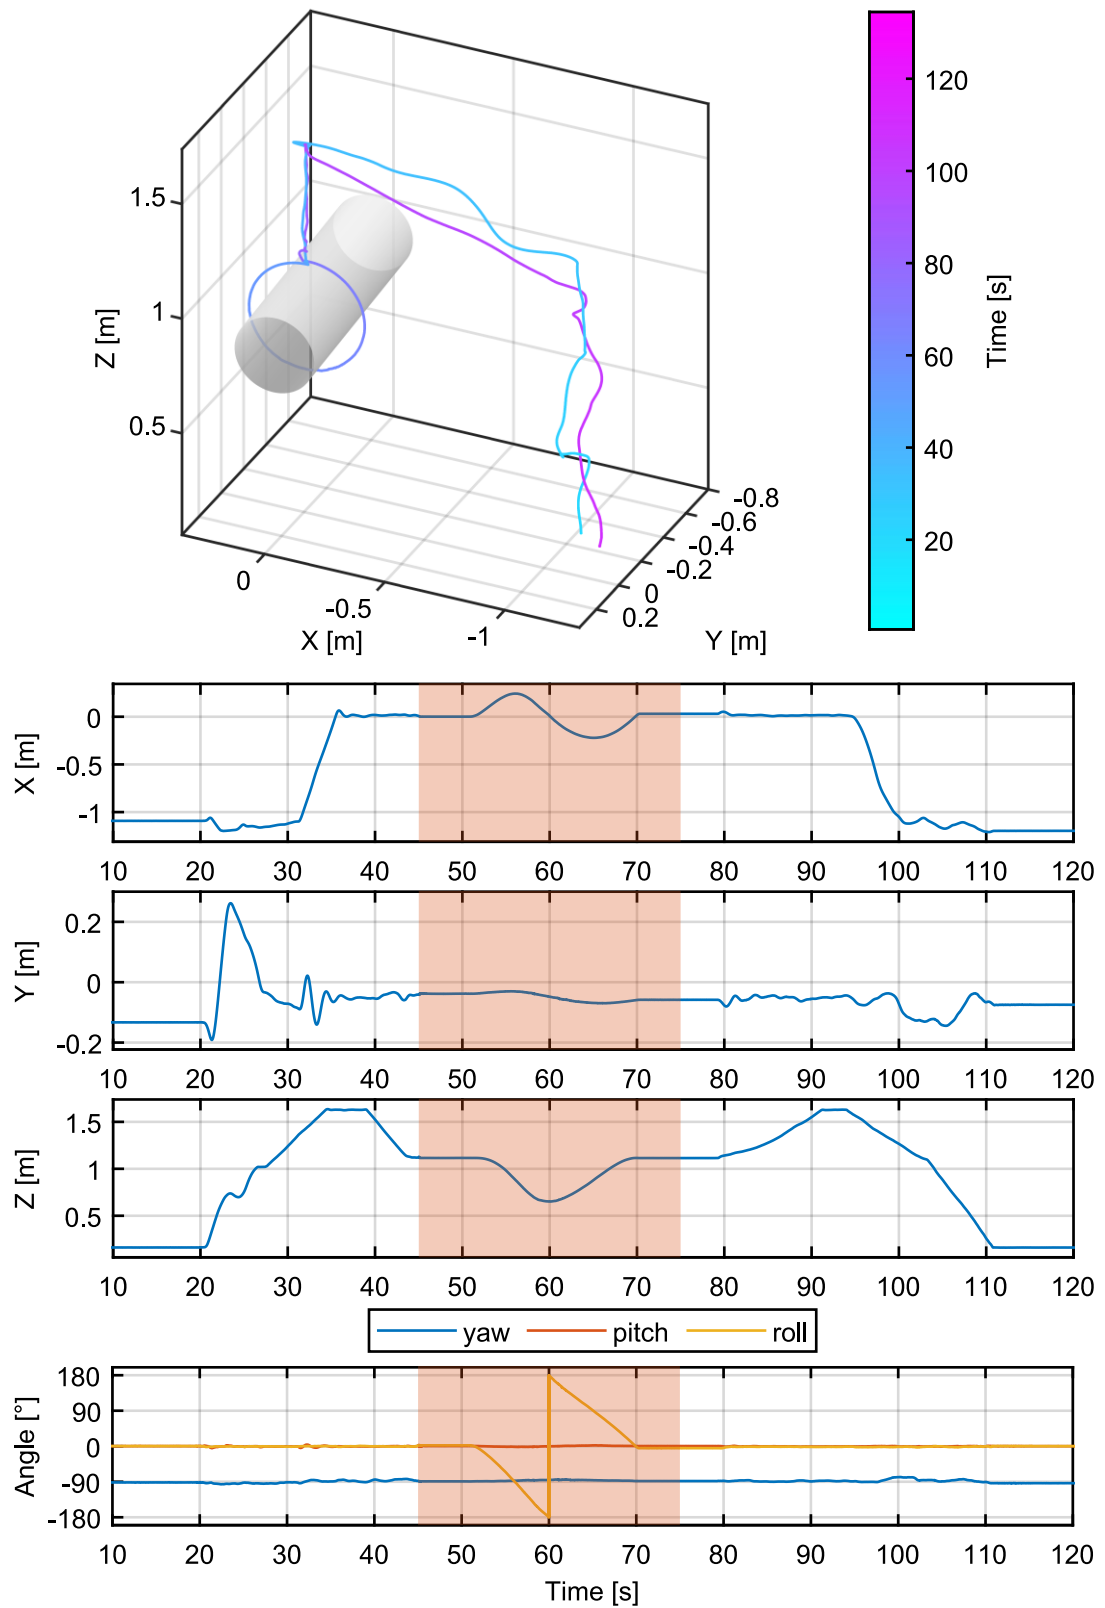


Figure 4-18: 3D position and attitude trajectory of the vehicle in $\{W\}$ as it conducts a flight-access circumferential scan. Orientation during pipe contact is given in the ZYX Euler angle convention. The highlighted region marks the times when the vehicle is in contact with the pipe.




Figure 4-19: Synchronised photographs of the fly-crawl-fly inspection. These show: (a) the vehicle and (b) an unprocessed inspection image captured while in motion passing the -90° clockface angle.




Figure 4-20: Speed of the vehicle while scanning the pipe circumference. The dashed line shows the mean speed. The dash-dotted lines show normalised RC command signals controlling the crawling motion.

4.6.5 Surface Area Inspection

This section concludes the empirical assessment with an examination of how the system may deliver visual inspection across an extended section of the pipe surface area without leaving surface contact. Such investigation comprises two aspects: a validation of the mechanism for helical translation about the pipe, outlined in Section 4.4.4, and a feasibility case-study of an inspection process using the radial orthoimage of the unrolled surface, derived as described in Section 4.5.3.

Both the spatial and imaging components of this NDE process are examined through a single experiment. The vehicle is initially positioned atop the pipe sample. With its position continuously recorded by the Vicon tracker, it is then piloted through five circumferential passes. A constant d_{sp} crawling command signal is applied to keep the vehicle at its maximum yaw angle relative to the pipe and so generate maximum translational motion. In this case, the pilot is permitted to adjust only the c_{sp} command signal to regulate the vehicle speed. Images are captured using the onboard camera at over 16 Hz throughout this process, being streamed to an operator ground station via Wi-Fi and saved for later assessment. The vehicle motion profile and imaging results are then generated through post-processing analysis of the recorded data.

4.6.5.1 Helical Translation

As such, the trajectory of the vehicle is given in Figure 4-21. This shows successful completion of the helical path, travelling a length of over 5.09 m around the pipe surface in a time of 79.3 s under adaptive thrust support. The vehicle was held in continual surface contact throughout this time, further validating the control algorithm. Consistent forward motion is visible in the $\{P\}$ x-axis, indicating that the vehicle was able to effect 2D motion relative to the pipe surface using the proposed differential steering mechanism. As such, the vehicle moved a total length of 124.6 mm along the pipe, an average of 24.9 mm per circumnavigation. Operational characteristics for surface area NDE are thus determined, identifying the accessible regions relative to the landing point and informing full coverage inspection strategies.

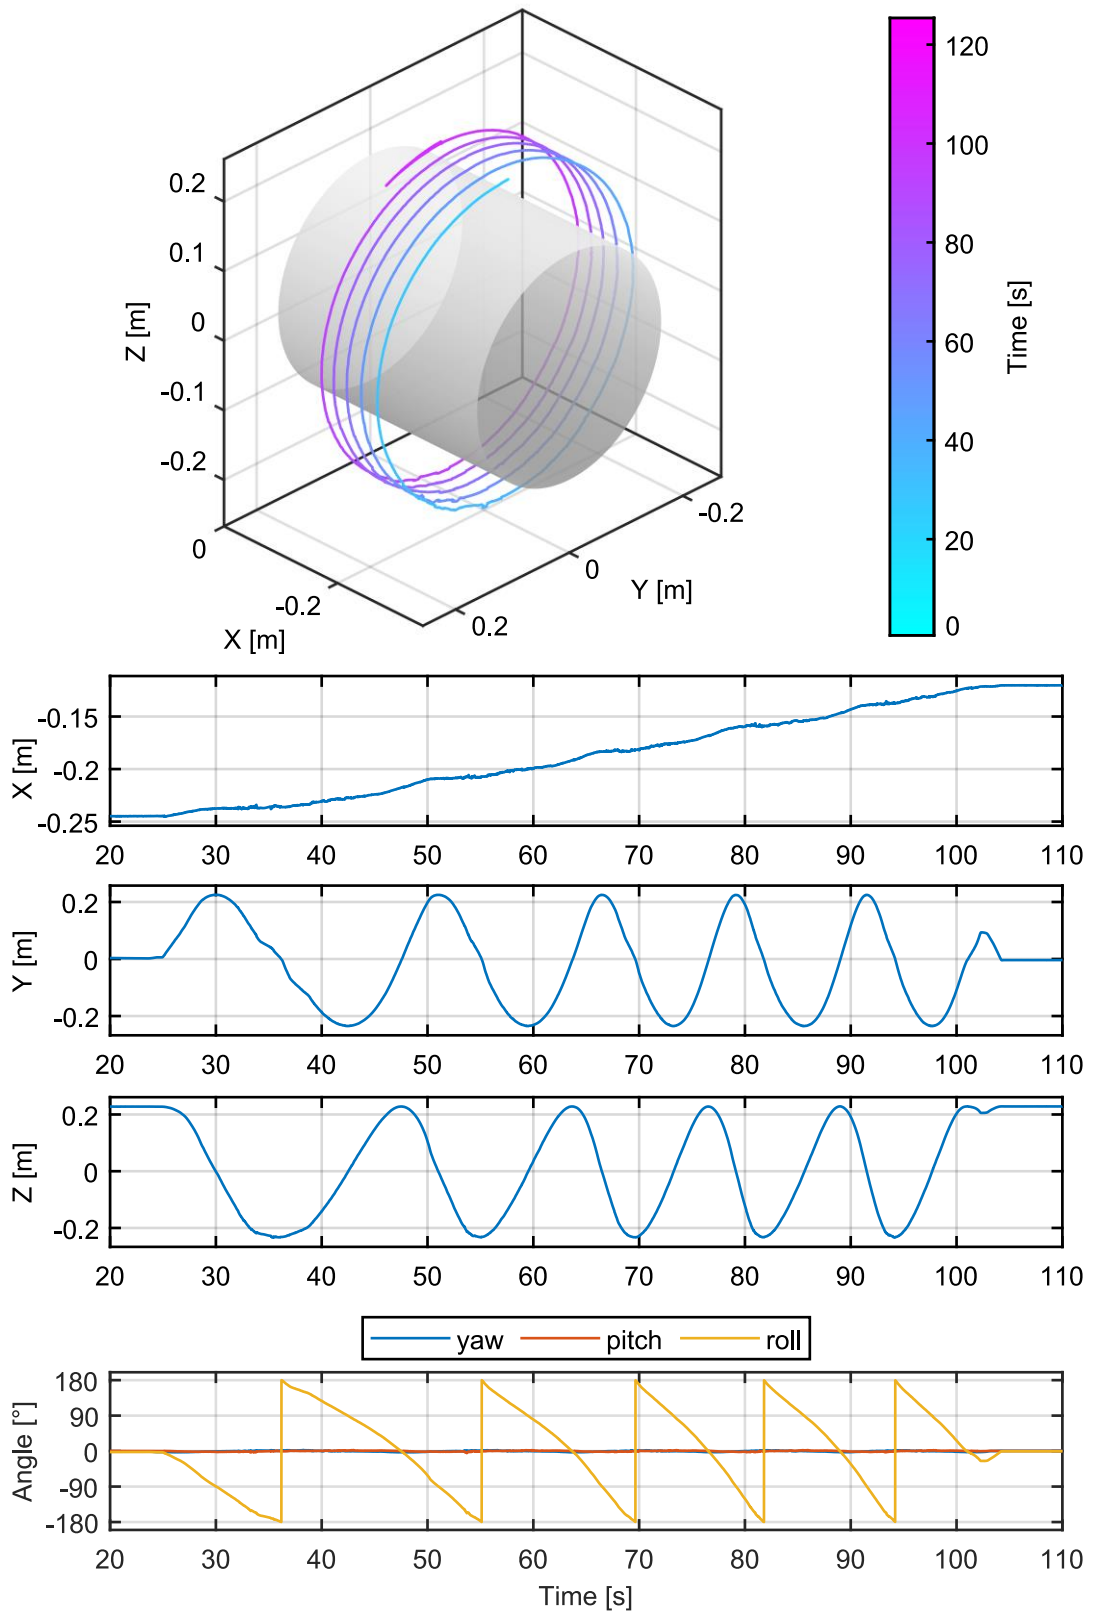


Figure 4-21: 3D position and pose trajectory of the vehicle in $\{P\}$ as it performs a helical translation along the pipe surface. Orientation during pipe contact is given in the ZYX Euler angle convention.

These results also demonstrate the ability of the vehicle to cross small surface terrain features. In moving through the path displayed in Figure 4-21, the front tyre of the vehicle has crawled across the weld cap, a stepped-ramp feature which protrudes a marginally beyond 5 mm from the pipe surface over a width of 40 mm. Here the small component of the tyre motion acting in direction perpendicular to the obstacle causes it to present a significant barrier to helical motion in the $\{P\}$ x-axis. Additional surface irregularities are also introduced in the $\{P\}$ yz-plane due to the sinusoidal weave pattern of the weld, visible in Figure 4-19(b). Successfully navigating these features therefore demonstrates robustness of the differential steering approach and indicates that it is well suited to intended applications of the vehicle where the in situ pipework exterior has been mechanically deformed, has developed protruding corrosive blister defects, or contains obstructive surface furniture such as the seams of spiral ducts and cladding.

This helical translation mechanism is further characterised by the plot of vehicle speed given in Figure 4-22, determined from vehicle pose data as previously by applying a median filter to the axial position derivatives, then computing the Euclidean length of the vector.




Figure 4-22: Vehicle speed while helically translating over the pipe surface. The dashed line shows the mean speed. The dash-dotted lines show normalised RC command signals controlling the crawling motion. Box annotations show peaks identified as cross-over effects in green and slip in orange.

This shows an average speed of 0.093 m/s, with a maximum of 0.169 m/s occurring at θ of -78° , and a standard deviation of 0.029 m/s. A gradual increase in speed is visible throughout, in agreement with the trend of the c_{sp} command signal. Gravity also imparts an approximately sinusoidal offset to the speed profile as the vehicle moves around the pipe, as previously shown in Figure 4-20. However, additional disturbance modes are observed during helical translation versus the purely circumferential path. These increase the standard deviation of the speed profile despite efforts made by the pilot to regulate the motion. Such disturbances appear as a number of peaks highlighted in the speed profile Figure 4-22.

Referencing the angular trajectory of Figure 4-21, a subset of these peaks in vehicle speed is noted to occur at the bottom of the pipe. Their causal mechanism is therefore found in the interaction controller, which induces a non-zero minimum leg supported torque to prevent roll chatter when in this location, as in discussed in Section 4.3.2. This torque must invert polarity at the cross-over point where θ passes $\pm 180^\circ$, transitioning the vehicle from resting on the legs of one side to the legs of the other. Due to the tolerance between the enclosed leg diameter and the pipe, the vehicle centre of mass therefore rolls through a short arc, rotating in the $\{B\}$ x-axis about the wheel point of contact. Here, the inclusion of the θ_D transition region in torque output prevents severe angular jerk due to instantaneous switch in propeller thrust and allows a passive shift lasting under 0.2 s. A corresponding speed component is added to the base motion around the pipe, yielding the spikes of Figure 4-22. This behaviour is confirmed by its deliberate induction when changing the vehicle direction at time of 102 s, to correct for overshoot from the desired stopping point and return to the top of the pipe. These cross-over effects are therefore well understood in the context of the controller and do not negatively impact NDE data capture.

A second subset of speed peaks shown in Figure 4-22 repeatedly occurs during decent, at θ around -67° , where the vehicle slips relative to the pipe surface. This arises from increased vulnerability to slip introduced by the kinematics of differential steering. In applying ω_z , the vehicle tyres rotate about the surface normal axis, inducing motion between the tyre and pipe surface. The static friction assumed in the tyre model is thus replaced by a significantly weaker kinetic friction mechanism, reducing the safety

margin imposed on the stabilisation thrust in the $\{B\}$ y-axis. This is exacerbated by constraint of the vehicle legs, preventing additional rotation and inducing skidding across the full tyre surface, and the vehicle yaw angle relative to the $\{P\}$ x-axis, unbalancing the symmetry between the front and rear leg reaction forces. The resulting yaw moment about the $\{B\}$ z-axis is opposed by the tyre friction forces, further decreasing their ability to support vehicle weight. The vehicle therefore slips a short distance when passing a region of the pipe with reduced grip due to localised differences in surface finish that also lie near the point of maximum applied leg forces, per Figure 4-5.

However, the slip exhibited at this point is within tolerance of the application, having minimal effect on the captured images, as examined below. The vehicle also demonstrates repeated completion of the ascent stage its circumferential pass, showing effective management of the slip conditions outside the low friction region. In practical deployment, operation without the constant ω_z included here to assess maximum $\{P\}$ x-axis translation may provide further robustness against this slipping behaviour and its detriments to scanning motion uniformity.

The slip case examined here therefore serves to aid understanding of the vehicle behaviour and the factors that must be considered ahead of any in situ application. Ultimately, the results demonstrate that the differential steering crawling control method, outlined in Section 4.4.4, can successfully effect 2D motion about the pipe surface amid the restrictions imposed by the interaction. The hybrid vehicle may thereby serve as a means for fine-scale repositioning of NDE sensors while retaining contiguous contact to directly support area coverage inspection modes.

4.6.5.2 Unrolled Surface Imaging

In completing the helical path depicted in Figure 4-21, a total of 1279 images were captured, representing a jpeg image set with file size of 2.53 GB and covering the full surface of the examined pipe section at a resolution similar to that shown in Section 4.6.2. These images are supplied to Agisoft Metashape to assess their suitability for photogrammetric reconstruction. As introduction of visual artefacts such as that of Figure 4-10 would provide distinct visual features and bias reconstruction

results, initial indicative evaluation of image quality is conducted using a proprietary algorithm within Metashape based on generic feature sharpness. Under this scoring system, a value above 0.5 suggests an image well suited to photogrammetry [363]. Here, the image set has mean image quality and standard deviation of 0.5845 ± 0.0833 , with a total of 82.96 % of images meeting or exceeding the threshold criteria. Images captured from the hybrid vehicle during helical translation are thereby validated ahead of use in 3D reconstruction.

Proceeding with the Metashape workflow, the textured 3D model was generated in 23 h 40 min 21 s on a desktop PC with 16 GB of RAM, using an Intel i7 6700k processor, and Nvidia GTX 1060 GPU. Of this processing time, image feature matching formed the largest component at 19 h 49 min, reflecting the challenges posed by self-similarity in industrial assets. This may be reduced by pre-selection of images to reduce redundant overlap but is completed here at full scale to better inform process feasibility and define upper imaging performance boundaries. The resultant high-resolution mesh has 5 408 689 vertices and 10 814 176 faces and is stored in a 551 MB ply format file. Following post-processing via the method of Section 4.5.3, the unwrapped jpeg image has a reduced file size of 24.5 MB. Both the mesh and unrolled orthoimage are presented in Figure 4-23.

In analysis of this output, successful generation of the 3D mesh without missing regions, or erroneous geometry and texture, further indicates that the image set obtained using the hybrid vehicle contains sufficient, well-resolved, unique features to permit image matching and camera registration via the SfM process. Here the benefit of immediate proximity full-circumferential visual inspection is shown. From a viewpoint at higher standoff, these features may not be resolved, leading the asset surface to appear uniform in texture and causing reconstruction to fail in localised low-feature regions. Further, the moving viewpoint has ensured complete surface coverage where manual inspection may be occluded by the asset or surrounding objects.




Figure 4-23: Images of the reconstructed exterior pipe surface area. (a) 3D model isometric view. (b) Unwrapped surface texture orthoimage. (c) Inspection features of interest in the unwrapped texture at 10x zoom. Coloured boxes show their corresponding locations in the unwrapped surface image.

As this reconstruction is obtained using monocular images in a multi-view SfM process without known physical tie point positions, dimensions of the reconstructed mesh are internally consistent but possess only relative scale. To enable sizing of defect features and overall pipe geometry, a scale factor must therefore be provided

based on known dimensions of an item within the reconstruction or physical measurement from another source. In this case, the manufacturing tolerances of the pipe specify its outer diameter be within no more than $\pm 1\%$ of nominal [364], and so allow mesh physical dimensions to be well approximated through scaling based on the best-fit radius in $\{C\}$ and nominal pipe outer diameter. As such, the unrolled mesh covers a total surface area of over 1724 cm^2 , extending beyond the length of the helical translation to cover a 193.3 mm span along the $\{P\}$ x-axis. This corresponds to an area coverage rate of over $21.7\text{ cm}^2/\text{s}$ using the helical translation process.

A further scaling factor between the unrolled surface image and mesh geometry can be computed with high accuracy based on the mean pixel distance between the unrolling seam edges and the best-fit radius. This gives an overall scaling factor of approximately 19.4 px/mm in the unrolled circumference texture of Figure 4-23(b), which may be further increased via the quality parameters of the reconstruction and unwrapping process.

Within the unrolled texture orthoimage, a number of features relevant to visual pipe inspection under ISO 4628 coating inspection standards [365] are easily detectable, with example instances given in the images of Figure 4-23(c). Scaling as above, each of these images depicts a square area with side-length of approximately 22 mm and containing an example feature of NDE interest. Instances of deterioration are clearly shown in features like rust, surface chipping, and scratching. Features of the manufacturing process are also visible including the weld cap, its start-stop point, porosity, and in some cases even variations in weld crystallisation along isotherm lines. A summary of the features shown in each image is given in Table 4-6.

The range of features visible within these images provides robust evidence in support of the practical feasibility of the UAV-crawler system for visual NDE via the unwrapping method of Section 4.5.3. Note that while annotative text is visible within the reconstructed surface, this has minor impact on the photogrammetric process as the vast majority of source images show only the unadorned pipe surface. Successful resolution of submillimetre weld spatter and crystallisation effects on the isotherm lines indicates that the vehicle scanning motion at has not significantly degraded the ability to resolve features at the scale of interest during in situ inspection versus the

static imaging case of Section 4.6.2. Further, the features of Figure 4-23(c) remain clearly visible without distortion arising from the photogrammetry and texture unrolling process.

TABLE 4-6: NDE FEATURES VISIBLE IN FIGURE 4-23(c)

Feature	Image											
	1	2	3	4	5	6	7	8	9	10	11	12
Discolouration	X	X	X	X	✓	X	X	X	✓	X	✓	X
Porosity	X	X	✓	X	✓	X	X	X	X	X	X	X
Rust	✓	X	✓	✓	✓	✓	✓	✓	✓	✓	✓	X
Surface Chipping	X	✓	X	X	X	X	X	X	✓	X	X	X
Surface Grinding	X	X	X	X	X	X	X	✓	X	X	X	X
Surface Scratch	X	✓	✓	✓	X	X	✓	X	X	X	X	X
Weld Cap	✓	X	✓	✓	✓	✓	✓	X	X	✓	✓	✓
Weld Crystallisation	X	X	✓	✓	X	✓	X	X	X	✓	X	✓
Weld Spatter	✓	X	✓	✓	✓	✓	✓	X	X	✓	✓	✓
Weld Start-Stop	X	X	X	X	X	X	X	X	X	✓	X	X

Notably, the unwrapped texture also includes small reflective regions where excessive weld spatter has been removed from the pipe surface via abrasive grinding, as in Figure 4-23(c)-8. These exhibit a complex mixture of specular and diffusive reflection properties and as such are typically very challenging for photogrammetric reconstruction. Their tendency to reflect environmental light can lead to localised overexposure and mask any unique features of their surface. Here, however, reconstruction is aided by the immediate proximity, oblique lighting, and the changing viewpoint provided by the hybrid crawler platform. Collectively, these factors allow individual score marks within the ground area to be made visible: increasing their feature scale within the image, exaggerating contrast due to shadows of fine geometry, and exploring viewpoints where specular-reflective lighting effects are minimised. As a result, the photogrammetric reconstruction is able to robustly resolve many such areas over the pipe surface, producing a texture free of holes or heavily interpolated regions and demonstrating applicability of the process in scenarios where pipe or cladding surfaces retain isolated reflective patches.

However, the unrolled texture of Figure 4-23(b) does exhibit certain minor visual artefacts as a consequence of image capture with the hybrid vehicle. Most prominently, brightness appears to vary in bands around the pipe circumference. This arises due to variation in relative surface illuminance from the onboard LED lighting array. While these LEDs are orientated such that their centrelines of maximum illuminance intersect in the centre of the imaged region, their angular illuminance profile means that light may also follow a specular-reflective path to the image sensor. As visible in Figure 4-14, light transmissibility along this path is markedly higher than the scattering path, leading to brighter regions at either side of the image despite the lower flux from the LED in that direction. Increased proximity of the metallic surface to the LED at the image sides exacerbates this effect, further increasing the illuminance in these regions. A mask is applied to exclude the most severe regions when merging the images to form the final mesh texture, but this must be balanced with retaining image area and so the effects cannot be fully eliminated. In turn, the bright strips correspond to the path of the camera around the pipe circumference. Similar discontinuity in brightness is also visible in the 0° to -45° region in the left of Figure 4-23(b), stemming from repeat coverage of this area due to overshoot past the 0° stopping point by the pilot which increases the relative weighting of the specular-reflective bright regions in the final texture. Nonetheless, the overall exposure in the unrolled texture remains within the usable range for clear representation of the features of NDE interest.

4.7 Discussion

Fitness for purpose of the hybrid UAV-crawler approach is thus proven above, highlighting stability of the feed-forward control algorithm in provision of adaptive thrust support, and permitting flexible access to in situ structures.

In the context of other state-of-the-art works, inclusion of radial adhesion thrust in a novel, flying, vehicle for full-circumference crawling significantly advances previous stabilization of a non-flight-capable cart within $\pm 20^\circ$ of the pipe top [279]. Further, versus the earlier NMPC approach [279], Algorithm 1 is highly portable with low embedded compute and sensor requirements, also utilizing propeller torque effects for energy-usage minimization.

By explicitly modelling contact forces, Algorithm 1 recognises when full support is feasible without propeller action, allowing them to be selectively and automatically disabled. Energy usage atop the pipe is lowered versus other strategies using active stabilization at 20 % propeller thrust [277], or omitting contact effects entirely [280]. Moreover, where [280] governs angular position of a benchtop test rig purely by torque, the lateral thrust support used here decreases actuator utilization on pipe sides where the vehicle weight moment is largest.

Here, sensor deployment from a surface crawling static airframe significantly reduces vehicle mass and mechanical complexity versus NDE performed from atop the pipe via manipulator arm [277], [278]. Compensation for moving arm CoM by battery repositioning [278] is precluded and there is no risk of the effector accidentally pushing the vehicle off the surface. Further, surface crawling removes scalability limits due to manipulator arm length [277], specialised effector tool size [278], or circumferential grasping [274], [280], expanding airborne inspection to large diameter pipes, tanks, or vessels.

Within the application context, the consistent CoM simplifies interaction behaviour and NDE sensor placement. Parametric expression of the desired radial contact force presents strong utility for compressive coupling, as with dry-coupled ultrasonic transducers [263]. Similarly, 95th percentile absolute position deviation under 2.4 mm, per Table 4-3, is far lower than that of free-flying visual inspection UAVs [234] and of small scale relative to the NDE image area. Higher sensor stability is shown in discrete magnetic sensor packages placed on a 6-inch pipe by a UAV [273], but therein manually piloted mean landing accuracy of 130 ± 40 mm at a 60 % success rate restricts inspection targeting. Here, under automated landing control, a 100 % success rate at mean position accuracy of 12.68 ± 12.20 mm is shown using direct thrust in the $\{B\}$ y-axis, overcoming issues of inconsistent ground-effect lift and unmodelled collision with the pipe posed to conventional UAV systems with attitude-based translation [348]. Moreover, the mean landing position accuracy is smaller in scale than the 24.9 mm distance along the pipe that may be translated in a single circumnavigation, showing how surface crawling may augment flight positioning performance for fine-scale NDE applications.

When in contact with the target, this surface crawling mechanism presents further distinct advantages versus free-flying UAV systems. Conventional multirotor craft can struggle to enter the sustained repeatable surface contact necessary for inspection, showing deviation of over 150 mm from a target location in instantaneous contact [251]. As in Chapter 3, aerial manipulator UAVs have been shown to significantly improve this aspect, with mean distance to the target contact point under 73 mm across multiple interactions, each sustained for duration over 2 s [263]. However, such vehicles can encounter fine-scale sensor repositioning challenges due to long-reach arms or complex interaction dynamics and so may commonly rely on a strategy exiting and re-entering contact, incurring high sensitivity to wind disturbances. By contrast, the 2D surface crawling shown by the hybrid vehicle permits highly accurate sensor deployment through active pose refinement during continuous contact and so effectively eliminates positioning error as an issue within the NDE process.

Such position control is typical of surface crawler robots. By fully enveloping the pipe, previous examples have shown the benefits of mechanical clamping, demonstrating excellent stability and indefinite deployment duration [267], [268]. Leveraging this strong interaction, they may also rapidly translate both along and around the pipe axis at maximum speeds of up to 150 mm/s and 280 °/s using an omnidirectional mecanum wheel drive [267], or 14.2 mm/s and 13.8 °/s using an inchworm mechanism [268]. These drive mechanisms are significantly faster than the average speed along the pipe axis of 1.57 mm/s shown by the hybrid vehicle during helical translation owing to their more direct kinematics but are comparable to its maximum circumferential speed of 59.8 °/s where the differential drive better aligns to the direction of travel. In this case, the lower axial translation speed is compensated for by the ability to fly between targeted inspection points, bypassing radial obstacles that would obstruct all access by the crawler systems. Further, the thrust-based adhesion strategy does not preclude alterations to the servo drive train and so may be upgraded to adopt an omnidirectional drive mechanism if warranted by the particular NDE application. The thrust support developed here may then meaningfully extend access capability of wheeled UAVs currently restricted to resting atop the pipe [277], [278] to include full circumferential deployment.

Regarding visual NDE performance, empirical assessment has been conducted to quantify performance relative to existing practices and identify any negative influence of the novel hybrid UAV-crawler platform. Results from the USAF MIL-STD-150A assessment presented in Table 4-4 show an average 95th percentile spatial resolution of 4.915 lp/mm in images taken while the vehicle holds static pose around the pipe surface. Line features of thickness above 102 μm are thereby accurately resolved. As such, challenges to visual NDE of fine-scale surface defects, like porosity and cracking, posed by the degrading effects of surface standoff and motion blurring intrinsic to free-flying UAV systems [234] are directly and successfully addressed. Further, the system spatial resolution is significantly above the minimum acceptable value of approximately 3 lp/mm for direct, in person, visual assessment under ASME BPVC Section V Article 9 standards [56], corresponding to that of a human eye at 600 mm standoff [354]. The minimum resolution of 1 lp/mm for evaluation of blistering, rusting, cracking, and flaking under ISO 4628 standard coating assessment [354], [365] is also well surpassed, showing performance levels with strong applicability to existing NDE practices.

These results are further supported by feasibility study of a visual NDE approach combining images captured by the hybrid vehicle to form a single unwrapped orthoimage of the entire target surface region under examination. While a 3D mesh provides intrinsic location context to defect features relative to the larger asset and is highly informative within digital-twin processes, the unrolled format also provides distinct practical benefits. In addition to intrinsic localisation context, the image is better suited to analysis within existing image inspection workflows and easily integrated with existing conventional and machine learning based defect feature identification algorithms [366]–[368] to speed review under human expert supervision and provide cost-effective tracking of defect growth over time. File sizes are also significantly reduced. For example, the unrolled image of Figure 4-23(b) requires 4.45 % of the storage space used by the 3D mesh and 0.95 % of that used by the base image set, eliminating significant redundant overlap. Further, within the reconstruction process, mesh fidelity and unwrapped texture quality settings can be altered to balance file storage requirements with NDE resolution targets and optimise results for the application.

The utility of similar approaches has been shown in the context of pipe interior inspection using borescopes [114], custom designed probes [113], [136], as well as Pipeline Inspection Gauges (PIGs) and robotic rovers [369], but is not widely considered for pipe exteriors. Whereas these methods rely on the omnidirectional viewpoint at the pipe centre to synchronously image the full interior circumference, results here show that similar outcomes may be derived from multiple sequential images captured in mechanical scanning of the exterior. As such, the camera pose is well controlled at constant standoff from the surface, removing the need for the centralisation or compensation strategies used inside the pipe [113], [136]. Moreover, this provides robustness against issues of inconsistent surface relative pose that can blur and distort images without an active sensor-based control loop [114]. Reconstruction then may be completed without missing texture patches, fully covering the area under test at consistent quality and resolution. This serves as strong supporting evidence for the utility of the developed hybrid vehicle.

4.8 Summary

In summary, this chapter has examined modelling, control principles, and quantitative performance assessment of a novel hybrid UAV-crawler vehicle. In a multi-point interaction approach, an MDT array of bi-directional fixed-pitch propellers is used to effect thrust-based adhesion to horizontal-axis cylindrical targets. The proposed offline optimization, feed-forward, control algorithm permits stable pose holding and translation about the asset outer circumference, with performance bounds established by laboratory testing. A fly-crawl-fly inspection mode is examined in detail, allowing rapid deployment and navigation of obstacles posed to conventional crawler access by surface furniture or other geometry. The proposed vehicle is thereby fully validated as a unique means for flight access to the outer surface of assets such as pipes, storage tanks, or pressure vessels common throughout the energy sector.

This is supported by detailed assessment of visual NDE performance in multiple static point inspections, showing resolution comparable to in-person processes under relevant industrial standards while proving full circumference visibility. A feasibility case study extends this to a surface scanning workflow, demonstrating 2D translation

while in continuous contact with a pipe surface. Merging the images captured during this process, a strategy to create a single unrolled orthoimage via photogrammetric reconstruction then highlights the benefits to inspection data localisation and contextualisation offered by this highly novel NDE platform with strong adaptability to existing analysis procedures. Ultimately, this work is expected to support increasing industrial adoption of aerial NDE systems, reducing human risk factors while expanding asset management insight and cost-efficiency.

CHAPTER 5

Conclusion

5.1 Summary

This thesis has presented a detailed examination of the current state of remote non-destructive evaluation for in situ industrial structures using unmanned aerial vehicles. Focus is now placed on how these airborne platforms can interact with the target asset to obtain meaningful quantitative measurements of its structural health, expanding insight currently provided through non-contact visual inspection by commercial systems.

Chapter 1, began with an establishment of motivations for these objectives. This considered the needs of the UK energy sector for inspection of critical infrastructure as part of its standard operating life and following plant decommissioning. An overview of the future electricity generation strategy identified the wind, nuclear, and oil and gas sectors to be of particular importance, constituting the majority of planned capacity until 2050 under low carbon objectives [9]. The inspection challenges facing each of these sectors were reviewed, identifying the most destructive degradation mechanisms and problems encountered in attempts to monitor their progression. Recurrent problems due to corrosion, erosion, and fatigue affecting all sectors were highlighted in large structures exposed to harsh environments such as storage tanks, pressure vessels, pipework, and offshore platforms, alongside other more field specific hardware. These cases each posed issues for manual inspection in their requirements for skilled work at height, often above the open sea; work amid confined spaces, encountering hazardous atmospheric conditions; operation within safe radiation exposure limits; and the cost of lost production during assessment of active facilities. In line with strong industrial demand to alleviate these issues, the benefits of current

remote inspection practices using multirotor unmanned aerial vehicles were examined. Supported by a number of case studies, this illustrated how UAV inspection can provide a highly mobile, elevated viewpoint, capable of flexible deployment with minimal supporting infrastructure. As such, it has resolved many access challenges, delivering informative and cost-effective visual survey free from the hazards of direct human presence.

Fundamental research motivation was then established in the desire to expand these benefits to include quantitative, contact-based sensory mechanisms for the above applications. This would augment airborne inspection capability to incorporate reporting of structural health status based on more discerning indicators located beneath the outer surface, further removing needs for resource intensive manual follow-up. To meet this goal, specific research aims for the thesis were defined. These were accompanied by a list of contributions to knowledge made in its completion, an outline of thesis structure, and a record of academic publications arising from this work.

In Chapter 2, a detailed review of academic literature and industrial activity was conducted to support the technical background necessary to further develop systems for airborne non-destructive evaluation. From among the range of existing NDE methods commonly applied to in-service assessment [95], [99], visual and ultrasonic testing were identified as candidates best suited to potential airborne implementation based on a first-principle assessment. They presented broad and complimentary applicability to the target applications, with minimal obstacles to UAV deployment posed by their functional requirements or their small, lightweight, instrumentation.

Accordingly, key operating principles for visual NDE and its extension to 3D reconstruction via photogrammetry were examined, as supported by underlying theory reviewed in Appendix A. This highlighted the ability of photogrammetric mesh reconstruction to retain the situational context of remote visual inspection images, despite their marked self-similarity, while also noting that orthophotography provides similar advantage within an existing 2D workflow. Comparable introduction to the operating principles of ultrasonic inspection followed, establishing how internal structure is inferred during thickness measurement and other extended NDE, as

described in Appendix B. Understanding sufficient to replicate these procedures from an airborne platform, and thereby assess the targeted energy sector defect modes, was thus developed. A discussion of the operating principles surrounding multirotor UAVs then completed the background material. As such, the aerodynamics underpinning the collective force-torque actuation principles of a propeller array were established. An overview of the control algorithms developed to exploit this behaviour and permit stable flight was then given, supported by a comparative analysis of salient aspects in Appendix C.

Trends for how these technologies were collectively employed in practice were then identified, noting that photogrammetry is given little consideration for close-proximity inspection of small-scale assets with current works favouring its use at infrastructure and landscape scales or adopting it for spatial registration of other data, such as in radiological survey applications. Factors limiting commercial adoption of ultrasonic and contact-based inspection were also apparent in the dynamics of conventional underactuated multirotor craft, requiring motion of the full platform to counteract disturbances. Solutions to address this were presented from within the robotics community, developing aerial manipulator UAV platforms incorporating multi-direction or actively vectored thrust to directly generate interaction forces and stabilise interaction. These works, however, are noted to often overlook the influence of the floating platform on the NDE task it performs. Opportunities to improve sensor stability were then identified for further exploration in systems contacting the target through adhesion, grasping, perching atop key features, or otherwise utilising the structure to support a stronger interaction. Themes to be explored through the remainder of the thesis were thus identified based on perceived gaps in the current state of the art.

In Chapter 3, a novel airborne inspection system was developed to perform ultrasonic thickness measurement of industrial structures using a dry-coupling wheel probe. The conceptualisation for this system was first explored, with a summary versus established works detailed in Table 3-1. In recognition of the recent innovations in aerial robotics, the system thus was developed to leverage the advantages offered by a thrust-vectoring over-actuated UAV. This provided omnidirectional thrust generation,

allowing for hovering at non-horizontal attitudes and stable interaction with arbitrarily oriented targets. System applicability was then greatly advanced versus previous approaches employing unidirectional thrust and mechanical stabilising structures, capable of interaction in only a single orientation. A more efficient vectored thrust paradigm, combined with the removal of need for an embedded manipulator arm, also reduced overall form factor and so aided operability within confined spaces versus previous multidirectional thrust platforms. Including the dry-coupling wheel probe then simplified the required ultrasonic hardware, removing the need for an onboard couplant gel reservoir and dispenser. This would also preclude issues with changing mass distribution or post-measurement clean-up in scenarios where couplant gel may represent a degressive contaminant. An affinity for linear scanning ultimately made this wheel probe a strong choice for area coverage applications, particularly the defined objectives of corrosion monitoring in large energy-sector structures.

Empirical assessment followed, using an analysis framework developed to validate the integrated system and quantitatively characterise any influence of the floating platform on measurement performance. As such, testing was conducted in reference to samples of known geometry, containing machined regions representative of corrosive wall loss. Measured thickness was determined by an autocorrelative ultrasonic signal processing algorithm, extracting inter-echo time of flight in the presence of variable first-echo delay due to tyre compression. Projection of the measurement point onto the sample surface then allowed direct numerical comparison of thickness with reference geometry, repeated under differing conditions to isolate influential factors.

A full application focussed characterisation of the system was thereby conducted, demonstrating feasible static point measurement of a vertically planar wall and the underside of an overhanging surface, as shown in Figure 3-17 and Figure 3-21. Ultrasonic thickness measurement accuracy in both cases was comparable to the manual process, showing mean absolute error below 0.095 mm and 0.093 mm, respectively, and so indicating minimal influence of UAV orientation. Through repeat interaction, measurement location uncertainty was also characterised at a worst-case radius of 16.55 mm, as in Table 3-7. Testing likewise determined that the probe could be positioned within a mean distance to a target point of 37 mm and 12 mm in the

horizontal and vertical axes, respectively, with minimal pilot effort, qualifying the ability of the system to repeatably access known defect locations. A final experiment quantified the system ability to perform rolling scans and gather continuous measurements along a linear path. All features of interest were successfully recorded in two passes across the sample surface, but the scanning motion was irregular and ultrasonic measurement accuracy was reduced to mean absolute error below 0.279 mm, as in Table 3-9. Investigation attributed these effects to error introduced by the uncertainty in the projected measurement location, causing measurements to be incorrectly registered to adjacent reference features, while sudden variation in wheel probe rolling resistance posed challenges to retaining a smooth continuous motion. Informed by this empirical assessment, expectations for the airborne ultrasonic thickness measurement system may be set and used to justify its confident adoption within standard in situ inspection practice.

In Chapter 4, a second novel vehicle was designed to explore the benefits of an alternative means for contact-based airborne inspection, hybridising aspects of surface crawler robots with multirotor aircraft. As with the previous UAV, system conceptualisation first established the unique aspects of this approach and its anticipated benefits, as summarised in Table 4-1. Accordingly, the vehicle was designed to deliver NDE while in strong, short-reach, contact with the target. Retaining flight capability so as to navigate around obstacles, it would then expand applicability versus crawler robots while mirroring their ability to perform fine-scale position refinement during continuous contact. Supported by its own thrust, the vehicle would be free from the target material constraints of magnetic or vacuum adhesion, while also removing restrictions to scalability incurred by mechanical grasping strategies. The hybrid vehicle would thus serve as an effective platform for inspection of the most common large-scale energy-sector infrastructure, here deploying immediate proximity visual inspection to target pipes, storage tanks and pressure vessels.

Development of a vehicle incorporating these features then entailed construction of a rigid-body Newton-Euler model, mathematically describing its interaction with a target asset. From this, the propeller actuation principles of Section 2.4.1 were expanded to define generic feasibility criteria for the minimum necessary force-torque

output to stabilise the vehicle in any orientation about the target surface. Employed to validate prospective designs, these metrics enabled development of a prototype utilising six rotors in a multidirectional thrust array, as depicted in Figure 4-1. A control scheme minimising energy usage in each pose via the least-squares sum of individual motor thrusts was also formed via the interaction model and implemented in real-time using a heuristic approximation function. Accompanied by free-flight control and differential steering of two motors at the vehicle front and rear, as depicted in Figure 4-7, the interaction control system was then capable of supporting 2D crawling around a target surface. A novel vehicle was thus designed from first principles to perform thrust-supported, surface-crawling, inspection.

Again, a series of laboratory experiments validated key aspects of functionality in the constructed prototype and quantified its performance in an application context. Successful circumnavigation of cylindrical targets was demonstrated by inspection of a 12.75 inch outer-diameter pipe section via a fly-crawl-fly inspection strategy. This saw the vehicle transition from flight to contact-based inspection, complete visual assessment in continuous motion around the pipe circumference, and exit contact to fly back to its starting position on the ground. When stopping at inspection points distributed in 45° increments around the circumference, the vehicle was stable to within a mean absolute deviation of under 1.03 mm, as shown in Table 4-3, a marked enhancement versus free-flying systems. Assessed via USAF MIL-STD-150A procedure [353], the vehicle was then able to image the target surface with resolution up to an average 95th percentile value of 4.915 lp/mm, resolving lines of 102 μm thickness or above and exceeding ASME and ISO standard requirements for direct visual inspection by a human [56], [354], [365]. A final experiment then assessed area coverage inspection, with the vehicle moving in 2D about the pipe surface via a helical path. While the differential steering increased vulnerability to slippage, the overall motion profile was within application tolerance and imaging results were largely unaffected. Combined to a single unrolled orthoimage via a photogrammetric algorithm, as shown in Figure 4-23, visible defect features indicated fine-scale rust, scratches, chips, weld spatter and even crystallisation structures with full location context, demonstrating the value of the immediate proximity inspection process. Strong potential utility is thereby shown for employment of these technologies within

the energy sector where frequent inspection of elevated or otherwise inaccessible pipework entails significant resource commitments.

In short, it has been demonstrated that multiple systems for airborne, contact-based, NDE can capture inspection data of sufficient quality to provide actionable insight for structural health management and repair. As characterised, these novel systems match or exceed the quantitative performance levels of manual inspection, also enabling flexible access with a reduced hazard profile. Without the extensive resource investment needed to support human entry entailed by scaffolding installation or dissipation and clean-out of harmful substances, remote airborne inspection may then be performed more frequently and at lower cost. Such attributes may then ultimately enable more cost-effective monitoring and NDE across a range of assets critical to continued operation of the energy sector.

These outcomes have been disseminated through several journal publications, conference proceedings, and presentations, as recorded in Section 1.6.

5.2 Key Findings & Implications

Through the examination and empirical characterisations of both systems for contact-based airborne inspection presented herein, the following key learnings are obtained. These provide succinct, actionable, insight for operators of industrial facilities, inspection service providers, and the wider research community wishing to implement contact-based airborne NDE systems.

- Using an over-actuated aerial manipulator UAV, static contact with target surfaces could be repeatedly entered and maintained with a 100 % success rate across multiple environmental interactions during laboratory testing. This system was demonstrably capable of exerting compressive force over the minimum 10 N necessary to enable dry-coupled ultrasonic thickness readings with mean SNR of 36 dB on uncoated aluminium samples.
- Versus ground truth sample dimensions, airborne ultrasonic thickness measurements consistently showed mean absolute error of 0.095 mm and

below against a vertical wall and 0.092 mm on an overhanging surface. Measurement performance for this system is then comparable to performance of handheld instrumentation [319], [334], [335] and largely independent of target orientation, where unobstructed access is possible.

- Projection of the UAVs onboard pose estimate to the point of environmental contact demonstrably encoded the 3D position of ultrasonic measurements to within a worst-case uncertainty radius of 16.55 mm in laboratory conditions, owing to the projection distance and the fusion of position and attitude noise. Location uncertainty spanning sharp feature boundaries can also incorrectly assign data to adjacent thickness regions, degrading mean absolute measurement error to 0.279 mm in the stepped bar sample. This uncertainty will impose a feature size limit on effective defect localisation and tracking across repeat inspections beyond what can be observed in individual ultrasound signals.
- Nonlinear friction effects and a long-reach interaction make probe coupling and alignment challenging to maintain during rolling scan thickness measurements, causing an intermittent motion profile with a low average speed of 8.5 mm/s, standard deviation of 12.7 mm/s, and frequent stops in the system under test. Though signal acquisition success rate dropped to 86.44 %, sufficient data were captured to resolve all features of the test-piece cross-section. Future work should then target these aspects to improve inspection coverage rate.
- Formulating a detailed rigid-body model of the full vehicle-target interaction has enabled design of a novel UAV-crawler hybrid platform for fly-crawl-fly inspection. Behavioural insights from the model have also yielded a control strategy minimising pose stabilisation power requirements and requiring only standard flight controller IMU sensors. Similar insights form the basis for an approximation function permitting real-time implementation of the optimised control function at 200 Hz within a low-cost single-core microcontroller. This strategy may be beneficially extended to other airborne contact applications.

- Entering into strong, multi-point, surface contact provides significant platform stability for airborne sensor deployment. Utilising partial support by the target pipe at a short distance from its centre of mass, the hybrid vehicle has held arbitrary circumferential positions with a mean absolute deviation under 1.03 mm. This represents an order of magnitude improvement over previous free-flying visual inspection UAVs [234].
- Ability for precise repositioning during surface contact can markedly increase the airborne deployment efficacy of static NDE sensors. Entering contact with a mean absolute deviation of 23.60 mm from the grouping centre, the over-actuated UAV encounters limits to probability of detection and follow-up assessment for smaller scale defects. Though the hybrid vehicle has entered contact with a mean absolute deviation of 12.68 mm from its target, its well-controlled 2D surface-relative locomotion means this may be corrected after the fact and so will not impact sensor positioning.
- Similarly, surface relative actuation using contact forces has been observed to offer improved sensor motion control, with the hybrid vehicle demonstrating consistent average circumferential scan speed of 75 ± 12 mm/s and mean helical path speed of 93 ± 29 mm/s, the later being subject to changing control input and minor induced slippage. This is faster and more consistent than the intermittent thrust-driven motion of the rolling ultrasonic scan performed by the over-actuated UAV, with an average probe speed of 8.5 ± 12.7 mm/s.
- Under thrust-stabilised contact, immediate-proximity imaging by the hybrid vehicle has shown a 95th percentile resolution of 4.915 lp/mm when evaluated by USAF MIL-STD-150A procedures. This sufficient to resolve cracks or linear defects with thickness over 102 μ m and exceeds requirements of manual visual inspection set by industrial standards [56], [354], [365], indicating significant potential to reduce human presence requirements and associated costs.
- Oblique onboard lighting and a short camera standoff aiding visibility of fine-scale features have permitted full photogrammetric reconstruction of a self-similar pipe sample with localised specular reflections via commercial software

using 1279 images captured in 79.3 s. The resultant textured mesh provides targeted 3D geometry measurement and deformation tracking over a 193.3 mm long pipe section. Unwrapped into a single orthoimage, the unified view retains full situational context of small defect features at a resolution of approximately 19.4 px/mm and file size over 105 times smaller than the original image set. Lack of perspective distortion then enables detection and accurate sizing of discontinuities within existing visual inspection workflows.

Within the context of the current state of airborne NDE technology and industrial practice, the quantitative outcomes summarised above present broader implications for adoption and ongoing development. These are outlined as follows.

- Fundamentally, both the over-actuated UAV and the UAV-crawler hybrid vehicle have demonstrated successful airborne deployment of NDE. Remote dry-coupled ultrasound and immediate-proximity visual inspection are now proven feasible within quantified performance boundaries. As such, several applications for contact-based airborne inspection within the energy sector have become more tractable means to reduce access costs and human hazard exposure. These include the regular assessment of environmental degradation in wind turbines, structural health in active or enclosed nuclear infrastructure, and corrosion in storage tanks, pipework, or other large steel structures.
- Through this thesis and work in the wider field of aerial robotics, the technology readiness level of remote airborne contact-based inspection has risen significantly in recent years. Approaches using aerial manipulator platforms with multidirectional [258], [260] or vectored thrust [259] are fully capable of rejecting aerodynamic disturbances to position NDE sensors with markedly enhanced accuracy and stability versus conventional UAVs [249], [251], [370]. Hesitancy in industrial adoption is thus expected to reduce as such UAV systems pass formal NDE qualification and enter commercial service, displacing perceptions of limited stability, precision, and repeatability of sensor placement formed around earlier under-actuated platforms.

- Adaptive counteraction of dynamic and unpredictable disturbance forces is then highlighted as an essential component of stable airborne interaction for NDE. Where contact occurs weakly, with low forces exerted at a single point or at long reach from a free-flying UAV, this necessitates highly reactive thrust control not possible in underactuated craft and challenging in vector thrust designs with slower dynamic responsiveness. Strong environmental interaction, utilising multiple contact forces at short reach with magnitude far above the disturbance, offers more intrinsic mechanical stabilisation and simplifies the control problem. This presents a more attractive solution for NDE sensor positioning but can become reliant on specific target attributes to enable implementation, such as magnetic susceptibility or graspable geometry, and so has limited generality. Accordingly, both strategies should continue to be actively pursued to maximise the applicability and efficacy of contact-based airborne NDE.
- Given the adaptability of the omnidirectional system to various plant geometry and the streamlined nature of the process when utilising dry-coupled ultrasound, a compelling case can be made for its industrial adoption in the immediate future. With improvements to the rolling motion control and position encoding accuracy, the system may enable at-scale deployment of rapid thickness profiling across the large surface area of tanks, monopile platforms, pipework, and other energy sector assets. Going further, the system architecture tightly integrating NDE within the robotics platform to leverage onboard data transmission, pose estimation, and control feedback may serve as a strong archetype for generic airborne inspection.
- Likewise, the UAV-crawler hybrid vehicle demonstrates significant utility within existing industrial processes where scaffolding, rope access or other hazardous and resource intensive methods currently support work at height. With the visual payload considered here, the resultant textured meshes and high-resolution orthoimages provide full geometry capture and easily reviewed imaging of short pipe segments, suited to digital-twin production and visual screening. While the thrust-supported crawling approach is less mature than

free-flying aerial manipulator strategies, it has demonstrated substantial advantages for NDE sensor placement stability and motion control. With expansion to permit operation outside of the specific pipe inspection use-case shown, such hybrid systems may then provide marked benefits in deployment across the energy sector.

- Future commercialisation of these technologies may beneficially adopt trends currently exhibited by UAV visual inspection service providers, following enhancements to digitised reporting and autonomous operations. In particular, operation of airborne contact-based NDE systems beyond visual line of sight and as site resident infrastructure would provide significant value for routine inspection of remote facilities such as offshore wind turbines or oil and gas platforms.

5.3 Limitations & Future Work

Considering the findings presented throughout this thesis and the potential disruptive influence they offer within industrial applications, means to address limitations and further advance airborne inspection technologies are proposed. These are elaborated upon as follows:

5.3.1 Sensor-Based Interaction Control

Information regarding in situ industrial assets can often be imprecise, with significant changes between their specified design and present condition accrued throughout their lifespan. Accordingly, variation in their position, dimensions, surface geometry and texture, or even installation or removal of auxiliary sub-structures may be expected. Similarly, mathematical models of the UAV agent used to derive their interaction control laws often differ from the physical system and include uncertainty. Combined, these aspects can pose challenges to robust NDE deployment where the interaction system differs significantly from expectations.

In the case of the Voliro UAV, the most influential factors stem from discontinuous friction effects when dynamically scanning the wheel probe across the target surface.

Sensor-based feedback of interaction force and surface relative probe odometry into an application specific control algorithm would allow a smoother rolling motion while guaranteeing satisfaction of the probe coupling criteria. This would increase scanning speed and coverage reliability, also improving measurement position registration and minimizing redundant passes for area thickness mapping.

Regarding the UAV-crawler hybrid, incorporation of feedback from tactile sensors may relax system model requirements and increase its ability to adaptively correct for unmodelled or changing conditions. Mechanisms to adapt to factors such as variation from the expected asset diameter, changes in surface friction outside of tolerated values, or differences versus desired thrust output would enhance robustness in practical deployments with minimal prior knowledge of the structure or its environment. This feedback may also permit further optimisation of energy usage, replacing fixed safety factors, imposed to assure interaction stability, with dynamic disturbance compensation.

5.3.2 Environmental Influence & Field Trials

Experiments used to characterise performance of the airborne NDE systems detailed herein were conducted in a laboratory setting, providing strong control over environmental influences and other disturbances. This allowed factors of interest to be isolated and examined independently, as necessary to support initial development and characterisation. However, practical field deployment will typically expose a UAV to a range of effects with the potential to alter system performance. Further study is suggested to repeat the experimental assessments of Section 3.4 and Section 4.6 while subjecting the UAV to controlled disturbances of increasing severity; replicating wind gusts, near surface aerodynamic turbulence, precipitation, changes in asset surface characteristics, sudden degradation in pose estimation accuracy, and other such aspects.

Combined with field trials in representative industrial facilities, results would then qualify expectations of NDE performance in sub-optimal conditions, building confidence in the system ability to gather meaningful data across a variety of real-world scenarios. Practical operability limits may also be precisely established,

informing planning strategies for UAV contact-inspection works amid current site weather conditions, forecasts, and anticipated workability windows. Detailed examination of the causal mechanisms behind any performance degradation can also inform further system development to mitigate or remove such factors, ultimately improving practical utility.

5.3.3 Applicability to Generic Asset Geometries

To some extent both systems for airborne contact-based NDE described herein rely on the presence of certain features in the target asset to perform the interaction necessary for data capture. Existence of these features cannot be guaranteed in all applications, imposing feasibility limits on particular use-cases.

Concerning the over-actuated UAV, its capability for omnidirectional interaction incurs only loose requirements. Provided there is sufficient clearance around the asset to permit safe unobstructed access, single-point thickness measurement is feasible in most scenarios where it may be otherwise conducted. However, the interaction controller as described in Section 3.2.1.2 is unaware of the target surface orientation relative to the UAV and so its attitude must be manually controlled, introducing dependencies on pilot skill level. Accordingly, inclusion of environmental perception to inform appropriate setpoint generation within the attitude controller is proposed to resolve challenges encountered via dynamic scanning along highly non-planar surfaces and further extend applicability.

By contrast, the UAV-crawler hybrid vehicle imposes more stringent requirements on the asset to enable useful interaction. While its smaller form factor allows for operation in more confined regions, and high pose stability is demonstrated under thrust-supported contact, its current implementation is strongly dependent on cylindrical asset geometry. This arises as a consequence of using leg reaction forces to provide partial support and the decision to constrain MDT thrust generation to the body-relative vertical plane. Doing so, thrust efficiency was maximised when circumnavigating horizontal-axis pipe-like structures and the interaction model could be simplified to a 2D solution for minimum-energy support during initial investigation. However, this current implementation precludes vehicle deployment to cylindrical

structures with inclined or vertical axes, where the component of system weight acting parallel to the pipe axis exceeds the safety margin for support by tyre friction.

Therefore, an extension is proposed to revise the propeller array or investigate fully vectored thrust so as to enable 3D thrust generation. The interaction modelling and control strategy could then be updated to remove their dependency on the horizontal pipe assumption. Further modification to the flight control structure could also enable both hovering with arbitrary orientation and entry to thrust-supported contact from such poses. Fly-crawl-fly interaction with arbitrary surfaces including convex, planar, and concave elements in all orientations may then become feasible, providing a stable platform for remote airborne NDE and other tasks. Optional use of a tethered power supply may also increase maximum payload mass and operational duration, enhancing practical utility where limited range is a secondary concern and risk of environmental snagging is minimal. Lastly, upgrades to remove the non-holonomic constraints of the existing differential drive system via mecanum wheels or swerve drives would provide direct bi-directional control over locomotion while in contiguous contact with the asset surface. This would increase the vehicle ability to precisely deploy an NDE transducer relative to an arbitrary target point and perform complex scanning motions, supporting in-depth inspection of wind turbines, nuclear facilities, oil and gas platforms, and other high-value structures.

5.3.4 Additional NDE Sensory Modes

As introduced in Section 1.1, industrial inspection must detect an extensive range of flaw and defect types, occurring in various components and structures, and so employs numerous sensory technologies. Both the thrust vectoring aerial manipulator and the hybrid vehicle airborne inspection systems developed herein are demonstrated and characterised with a single sensor mode, performing ultrasonic thickness measurement and photogrammetric inspection, respectively. The characteristics and benefits of many other techniques used in wider NDE therefore remain unexplored.

Several opportunities to expand the ultrasonic measurement capability offered by the Voliro UAV. Development and integration of an algorithmic correction for V-path uncertainty would increase thin-wall measurement confidence and reduce the need for

application specific empirical calibration and compensation strategies, similarly improving practical utility. Incorporation of other transducers such as an EMAT sensor or linear phased-array would permit generation of alternate wave-modes and imaging techniques, as indicated in Appendix B. These could enable rapid measurement of extended surface areas and increase the detectable range of defect modes. The Voliro UAV is highly extensible in this regard, with sufficient maximum payload to accommodate a variety of NDE instrumentation [310].

Despite a lower payload, the UAV-crawler hybrid vehicle could also serve as a platform for additional NDE capability. Adaptation to support ultrasonic sensors could extend capability to volumetric inspection, including detection of material loss and internal crack formation. Beyond this and extending the other more general modes of NDE considered in Section 2.1.1, Pulsed Eddy Current (PEC) transducers present a synergistic avenue for further investigation, benefiting from the stability of the hybrid vehicle during their longer signal capture times while operating without significant contact forces that would alter the interaction model. PEC sensors are also noted for their established applicability to inspection of the severe degradation mechanisms commonly encountered by insulated pipes, storage tanks and pressure vessels [61], [66]. Providing means for relative wall thickness measurement without necessitating removal of non-conductive cladding or insulation, they may effectively monitor both flow accelerated corrosion and corrosion under insulation and in turn confer significant additional value when integrated with the hybrid vehicle.

Following appropriate verification and characterisation studies, both vehicles may thereby be adopted as highly capable means for practical NDE deployment, driving a reduction in cost to inspect, increasing asset insight, and improving operational safety.

5.3.5 Enhanced Automation

Capture of NDE data suitable to ascertain target status and rule out presence of critical flaws often requires expert knowledge of the asset, its degradation mechanisms, and likely failure points. Additionally, the NDE transducer will present requirements for successful data capture, such as camera focal distance or ultrasound probe orientation. Measurement locations must also be carefully recorded to ensure complete coverage

of the target region and guarantee the absence of flaws that can critically affect mechanical integrity. However, tracking all these aspects while piloting a UAV amid changeable environmental conditions represents a significant cognitive load.

Improvement in this regard may be found by incorporation of assistive automation features within the airborne contact inspection process, as shown by Omari *et al* [235], Quenzel *et al* [236], and Zhang *et al* [234] in Section 2.5.1 concerning free-flight visual inspection. Augmentation of the Voliro UAV flight controller to automatically maintain ideal contact force and orientation conditions for NDE data capture while the pilot directs the high-level flight path would simplify the process significantly. Further, online tracking of captured measurements versus a pre-defined inspection plan and introduction of on-the-fly adaptations as information is obtained could ensure complete coverage with minimal redundant data and reduce the need for repeat inspections with both UAV systems developed herein. Further, a fully automated solution incorporating such aspects within an exploratory SLAM algorithm and could enable full airborne contact-based inspection coverage with minimal prior knowledge, reducing onsite planning requirements. This could substantially improve the efficacy of airborne NDE operations, ultimately increasing cost efficiency for both the inspection service provider and asset operator.

5.3.6 BVLOS Operation & Site Residency

As indicated in Section 2.5.4.2, aerial inspection service providers are now beginning to investigate the operational efficiencies that can be attained by remote control of a UAV agent from beyond visual line of site and at an extreme distance. While these systems are in their infancy, requiring technological and legislative development to enable integration with other airspace users, a small number of entities have developed operable products. Currently, these are used for routine visual inspection or rapid response to emergent situations, providing a flexible strategy for monitoring of expansive facilities with reduced human presence.

Similar BVLOS or site-resident applications may be developed for contact-based airborne NDE systems following sufficient abstraction and automation of the inspection task. This would greatly reduce travel time, hazard exposure, and facility

access costs associated with deployment of multiple highly skilled inspection personnel, particularly in offshore scenarios. Moreover, operation from a centralised location allows NDE analysis to be conducted in consistently optimal conditions and by highly specialised experts, mitigating influences of human subjectivity. Thereby, overall quality of inspection may also be improved.

5.3.7 NDE Data Fusion

Where photogrammetry may reconstruct a photorealistic visual facsimile of a target structure exterior, ultrasound is noted for its ability to image the internal structure. The two methods are therefore highly complementary within an inspection context and together may provide a far more detailed report on asset status than possible with either operating alone. With UAV application of both of these NDE modes successfully characterised throughout Chapter 3 and Chapter 4 of this thesis, an opportunity is now presented to combine these techniques with a highly mobile agent for airborne data capture.

This could be realised through a heterogenous swarm, or by deploying multiple sensors aboard the same UAV. As such, the interconnected systems may function synergistically. The onboard position and attitude estimates employed by the UAV during flight could provide a coarse initial value for the camera poses determined as part of photogrammetric reconstruction process. Processing times may then be substantially reduced by pre-selecting images with potential matching features based on spatial proximity. Moreover, the refined camera pose returned following reconstruction may also present an effective means to spatially encode synchronously captured ultrasonic readings relative to the 3D asset, improving accuracy versus real-time position estimates and aiding comparison across multiple inspections.

Amid growing adoption of digital twins and online reporting software, this could prove highly beneficial to workflow efficacy in commercial NDE. Where a single sensing mode presents an anomaly that is challenging to discern or characterise within reporting requirements, co-located data from another transducer could provide rapid clarification, also drawing on situational context from the 3D model. Benefits to

tracking asset state over time are also presented, granting a common reference for historical data and supporting statistical analytics across multiple plant items.

Ultimately, multiple avenues for further research and development exist that may aid realisation of the great promise noted in airborne contact-based NDE systems. Going forward, these technologies may be expected to help UAVs perform safe, detailed, and cost-effective in situ assessment across multiple applications within the energy sector and beyond.

APPENDIX A

Photogrammetric Theory

This appendix provides a consolidated reference concerning the operational theory behind the function of a camera, as below in Section A.1, and its application to photorealistic digital reconstruction of target objects and landscapes, following thereafter in Section A.2. Common aspects relevant to general UAV pose estimation and mapping strategies are also noted to provide full context as to the practical utility of these concepts.

Further information concerning the pinhole model commonly used to describe a camera may be obtained via the authoritative work by Hartley and Zisserman [107], or the more accessible introduction presented by Corke [108]. Similarly, a number of resources describe 3D reconstruction in detail beyond what may be included herein, such as the popular tutorial by Furukawa and Hernandez [371], the comparison of 3D camera properties by Giancola *et al* [118], or the thorough review of its application to geological survey is given by Smith *et al* [115].

A.1 The Pinhole Camera Model

A.1.1 Simple Linear Representation

At its most abstract level, a camera exists as a mapping function between a point in 3D space, $\mathbf{X} = [X \ Y \ Z]^T \in \mathbb{R}^3$, and a 2D location in the image, $\mathbf{x} = [x \ y]^T \in \mathbb{R}^2$, formally described as $\mathbf{X} \mapsto \mathbf{x}$.

This mapping is simplest in the case of central projection as depicted in Figure A-1, where the camera centre, i.e. the pin hole, lies on the origin of the world coordinate frame $\{W\}$ and the virtual image plane is centred on the z-axis at an offset of focal length, F , and parallel to the x-y plane. Note that hereafter the virtual image plane in

front of the camera will be used as a simplifying analogue for the true, inverted, image plane behind the camera. The principal axis is then defined as the line from the camera which perpendicularly intersects this plane, with the principal point defining the point of intersection.

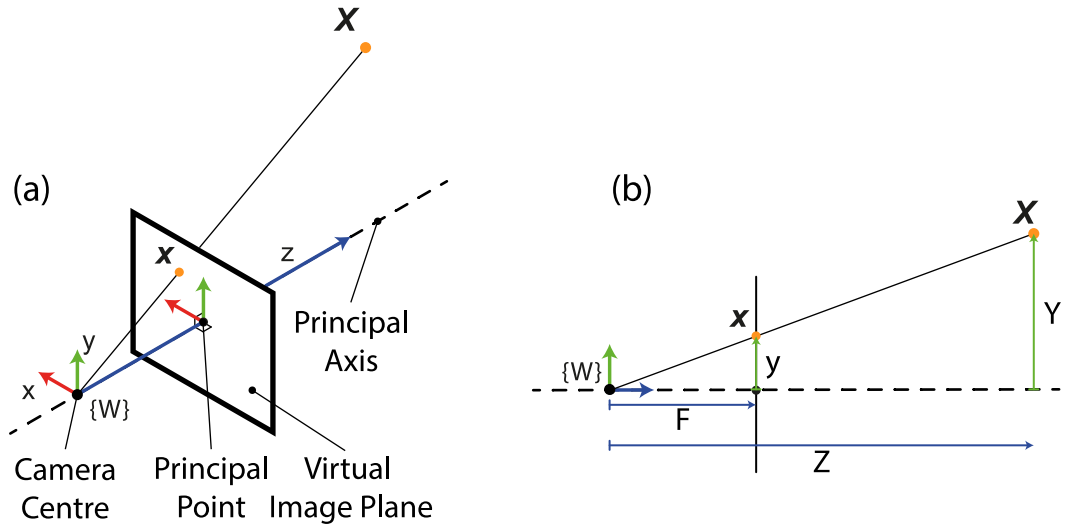


Figure A-1: The simplest central projection pinhole camera model, linearly mapping a 3D point to a position on the image using similar triangles. (a) Isometric view. (b) Side view. The definition of key terminology regarding the camera centre, principal point, and principal axis is also illustrated.

In recognition of the similar triangles formed through linear projection, the mapping operation $X \mapsto x$ may be represented by a matrix product with the camera matrix, P , utilising the homogeneous coordinate representation of x and X , as indicated by the \sim accent, as follows.

$$\tilde{x} = P\tilde{X} \tag{A-1}$$

$$\underbrace{\begin{bmatrix} FX \\ FY \\ Z \end{bmatrix}}_{\tilde{x}} = \underbrace{\begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_P \underbrace{\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}}_{\tilde{X}} \tag{A-2}$$

Note that homogeneous coordinates describe the location of a point in projected space, denoted \mathbb{P} [108]. As such, a point in 2D Euclidean space, $[x \ y]^T \in \mathbb{R}^2$, may be represented in homogenous coordinates as $[x_1 \ x_2 \ x_3]^T \in \mathbb{P}^2$, subject to

$$x = \frac{x_1}{x_3}, \quad y = \frac{x_2}{x_3}, \quad x_3 \neq 0 \quad (\text{A-3})$$

The additional degree of freedom introduced in \mathbb{P}^2 therefore imparts the useful property that $k[x_1 \ x_2 \ x_3]^T$ will appear in the same location as $[x_1 \ x_2 \ x_3]^T$ on return to \mathbb{R}^2 , where k is a non-zero scalar. This similarly extends to 3D points in space, such as applied to \mathbf{X} in (A-2). Returning the homogeneous result of (A-2) to 2D Euclidean space within the image plane, it is then apparent that the result agrees with that expected from the similar triangles representation.

$$\mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} FX/Z \\ FY/Z \end{bmatrix} \quad (\text{A-4})$$

Per this relationship, increasing the focal length, and thus the distance of the image plane from the camera centre, will shrink the field of view, as indicated by the size of plane at the rear of the 3D scene in Figure 2-1. With a consistently sized image plane this serves as the mechanism for optical zoom in a physical camera, though it should be noted that F in this case is changed by re-arranging internal lenses and not physical displacement of the sensor.

A practical camera, however, is subject to further influences that alter the expression of \mathbf{P} given in (A-2), such as those depicted in Figure A-2. These are reviewed below to establish a more representative model.




Figure A-2: A more practical pinhole camera model. Here the camera has arbitrary position and orientation relative to $\{W\}$, indicated by translation vector δ and rotation matrix R . The projected point is also expressed in pixel coordinates relative to the top left of the image and the principal axis does not pass through the centre of the image plane, instead existing at (u_0, v_0) .

A.1.2 Principal Point Offset & Pixel Space

First among these effects is the conversion from spatial coordinates to pixel coordinates within the image sensor. This is modelled via scale factors, defining the number of number of pixels per unit length in the both the horizontal and vertical directions as ρ_x and ρ_y , respectively, each having a different value if the sensor has non-square pixels. Additionally, owing to manufacturing tolerances, a spatial offset may exist between the centre of the image plane and the principal point. Images are also conventionally indexed such that the pixel at the top left is numbered $(0,0)$. An offset is therefore applied in the image x-axis and y-axis to account for these combined effects within the pixel space. The projected point x is thereby expressed in Euclidean pixel coordinates as

$${}^p\tilde{x} = \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \rho_x x + u_0 \\ \rho_y y + v_0 \end{bmatrix} \quad (\text{A-5})$$

A.1.3 Arbitrary Camera Pose

A second aspect of practical cameras is that they may be positioned arbitrarily relative to the world coordinate frame, as in Figure A-2. Consequently, it is useful to define a camera relative coordinate frame, $\{Cam\}$, such that the camera centre lies at the origin,

\mathbf{O}_{Cam} , and its principal axis aligns with the $\{Cam\}$ frame z-axis. The pose of the camera is then defined relative to the world frame by a 3x3 rotation matrix, ${}^W_{Cam}\mathbf{R}$, and a translation offset, ${}^W_{Cam}\boldsymbol{\delta}$. Together these may form a homogeneous transformation matrix, ${}^W_{Cam}\mathbf{T}$, expressing the camera position relative to the world coordinates.

$${}^W_{Cam}\mathbf{T} = \begin{bmatrix} {}^W_{Cam}\mathbf{R} & {}^W_{Cam}\boldsymbol{\delta} \\ \mathbf{0} & 1 \end{bmatrix} \quad (\text{A-6})$$

Note that the inverse of this transformation matrix may be used to express arbitrary positions in $\{W\}$, denoted ${}^W\mathbf{X}$, relative to the camera frame using their homogeneous form as follows.

$${}^{Cam}\tilde{\mathbf{X}} = ({}^W_{Cam}\mathbf{T})^{-1}{}^W\tilde{\mathbf{X}} = {}^{Cam}_W\mathbf{T} {}^W\tilde{\mathbf{X}} \quad (\text{A-7})$$

Thereby, the centre projection model may be employed in a camera with arbitrary pose in the parent world frame, as is advantageous for mobile robotic deployment.

A.1.4 Revised Camera Model

A more practical linear camera model is then found by combining the above effects within a camera calibration matrix, \mathbf{K} , used alongside the transformation matrix, ${}^{Cam}_W\mathbf{T}$, and the identity projection matrix, \mathbf{P}^* . The projected pixel space location, ${}^p\mathbf{x}$, of point in $\{W\}$, ${}^W\mathbf{X}$, may then be represented as follows, again utilising homogeneous coordinates [108].

$$\underbrace{{}^p\begin{bmatrix} u \\ v \\ Z \end{bmatrix}}_{p\tilde{\mathbf{x}}} = \underbrace{\begin{bmatrix} \alpha_x & s & u_0 \\ 0 & \alpha_y & v_0 \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{K}} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\mathbf{P}^*} \underbrace{\begin{bmatrix} {}^W_{Cam}\mathbf{R} & {}^W_{Cam}\boldsymbol{\delta} \\ 0 & 1 \end{bmatrix}}_{{}^{Cam}_W\mathbf{T}} \underbrace{\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}}_{{}^W\tilde{\mathbf{X}}} \quad (\text{A-8})$$

Here, the \mathbf{K} matrix combines the pixel coordinate conversion of (A-5) with the existing focal length diagonal terms of \mathbf{P} in (A-2) to produce scalars $\alpha_x = \rho_x F$ and $\alpha_y = \rho_y F$. A further skew term, s , is introduced to describe any non-orthogonality in the image sensors axes and commonly takes the value of zero owing to the precision of modern semiconductor manufacturing [108]. The overall mapping function of the camera $\mathbf{X} \mapsto \mathbf{x}$ is then described by the combination of its intrinsic effects, occurring within the

camera itself, and extrinsic effects, referring to its pose relative to the world scene, yielding the updated camera matrix \mathbf{P} , concisely written as follows,

$$\mathbf{P} = \mathbf{K}\mathbf{P}^*{}_{\mathit{Cam}}^{\mathit{W}}\mathbf{T} \quad (\text{A-9})$$

An arbitrarily positioned camera with representative, non-ideal, sensor characteristics may thereby be modelled and utilised as in (A-1).

A.1.5 Lens Distortion

Prior to this point, the camera model has represented a linear function, but a practical camera also introduces non-linear distortion due to lensing effects. This can take many forms including: chromatic aberration, where focal length varies with incident light wavelength creating colour fringing; spherical aberration, where focal length varies across the lens radial direction and blurs the image edges; and astigmatism, where rays from perpendicular planes have different foci, causing blur in one axis [108]. In turn, these aspects can have significant impact on visual NDE acuity, influencing both the visibility of surface texture artefacts and apparent target geometry.

To minimise their degressive effects on the inspection process, the displacement of the projected ray in the image plane under the previous model is thus modelled and used to correct the image. Commonly this is performed as a function of two distortion components: radial and tangential, as depicted in Figure A-3.




Figure A-3: Image distortion components. The black grid shows the original image, while the orange shows the result of the distortion component indicated. Barrel distortion is referred to as “negative” as the distorted location moves increasingly closer to the centre at larger radii. Similarly, pincushion distortion is termed “positive”.

The radial distortion occurs at the lens extremities, where light rays are more substantially deflected than those nearer the centre and is commonly modelled as a polynomial function as here, following the Brown-Conrady method [108], [372], or using a Fitzgibbon division model [373]. The change in projected location of a light ray is then expressed as a function of its radius from the principal point, $r = \sqrt{(u')^2 + (v')^2}$, expressed in pixel dimensions where $u' = u - u_0$ and $v' = v - v_0$.

$$\begin{bmatrix} du \\ dv \end{bmatrix}_r = \begin{bmatrix} u'(k_1 r^2 + k_2 r^4 + k_3 r^6 + \dots) \\ v'(k_1 r^2 + k_2 r^4 + k_3 r^6 + \dots) \end{bmatrix} \quad (\text{A-10})$$

This function may be expanded to arbitrary powers of r , according to the precision desired, but is commonly parameterised by three scalar coefficients, (k_1, k_2, k_3) .

In contrast, tangential distortion arises where the lens is non-parallel to the image sensor. It acts in a direction perpendicular to the radial line to shed the image from the sensor centre and is modelled as below [372] using two scalar coefficients (p_1, p_2) .

$$\begin{bmatrix} du \\ dv \end{bmatrix}_t = \begin{bmatrix} 2p_1 u' v' + p_2 (r^2 + 2(u')^2) \\ p_1 (r^2 + 2(v')^2) + 2p_2 u' v' \end{bmatrix} \quad (\text{A-11})$$

Their additive effect is then represented by the nonlinear function $\mathbf{dx}(u, v)$, combining (A-10) and (A-11) to yield the total displacement in the projected point after non-linear distortion, ${}^d \mathbf{x}$.

$${}^d \mathbf{x} = {}^p \mathbf{x} + \mathbf{dx}(u, v) \quad (\text{A-12})$$

$$\mathbf{dx}(u, v) = \begin{bmatrix} du \\ dv \end{bmatrix}_r + \begin{bmatrix} du \\ dv \end{bmatrix}_t \quad (\text{A-13})$$

The five parameters defining the distortion, $(k_1, k_2, k_3, p_1, p_2)$ are then considered as additional intrinsic parameters, alongside those of \mathbf{K} in (A-8), and as such may be identified by a camera calibration algorithm to permit image correction. Dedicated software tools for such procedures are widely available [374]–[377]. Factors within the camera affecting the final image quality and the accuracy with which it depicts the exterior scene are thereby identified to enable accurate sizing of single items, surface flaws, and large structure geometry in NDE applications.

A.1.6 Multiple Viewpoints & Epipolar Geometry

When examining two or more cameras via the model described in (A-8), an important consequence is their ability to triangulate a mutually visible point in 3D space, extending imaging capability beyond the 2D representation described above. This utilises epipolar geometry, defined by Hartley and Zisserman [107] as the intrinsic projective geometry between two views, independent of scene structure and dependent only on the internal structure and relative pose of the cameras. This also forms the basis for 3D sensing technologies such as RGB-D cameras, stereo vision, and fixed multi-camera arrays.

In the most general case, this geometric relationship between cameras is represented by the fundamental matrix, $\mathbf{F} \in \mathbb{R}^{3 \times 3}$, a rank 2 homogeneous matrix simplifying to the essential matrix, $\mathbf{E} \in \mathbb{R}^{3 \times 3}$, where the calibration matrix of each camera is known. A brief overview of the origin and consequences of this relationship follows.




Figure A-4: Epipolar geometry for multi-view point correspondence. Two cameras image the same arbitrary point in space, \mathbf{X} , appearing in their images at \mathbf{x}_1 and \mathbf{x}_2 respectively. Each camera centre, \mathbf{O}_1 and \mathbf{O}_2 is visible to the other camera, projected as \mathbf{e}_1 and \mathbf{e}_2 respectively. These intersections between the baseline linking the two centres and the image planes are termed epipoles. Epipolar lines \mathbf{l}_1 and \mathbf{l}_2 then link the projected point and the epipole, radiating from the later for arbitrary \mathbf{X} .

From Figure A-4, it may be observed that given only its projected point in Image 1, \mathbf{x}_1 , its location in 3D space, \mathbf{X} , cannot be exactly determined by reprojection, but will lie along the ray defined by $\tilde{\mathbf{x}}_1 = \mathbf{P}_1 \tilde{\mathbf{X}}$, passing through the camera origin \mathbf{O}_1 and \mathbf{X} . In homogeneous coordinates, $\tilde{\mathbf{X}}$ defines this ray and may be calculated using the pseudo-inverse [218] of the camera matrix as below, yielding a line on return to 3D Euclidean space.

$$\tilde{\mathbf{X}} = \mathbf{P}_1^\dagger \tilde{\mathbf{x}}_1 \quad (\text{A-14})$$

In the presence of a second camera also observing the 3D point, \mathbf{O}_1 and \mathbf{X} may be projected into Image 2 as follows. Note that the projection of a camera origin to another image is a special case resulting in an epipole, here \mathbf{e}_1 .

$$\tilde{\mathbf{e}}_1 = \mathbf{P}_2 \tilde{\mathbf{O}}_1, \quad \tilde{\mathbf{x}}_2 = \mathbf{P}_2 \tilde{\mathbf{X}} = \mathbf{P}_2 \mathbf{P}_1^\dagger \tilde{\mathbf{x}}_1 \quad (\text{A-15})$$

The ray projected from \mathbf{x}_1 will then appear as an epipolar line, \mathbf{l}_1 , stretching across Image 2 and containing both the epipole \mathbf{e}_1 and \mathbf{x}_2 as in Figure A-4.

To express \mathbf{l}_1 mathematically, note that the points \mathbf{O}_1 , \mathbf{O}_2 and \mathbf{X} may define a plane containing both the ray projected from Image 1 and the baseline between the two camera centres, termed the epipolar plane. As it contains the camera centres, the plane is visible as a line in either Image 1 or Image 2. Accordingly, \mathbf{l}_1 may be calculated by defining the plane via the vector cross product, with \mathbf{l}_2 expressed similarly.

$$\mathbf{l}_1 = \tilde{\mathbf{e}}_1 \times \mathbf{P}_2 \mathbf{P}_1^\dagger \tilde{\mathbf{x}}_1 = \mathbf{F} \tilde{\mathbf{x}}_1 \quad (\text{A-16})$$

$$\mathbf{l}_2 = \tilde{\mathbf{e}}_2 \times \mathbf{P}_1 \mathbf{P}_2^\dagger \tilde{\mathbf{x}}_2 = \mathbf{F}^T \tilde{\mathbf{x}}_2$$

From which the fundamental matrix, \mathbf{F} , is defined by grouping the cross-product terms.

$$\mathbf{F} = \tilde{\mathbf{e}}_1 \times \mathbf{P}_2 \mathbf{P}_1^\dagger \tilde{\mathbf{x}}_1 \quad (\text{A-17})$$

Furthermore, in lying on \mathbf{l}_1 , the projected point \mathbf{x}_2 will satisfy the plane equation $\mathbf{x}_2^T \mathbf{l}_1 = 0$. Substituting (A-16), then yields the expression below [107].

$$\tilde{\mathbf{x}}_2^T \mathbf{F} \tilde{\mathbf{x}}_1 = 0 \quad (\text{A-18})$$

This is referred to as the correspondence condition and must be satisfied by all points in multi-view images projected from the same point in the depicted scene. The utility of (A-18) is then that the fundamental matrix may be determined directly from the point correspondences within the images without prior knowledge of either camera matrix, proving highly advantageous for photogrammetric 3D reconstruction as examined in Section A.2.

Where the camera calibration matrices are known, a similar condition may be derived for the essential matrix using normalised projections in the homogeneous space such that $\hat{\mathbf{x}} = \mathbf{K}^{-1} \tilde{\mathbf{x}}$.

$$\hat{\mathbf{x}}_2^T \mathbf{E} \hat{\mathbf{x}}_1 = 0 \quad (\text{A-19})$$

Comparing with (A-18), the resulting essential matrix, \mathbf{E} , may be related to the fundamental matrix via the expression,

$$\mathbf{E} = \mathbf{K}_2^T \mathbf{F} \mathbf{K}_1 \quad (\text{A-20})$$

Thereby, it is similarly possible to determine the relative pose of two cameras from their corresponding image points, with dimensionality of the problem reduced by the application of known calibration matrices.

A.2 3D Reconstruction

Informed by the camera model presented above, extraction of 3D geometry from multiple images of a scene or object is then possible. This may be conducted via the Structure from Motion and Multi-View Stereoscopy (SfM-MVS) processes to generate a high-fidelity digital representation of the target, as is the primary interest in NDE applications [115], [371], or the mathematical concepts involved may serve additional wider purposes for navigation of mobile robotics. As such, a real-time position

estimate of an airborne UAV system may be obtained via a number of Visual Odometry (VO) [313], [378] or Simultaneous Localisation and Measurement (SLAM) [379], [380] methods. Together, these algorithms can then be regarded as a spectrum of machine vision technologies, as summarised in Figure A-5.

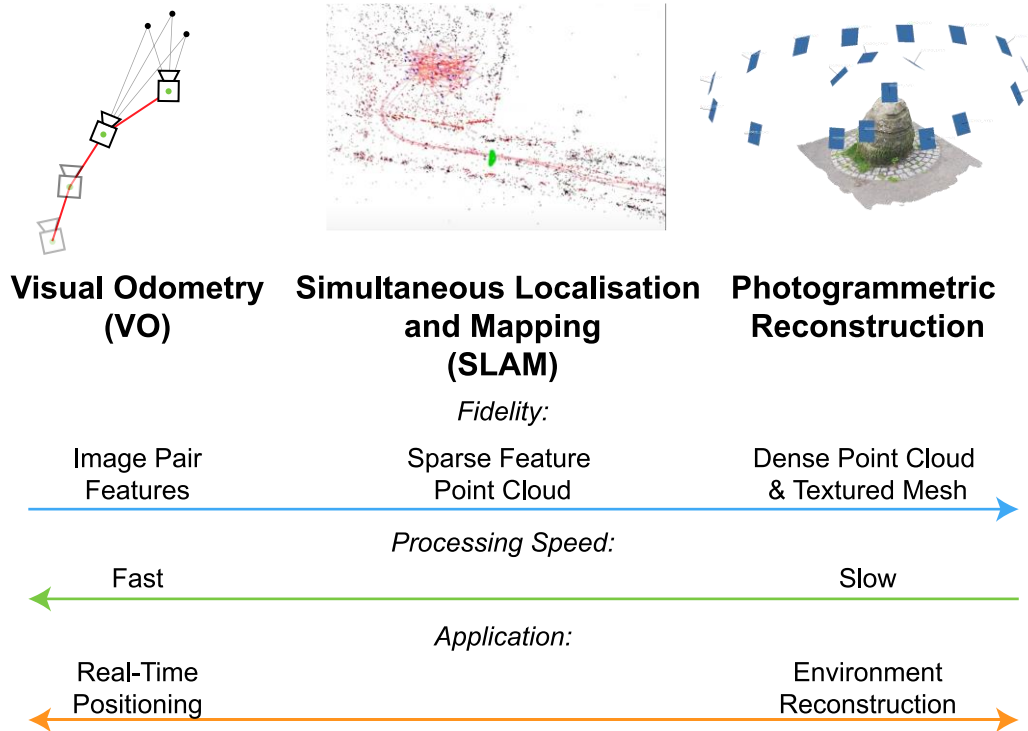


Figure A-5: High-level comparison of the types of machine vision used by UAVs. Visual Odometry accumulates the relative pose between sequential images. SLAM [380] extends this, mapping the environment to identify previously visited locations and optimise the camera trajectory, negating drift via loop closures. Photogrammetry further densifies the scene reconstruction to provide photorealistic output for use in other applications but does not operate in real-time.

With a primary focus on NDE applications, this section accordingly presents an overview of the SfM-MVS workflow as depicted in Figure A-6. Crossover with VO and SLAM application is then highlighted where appropriate to inform the practical distinction between the techniques.

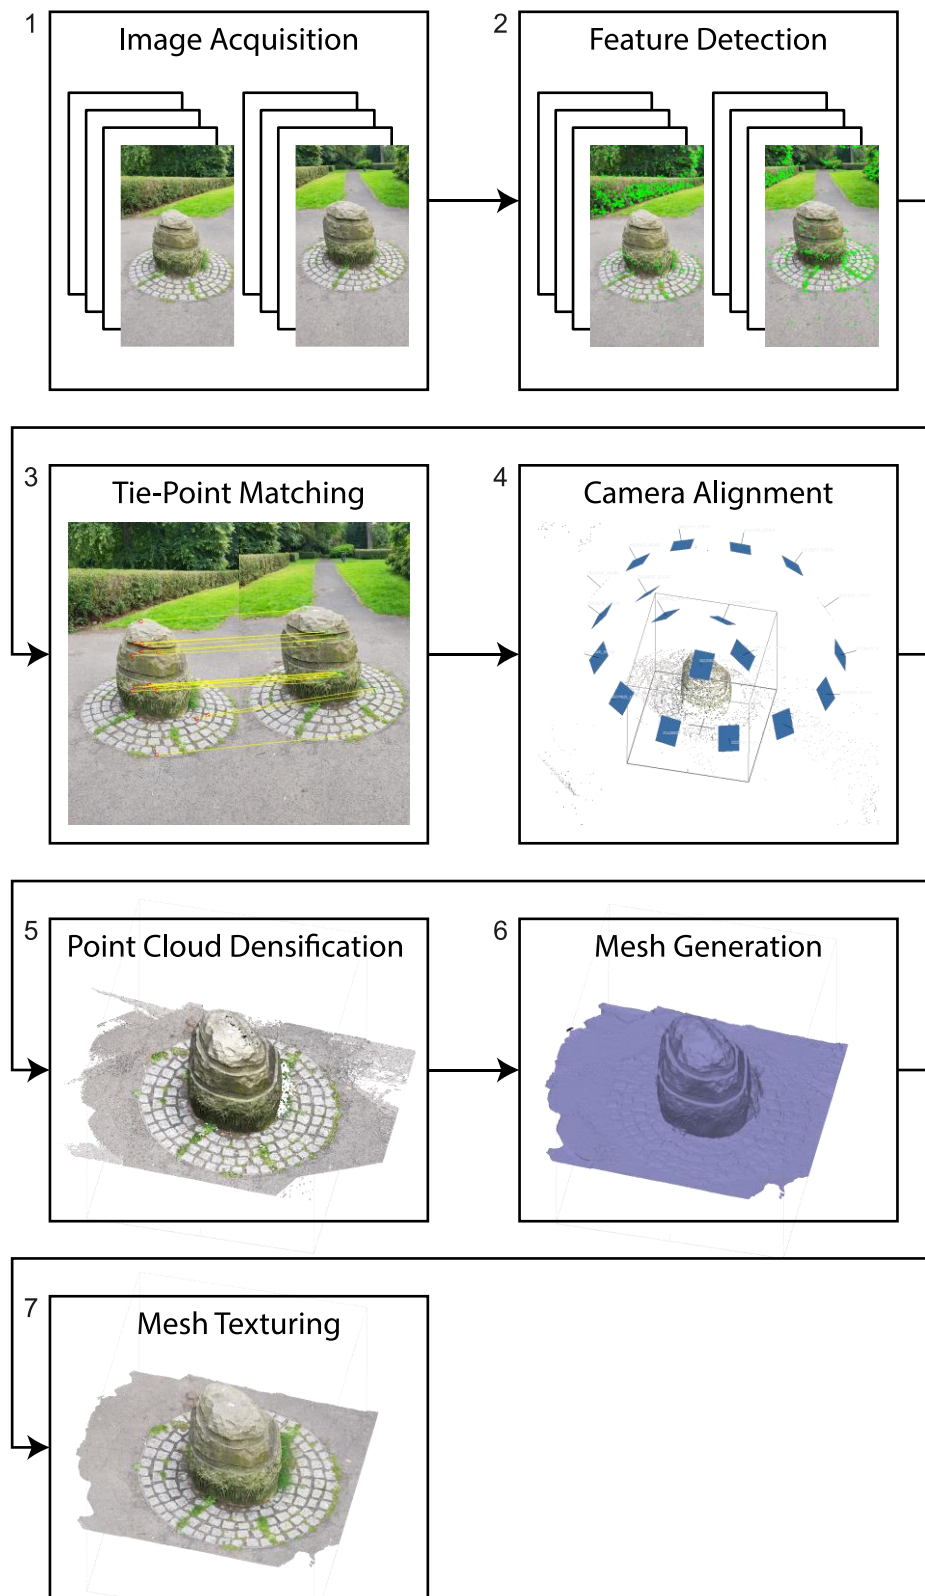


Figure A-6: A generic photogrammetric reconstruction workflow. Numbers indicate the sequence of sub-processes undertaken to generate a final textured mesh output from a series of input RGB images.

A.2.1 Image Acquisition

Initially, a number of images are taken of a scene or object to be reconstructed from different viewpoints. The reconstruction algorithm is not influenced by the mode of their capture, but image content must include the complete outer surface of the reconstruction target with sufficient overlap between multiple images to match visible features. When deployed from a robotic arm or UAV for reconstruction, this step commonly involves a prior path-planning stage to ensure complete coverage. Multiple path planning algorithms exist as a separate field of research, including mobile robotic operator/pilot lead heuristics and offline automated strategies relying on approximately known target geometry and position, as reviewed in detail by Aggarwal and Kumar [212]. In both VO and SLAM processes image acquisition will occur continuously, dropping the requirement for full surface coverage to instead emphasise real-time camera pose updates.

A.2.2 Feature Detection

With a completed set of images, typically numbering in the hundreds for small objects, or thousands and tens of thousands for high-fidelity building and extended site representations, photogrammetric reconstruction requires a means to identify which images correspond to the same physical scene within the target. Accordingly, the image set is processed to identify salient features and characteristic regions within each image such as edges, corners, and unique surface textures.

Any of a number of feature detection algorithms may be used, include the Scale-Invariant Feature Transform (SIFT) [381], Speeded-Up Robust Features (SURF) [382], Binary Robust Invariant Scalable Key-points (BRISK) [383] and Oriented FAST and Rotated BRIEF (ORB) [384] descriptors. Note, ORB is a combination of the Features from Accelerated Segment Test (FAST) [385] feature detection method with the Binary Robust Independent Elementary Features (BRIEF) descriptor [386].

When applied to an image, each feature description algorithm will return the pixel location of each detected visual feature alongside a descriptor used to identify it. In the case of SIFT, this is a 128 element vector of integers. SURF uses a 64 element floating

point vector, while ORB and BRISK both use a 512 bit binary string to enable faster feature comparison. Step 2 of Figure A-6 depicts an example of SIFT feature detection, showing the most distinctive 500 features identified in the images as green dots.

In the SfM-MVS context, key performance attributes are then uniqueness of these feature descriptors and their consistency across changes to the object position, scale, rotation, illumination, and viewing direction. This can present a challenge in NDE applications owing to the self-similarity of high texture surfaces, such as brickwork or concrete, or due to feature sparsity in large asset surfaces. As reviewed by Tareen and Saleem [387], SIFT produces the highest overall feature matching accuracy in the presence of geometric transforms across multiple test image sets, and so is best suited for photogrammetry, but ORB and BRISK are most time efficient per identified feature.

Features descriptors may thereby be selected to balance reconstruction quality with processing speed, also permitting real-time use in VO and SLAM applications where speed has higher emphasis [380]. Note also that projective RGB-D sensors need not rely on visual feature descriptors to the same extent as they actively emit a dot matrix that is directly visible in an infrared image.

A.2.3 Tie-Point Matching

Identifying and matching repeat views of the same visual features between images follows next in the SfM pipeline. This is conducted based on similarity of the feature descriptors in each image, evaluated using Manhattan or Euclidean distance in the case of the vector descriptors of SIFT and SURF, but using the faster Hamming distance with the ORB and BRISK binary strings [115]. Comparisons are then ranked by shortest distance, taking the nearest neighbour as a match or discarding the feature if the ratio of distances between the nearest and second nearest neighbour known to be from another object falls outside a prescribed threshold [381]. However, as there can be tens of thousands of features per image and thousands of images, exhaustive feature to feature comparison can quickly become intractable.

Optimisations are therefore applied to speed matching across the large image sets necessary for building reconstruction. A popular algorithm by Muja and Lowe [388] sees the global list of SIFT features pre-binned into multiple k-dimensional trees using randomised descriptor elements. Only bins nearest the descriptor under test are then checked for nearest neighbours, abandoning the search after a specified limit, while other trees can be checked in parallel. This has been shown to offer up to a 100x speedup versus exhaustive search with under a 10 % loss in nearest neighbour accuracy [388], better enabling photogrammetry for NDE of large scale industrial assets.

The features matches across images are then used to produce a complete set of 2D image tracks, listing the images in which a consistent feature appears and its pixel coordinates in each [371]. These are represented visually in step 3 of Figure A-6 as the yellow lines linking the image pair.

It is at this stage the distinction between visual odometry and SLAM becomes apparent. Whereas VO attempts to incrementally estimate relative camera pose by ensuring local consistency between sequential images, SLAM instead seeks global consistency between a world map and its camera pose within it [389]. Accordingly, VO conducts feature matching over a minimal number of sequential images, quickly yielding a pose estimate subject to drift accumulation. SLAM, by contrast, performs intensive matching across a library of previously geolocated key features to identify locations visited earlier and thereby further refine its non-drifting camera pose estimate via loop closure. The choice of which to implement is then dependent on available compute power aboard a UAV agent and pose accuracy requirements.

A.2.4 Camera Alignment

The next stage in the reconstruction process concludes its SfM portion. Here the camera matrices are determined for each captured image, incorporating the full pose in a common coordinate system and calibration matrix parameters as given in Section A.1.4. A sparse feature-based point cloud is also generated, as in step 4 of Figure A-6.

As such, camera alignment is accomplished using the matching feature tracks obtained previously, with images sharing multiple tracks identified as stereo pairs. As examined in Section A.1.6, the epipolar geometry of the image pair then permits recovery of their fundamental matrix via the feature tracks [107]. However, owing to the similarity of visual features, it is common for direct visual feature comparison to return several false-positive “outlier” matches. These strongly influence the least-squared error metrics used to optimise the sparse reconstruction and must therefore be removed to prevent significant degradation.

Outlier rejection is commonly performed via RANdom SAMple Consensus (RANSAC) [390]. Leveraging the geometric correspondence condition expressed in (A-18), this iteratively takes the minimum randomised subset of 7 matching features necessary to define a candidate fundamental matrix between images sharing multiple feature tracks [107]. Feature pairs falling within a set tolerance of the epipolar condition are denoted inliers, agreeing with the candidate \mathbf{F} , while false matches fall outside this tolerance. Repeating these steps, the fraction of inlier features is maximised to within a pre-set confidence interval. Outlier tracks identified by the maximum inlier candidate \mathbf{F} are then rejected from the image pair, before the optimal \mathbf{F} is determined using Maximum Likelihood Estimation (MLE). Employing all remaining inlier matches, this uses the Levenberg–Marquardt optimisation algorithm to minimise reprojection error between the observed feature location in an image and its location when projected from a stereo pair image [107]. A similar approach can be taken using the essential matrix condition for pre-calibrated cameras, via the relationship given in (A-19). Repeating this process for each stereo pair across the image set yields a set of geometrically verified feature tracks and fundamental matrices, which may be decomposed to provide the camera matrices for each image and used to give an initial estimate of the 3D feature points via triangulation.

A bundle adjustment step then forms the final phase of SfM camera alignment [107], [115], [391]. Applied to refine and minimise the reprojection error across images, this provides non-linear least squares optimisation to the camera pose, intrinsic matrices, and any non-linear distortion effects, such as those in Section A.1.5. Its name arises from the bundle of rays re-projected from the cameras during this process. While a

number of approaches exist [391]–[393], as recently reviewed by Bianco *et al* [130], the consistent objective is to find the set of camera matrices and observed 3D points minimising total reprojection error versus the image features. This can be applied incrementally, starting with the stereo pair containing the largest number of geometrically verified tracks, then re-applying to register subsequent images with earlier parameters frozen, growing the reconstruction accordingly. A final global adjustment then refines the alignment solution across all parameters to produce the optimal camera matrices and sparse point cloud, as depicted in step 4 of Figure A-6.

Note, however, that the reconstruction is only complete up to a similarity transform, lacking scale information owing to the nature of projective geometry. For monocular images, this can be provided by control points in the scene with known absolute position or via the image capturing system, invoking the joint encoder pose estimate of a robotic effector or the GNSS position of a UAV. By contrast, the known calibrated extrinsic transform between images of each sensor in a stereo camera provides sufficient scaling information to allow direct reconstruction with accurate physical dimensions. Similar benefit is found in the depth measurements of RGB-D sensors.

In the case of VO and SLAM, the algorithm will typically cease at this point, having robustly determined the camera location, and commence the next iteration with a newly acquired image, forgoing more detailed scene reconstruction.

A.2.5 Point Cloud Densification

This stage is followed by application of a Multi-View Stereo algorithm, utilising the bundle adjusted camera pose and calibration parameters alongside the stereo pairs and reprojected 3D feature points from the SfM stage to densify the output point cloud, increasing spatial resolution and allowing the output reconstruction to resemble the 3D target more closely. These point cloud models may be applied at the scale of individual components, small objects as in Figure A-6, or entire landscapes as in Figure A-7.



Figure A-7: A dense point cloud of a countryside house produced using commercial reconstruction software, Pix4D [127].

As with SfM, multiple MVS algorithms exist and are selected according to the needs of the specific application. As reviewed by Seitz *et al* [394], examples include the early work by Goesele *et al* [395], the popular Clustered MVS and Patch MVS (CMVS-PMVS) toolchain by Furukawa *et al* [396], [397], and the more recent Shading aware MVS (SMVS) by Langguth [398]. Typically, these grow the reconstruction around the verified 3D feature points to produce depth map images, defining the distance from the camera centre to the object in a ray through each pixel and so mimicking the output of an RGB-D camera.

This process leverages epipolar geometry, as in Section A.1.6, to perform efficient stereo matching of image regions near the tie-features across stereo pairs, employing pixelwise matching [395], encompassing visual features less distinctive but more abundant than those of Section A.2.2, including edges and corners [396], or expanding this to use shading gradients [398]. Intermediate representations of the target surface using multiple planar [396] or bilinear patches [398], provide a further opportunity to filter results for visual consistency and reduce noise. These patch-wise surface representations are then polled to return the depth images and dense point clouds. This also forms the basis for 3D imaging using a stereo camera, explaining their high

reliance on surface colour and texture variation to derive a depth image, and limiting their functionality regarding industrial structures with uniform appearance.

Where spatial geometry information is the sole interest of the NDE process, reconstruction may stop following this stage and examine the dense output point cloud to track changes in localised geometry such as chipping and erosion or large-scale structural deformation. Such evaluation is readily supported by tools such as Cloud Compare [399].

A.2.6 Mesh Generation

In applications where the visual surface texture provides useful insight, as with staining, rust formation, or corrosion, further data processing may generate a photorealistic mesh to allow defect inspection with fully reconstructed environmental context. As such, this stage of the reconstruction process derives the untextured geometric mesh, consisting of numerous interlinked triangular faces that define a sealed hull around the target geometry, as in step 6 of Figure A-6, ahead of subsequent surface texturing.

The untextured mesh may be constructed using either the point cloud or camera depth maps generated previously via algorithms such as Volumetric Range Image Processing (VRIP) by Curless and Levoy [400], Poisson Surface Reconstruction (PSR) by Kazhdan *et al* [401], [402], Floating Scale Surface Reconstruction (FSSR) by Fuhrmann and Goesele [403], or more recently Global Dense Multiscale Reconstruction (GDMR) by Ummenhofer and Brox [404].

Whereas these differ in their specific optimisations for processing time, noise rejection, hole filling, and sample scale variations, their core function is broadly similar. Using the dense point cloud directly or projecting the depth map images into 3D space and aggregating them into voxel groups, surface normal vectors can be estimated via the observing camera pose. An indicator function is then derived as a scalar field, valid throughout the reconstruction volume. Its gradient typically correlates to the surface normal vectors, while its absolute value is positive in free space outside the object or negative within its surface. This may act as a Signed

Distance Function (SDF) giving the distance from any point in space to the nearest surface. The isosurface where the SDF crosses zero then defines the object exterior surface, allowing the surface mesh to be extracted via a marching-cubes-like algorithm [405], [406].

The overall process thus yields a mesh preserving fine-scale surface detail while representing complex arbitrary geometry with multiple concave regions and holes. Generally, the additional redundancy of the depth map approach offers a higher quality reconstruction, but using the point cloud ahead of mesh generation enables selective pruning of low-confidence points poorly depicted in the image and segmentation of individual assets for targeted NDE [407], [408].

A.2.7 Mesh Texturing

To generate the final photorealistic mesh, visual texture is applied based on the observations of the surface made in the input image set to produce a texture map. This links each vertex of the 3D mesh faces to a point in a 2D texture image containing the visual surface detail and is also referred to as “UV-mapping” in reference to the (u, v) pixel coordinates of the vertices in the texture image.

As such, the texture of each mesh face may be obtained by back projection from the aligned images in which it is observed, ignoring faces occluded behind others in the object. Where perspective effects leave gaps in the texture, multiple redundant views provide robust means to ensure complete coverage. Blending multiple views may also offer the potential to increase surface sampling resolution via averaging weighted by their angular offset from the face normal, as first proposed for dynamic rendering by Debevec *et al* [409] and later expanded to a static texture as in Wang *et al* [410]. However, misalignment in the camera pose, changes in lighting, and specular reflections can cause blurring, highlight artefacts, and discontinuities at image boundaries. As an alternative, mosaicking multiple texture patches each from a single image, as used by Lempitsky *et al* [411], avoids these effects at the expense of resolution by using only the image best aligned to the target location and minimising discontinuities across the seam with its neighbours. Any remaining seam artefacts or specular highlights can then be removed via Poisson image editing [412] as shown by

Chen *et al* [413]. Shadows and other lighting variation between images are then removed via a number of methods, examined by Xu and Mulligan [414], to present a texture with consistent illumination and improved visibility of surface features.

Transfer of the fully textured mesh surface to a texture image for efficient storage and rendering is then conducted via mesh parameterisation. As examined by Hormann *et al* [415], this generates a map between the 3D mesh vertices and the pixel coordinates of a unified texture image, effectively flattening the 3D mesh geometry into a 2D plane. Unlike primitive shapes such as cubes and cylinders, the complex geometry of reconstructed scenes is non-trivial to map to a single 2D texture patch without introducing regions of significant distortion and variable resolution ill-suited to NDE applications. Similarly, texture mapping via individual triangular faces results in excessive seams and accompanying visual artefacts. Consequently, the mesh is divided into irregular patches referred to as “charts” to facilitate texture mapping. As described by Lévy *et al* [416], these are generated to jointly minimise distortion and seam artefacts, ensuring continuous invertible mapping and positioning boundaries along high curvature features. These irregular charts are then packed into the texture image so as to minimise unused space and allow maximum texture resolution for given image dimensions [416], as depicted in Figure A-8.




Figure A-8: Illustration of reconstruction texture mapping. (a) The final textured mesh with magenta lines highlighting individual texture chart boundaries. (b) The resulting 2D texture image.

This completes both the mesh texture generation and the larger 3D reconstruction process depicted in Figure A-6, resulting in a high fidelity representation of the target

scene which may be readily deployed to aid spatial understanding during NDE processes.

As such, the geometry modelling advantages of point cloud reconstructions are further expanded. Transfer of the texture data captured in multiple redundant images of the same point to the mesh further permits flexible RVI where manual observation is obstructed, minimising artefacts due to viewpoint [414] and representing any surface discontinuities or defect features with their full location context. With the option to spatially encode other NDE data, photogrammetric reconstruction then provides a highly effective tool for quantitative inspection and asset management [111], [112].

APPENDIX B

Extended Aspects of Ultrasonic NDE

This appendix unites additional information concerning aspects relevant to practical ultrasonic NDE but beyond the scope of the main text, with a brief assessment of the potential of each topic to aid airborne inspection.

Accordingly, Section B.1 examines behaviours arising during ultrasonic signal transmission from a standard piezoelectric element. Section B.2 examines how separation between piezoelectric elements can influence thickness measurements acquired using a dual-element ultrasonic probe. Section B.3 then introduces how individual ultrasonic readings may collectively form cross-sectional and volumetric images, with Section B.4 discussing algorithms by which this is accomplished. Finally, Section B.5 considers modes of ultrasonic signal generation beyond conventional piezoelectric approaches.

B.1 Beam Spread

For practical application of ultrasound NDE, it is important to note that the ultrasonic wave emitted by a piezoelectric probe is not constrained to a purely linear ray. Instead, a beam is formed with non-zero cross-sectional area, exhibiting 3D variation in acoustic field intensity. As illustrated in Figure B-1, the shape of this beam is defined by a -6 dB, half-amplitude loss from that at the centreline, where the highest particle displacement amplitudes are consistently located.




Figure B-1: Ultrasonic beam spread from a single unfocused circular element piezoelectric transducer [140]. After propagating distance N , the near field cylindrical profile of the beam expands into a cone in the far field zone. Annotations indicate other variation of the ultrasound signal amplitude within the beam. Dimensions are not to scale.

The transmitted acoustic field is thus approximately constrained to the shape of a cylinder in the near field or “Fresnel” region, before expanding into a cone in the far field or “Fraunhofer” region. Transition between the two occurs after the near field distance, N . For an unfocused circular compression wave transducer this may be calculated by the element diameter, D , and ultrasonic wavelength in the propagating material [140].

$$N = \frac{D^2}{4\lambda} \quad (\text{B-1})$$

As shown in Figure B-1, the acoustic field also exhibits variation within its -6 dB boundaries. Notably, the amplitude along the centreline ray exhibits several peaks in the near field, with the location of the last of these maxima defining N . This variation arises owing to interference of the wave fronts transmitted from different locations about the piezoelectric element surface. It is therefore also present in the transverse directions, perpendicular to the centreline, showing significant variation in the near field before transitioning to resemble a bell curve in the far field.

These effects can introduce errors to amplitude-based measurements, such as those taken when sizing discontinuities within a test piece. In common procedures defining defect size by the probe position where the received signal drops by -6 dB from its maximum [417], near field transverse amplitude variation can significantly influence

results, masking the true defect response. For this reason, such measurements must be taken using the more uniform far field, selecting an appropriate probe frequency or adding a delay line such that target features of interest cannot occur in the near field.

Outside of amplitude sensitive applications, however, beam spread can also prove advantageous. It provides a tolerance for non-ideal probe alignment allowing reception of a degraded far-field signal where the centreline ray path would entirely miss the receiving element. This property may be meaningfully exploited in UAV deployment of ultrasonic NDE where variation in probe pose due to the floating platform can be anticipated.

B.2 V-Path Thickness Measurement

Ultrasonic thickness measurement is a common mode of NDE widely deployed to ensure safe operations and target maintenance activity across industrial structures sensitive to corrosive wall loss [100]. As described in Section 2.3.4, dual-element probes can be particularly well suited to this purpose per relevant industrial standards [171], [172]. However, dual element probes present certain additional nuances regarding signal propagation which can influence the conversion of wave-packet time-of-flight to a thickness measurement.

As depicted in Figure B-2, a primary influence on thickness measurement with dual-element probes is the V-shaped path travelled by the ultrasonic wave-packet while moving between the separate piezoelectric elements used for transmission and reception. Versus single-element methods, this V-path enhances sensitivity to echoes from the bottom of internal surface pitting, granting a clearer reading of the minimum remaining thickness. This arises due to a pseudo-focussing effect at the intersection between the beam paths from the transmit and receive elements, as set by their “roof angle” offset from horizontal [326]. Further, the use of distinct transmit and receive elements can reduce sensitivity to near-surface dead zone effects found in thin samples or where rough outer surfaces trap pockets of couplant and introduce persistent ringing to the signal. The net result is then a strong capability for robust measurement of corroded or pitted samples [160], [326], [418].




Figure B-2: Cut-away diagram of Single and Dual-Element probes showing the path of an ultrasonic wave-packet during thickness measurement. (a) Using a single element probe, the wave-packet propagates directly across the sample thickness and is reflected back along its incident path. (b) Using a dual-element probe, the wave-packet travels a longer V-path due to the horizontal separation of the transmit and receive elements. The blue arrows depict the propagation path for the first received back-wall echo. Both probes are typically comprised of the same functional components, as described in Section 2.3.3.

However, the V-path travelled by the ultrasonic wave-packet from a dual-element probe is fractionally longer than the path directly transiting the component thickness used by a single-element probe. For this reason, current industrial standards [171], [172] do not generally recommend the use of Mode 3 inter-echo thickness measurement for thin samples where the horizontal element separation can represent a significant proportion of the total path length and so meaningfully affect the thickness measurement. Specialised near-focus probes or appropriate compensation and signal processing are suggested to aid measurement in such circumstances. Properly conducted Mode 3 methods can then yield valuable asset integrity information in scenarios where Mode 1 or Mode 2 measurements are not practical.

To aid understanding of V-path signal propagation and its effects on measured thickness, a simple, geometry-based, model is analysed. In this model, a wave-packet propagates in an ideal plate of thickness, d , between two infinitesimal points on the top surface, Tx and Rx , with separation distance, w . To improve generality, this makes the simplifying assumption that the overall ultrasonic measurement is most significantly influenced by its path through the sample, and that factors relating to specific probe hardware and pulse propagation therein may be omitted or examined

separately. The multi-skip path travelled by the ultrasonic wave-packet may then be unrolled and represented by a series of triangles, as shown in Figure B-3.




Figure B-3: A simple model for V-path thickness measurement. This shows the path taken by multiple back-wall echoes to travel from the transmitter (Tx) to the receiver (Rx) piezoelectric elements and the equivalent unrolled path. Note that both elements are considered to exist at points on the sample front-wall surface for the purposes of this examination.

From the model of Figure B-3, the path length, L_k , travelled by the k^{th} echo signal received at point Rx , is thus expressible as the hypotenuse of a right-angled triangle with base w and height $2dk$, owing to the pulse-echo configuration.

$$L_k = \sqrt{w^2 + (2dk)^2} \quad (\text{B-2})$$

Equivalently, this may be represented as the length of the path making $2k$ transits directly across the true thickness plus an additional non-negative component, l_k , which varies for each echo.

$$L_k = 2dk + l_k \quad (\text{B-3})$$

Within the triangular path model, l_k is then the non-linear function giving the difference between the hypotenuse and height for the k^{th} echo path.

$$l_k = \sqrt{(2dk)^2 + w^2} - 2dk \quad (\text{B-4})$$

Accordingly, its partial derivative, $\frac{\partial l_k}{\partial k}$, describes the change in additional path length between successive echoes.

$$\frac{\partial l_k}{\partial k} = 2d \left(\frac{2kd}{\sqrt{(2kd)^2 + w^2}} - 1 \right) \quad (\text{B-5})$$

From this, note that the derivative is less than zero for all non-negative k and will asymptotically trend towards zero magnitude as the number of echoes grows larger.

Subject to these V-path effects, the arrival time of the k^{th} echo, t_k , is related to path length by its propagation velocity within the sample material, v .

$$t_k = \frac{L_k}{v} \quad (\text{B-6})$$

Given the additional path component l_k , the V-path arrival time t_k will always lag behind that of an echo making $2k$ transits directly across the sample width. The thickness measured using the Mode 1 time of flight, \widetilde{d}_{1k} , will then consistently over-size the test piece and can be then expressed by substitution of (B-2) into (3-11) as follows. Note that this tends towards the true thickness, d , when $w \rightarrow 0$ or $k \rightarrow \infty$.

$$\widetilde{d}_{1k} = \frac{vt_k}{2k} = \sqrt{\left(\frac{w}{2k}\right)^2 + d^2} \quad (\text{B-7})$$

The nature of this over-sizing becomes readily apparent when \widetilde{d}_{1k} is expressed via (B-3), as below.

$$\widetilde{d}_{1k} = d + \frac{l_k}{2k} \quad (\text{B-8})$$

The same thickness measurement can be performed using a Mode 3 method, provided the signal presents two or more backwall echoes. In this case, the latest inter-echo time of flight given k echoes, T_k , is proportional to the difference between path lengths travelled by successive echoes, ΔL_k , and may be calculated using (B-2) as below.

$$T_k = t_k - t_{k-1} = \frac{\Delta L_k}{v} = \frac{1}{v}(L_k - L_{k-1}) \quad (\text{B-9})$$

$$T_k = \frac{1}{v} \left(\sqrt{w^2 + (2dk)^2} - \sqrt{w^2 + (2d(k-1))^2} \right) \quad (\text{B-10})$$

Noting that the ultrasonic wave-packet transits the thickness of the sample twice between each received echo, the measured thickness using the k^{th} Mode 3 inter-echo time of flight is found using (3-11) and expressed as \widetilde{d}_{3k} . Again this tends towards the true thickness, d , as $w \rightarrow 0$ or $k \rightarrow \infty$.

$$\widetilde{d}_{3k} = \frac{vT_k}{2} = \frac{1}{2} \left(\sqrt{w^2 + (2dk)^2} - \sqrt{w^2 + (2d(k-1))^2} \right) \quad (\text{B-11})$$

In this case, the equivalent expression of \widetilde{d}_{3k} using (B-3) indicates that this Mode 3 thickness measurement includes an error related to the difference between the additional V-path length of successive echoes.

$$\widetilde{d}_{3k} = d + \frac{1}{2}(l_k - l_{k-1}) \quad (\text{B-12})$$

With the consistently negative change in additional V-path length l_k for successive echoes per (B-5), it can be inferred that $l_k < l_{k-1}$. Therefore, the reading from a Mode 3 inter-echo thickness measurement using a dual-element probe will consistently under-size the true dimension, despite the additional length of the V-path.

Measurements conducted using either the Mode 1 or Mode 3 time-of-flight extraction method then present different error characteristics which must be understood so as to inform effective practical NDE. To illustrate generic behaviour, the scenario where $Tx-Rx$ separation is related to the sample thickness by positive scalar factor, a , such that $w = ad$, is examined using the model of Figure B-3. The fractional errors in the

Mode 1 and Mode 3 thickness measurements made using the k^{th} back-wall echo, e_{1k} and e_{3k} respectively, are then represented as below, with a comparative plot given in Figure B-4, below.

$$e_{1k} = \frac{\widetilde{d}_{1k} - d}{d} = \left(\sqrt{\left(\frac{a}{2k}\right)^2 + 1} \right) - 1 \quad (\text{B-13})$$

$$e_{3k} = \frac{\widetilde{d}_{3k} - d}{d} = \left(\sqrt{\left(\frac{a}{2}\right)^2 + k^2} \right) - \left(\sqrt{\left(\frac{a}{2}\right)^2 + (k-1)^2} \right) - 1 \quad (\text{B-14})$$

Figure B-4(a) then illustrates how the Mode 1 thickness measurement will oversize the measured thickness due to the longer length of the V-path. Per (B-13), a +11.8 % error is presented when using the first back-wall echo with a probe giving $Tx-Rx$ separation equal to the component thickness. This falls to +3.1 % for the second echo and asymptotically trends towards zero as more back-wall echoes are available. Mode 3 thickness measurement, however, is valid only after two back-wall echoes are available and will consistently undersize the test piece, despite the longer propagation path length. Mode 3 error is also consistently larger in magnitude than Mode 1 per (B-14), presenting a -5.6 % error versus the true thickness for a signal with two echoes obtained from a the same probe as described above. Nevertheless, this error will likewise tend asymptotically towards zero magnitude when computed using later echoes.

Figure B-4(b) also indicates that both V-path thickness measurement methods are very sensitive to changes in probe element separation relative to the thickness of the part, as may be expected. Each shows an error reduction of the order of two decades for each factor of 10 reduction in the ratio a , independent of the number of echoes used. The significance of V-path effects on thin samples, where the horizontal distance between elements becomes a less significant component of the total ultrasonic path length, is thereby highlighted.

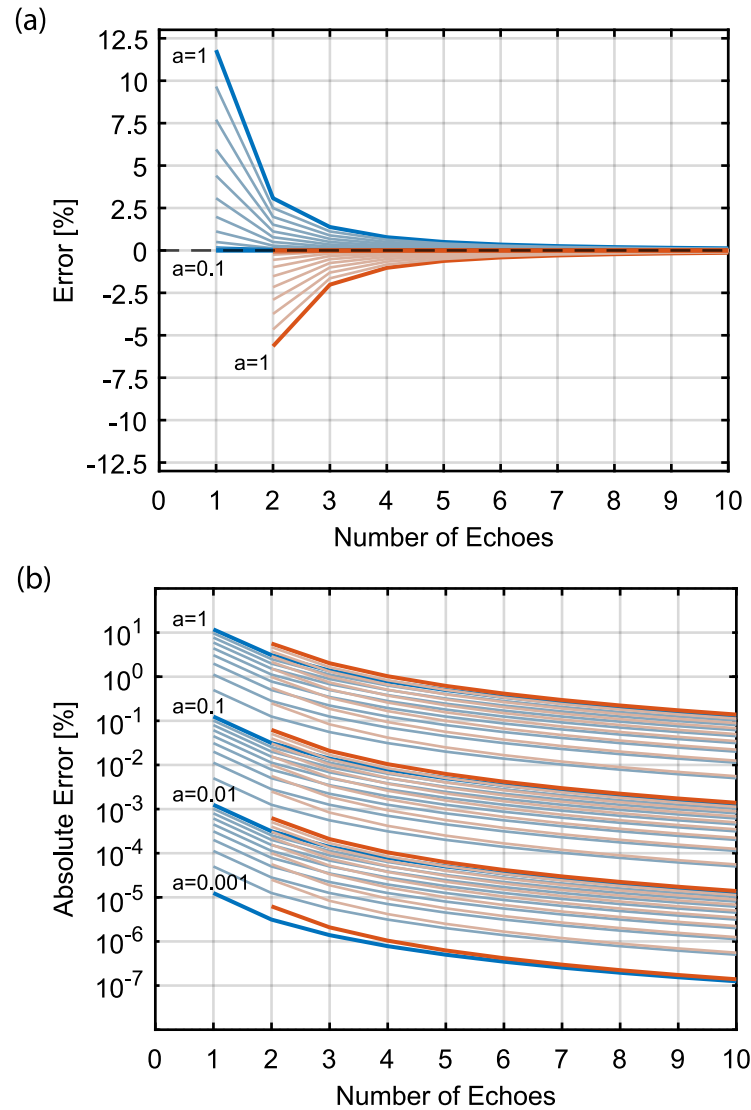


Figure B-4: Thickness measurement error due to V-path effects with varying element separation for a given number of available echoes. (a) Signed error displayed on standard axes. (b) Absolute error magnitude on log axes. Probe element separation to sample thickness ratio, a , is evaluated at log-space intervals. Mode 1 measurements are coloured blue. Mode 3 measurements are coloured orange.

In summary, the simplifying assumptions of the triangular propagation path model allow number of key implications for ultrasonic thickness measurement to be derived. Firstly, the V-path of the ultrasonic wave-packet through the plate sample when using a dual-element probe is always longer than the direct transit path of a single element probe, but this additional length becomes increasingly less significant as the sample gets thicker or measurement uses later back-wall echoes. Accordingly, Mode 1 ultrasonic thickness measurement conducted with a dual-element probe will consistently over-size the test piece. Measurement error magnitude can, however, be

reduced by using later arriving echoes provided they remain clearly distinguishable from any signal noise. By contrast, dual-element Mode 3 thickness measurement will consistently under-size the sample. The Mode 3 error magnitude is marginally larger than the Mode 1 error given a consistent number of back-wall echoes in the received signal, but again can be reduced by using later echoes. Using the inter-echo time difference also permits this method to be employed with an unknown or variable probe-transit time or delay line and through transmissive coatings, which cannot be achieved via Mode 1. As this under-sizing error also lies on the side of safety when measuring remaining wall thickness, the practical utility of Mode 3 inter-echo thickness measurement can be supported when using a dual-element probe in such cases, provided the influence of V-path measurements is well characterised and documented.

B.3 Cross-Sectional & Volumetric Imaging

Whereas single-element thickness measurement and flaw detection may provide beneficial insight to a small region of the test piece along a beam propagation path, further information regarding volumetric defect status can be obtained by combining signals to generate a cross-sectional representation of the structure. As shown in Figure B-5, this may be accomplished by physically moving a single probe, provided its position is accurately encoded, or by electronically scanning using the multiple static elements of a linear ultrasonic array, with the result of either approach termed a “B-scan”.




Figure B-5: Ultrasonically imaging a structure cross-section. This may be performed by (a) physically moving a single transducer, or (b) by electronically scanning through multiple transducers over the length of a physically static phased array probe.

In essence, an A-scan interrogates the material characteristics along a single line through its thickness. A B-scan may then be thought of as a horizontal stacking of A-scan lines to produce a planar cross-sectional signal set. By similar extension, a C-scan is then considered a stacking of B-scans, in the direction normal to their image plane, to produce a full volumetric signal set. Such techniques thereby generate more visually intuitive results than the abstract interpretation of A-scan signals reviewed in Section 2.3.4, allowing direct observation of the defect features in the context of their surrounding material and so better informing health monitoring and remediation activities.

Examples of such imaging techniques are presented in Figure B-6. Note that therein the A-scan has been demodulated to remove the carrier frequency of the piezoelectric element oscillations from the sample echo response. This may be achieved using the Hilbert envelope, exposing response amplitude for waveform peak detection and timing analyses [419]. Additionally, the C-scan image in this case depicts the ToF to the first back-wall reflection, indicating regions of bulk thickness loss across the scanned surface.

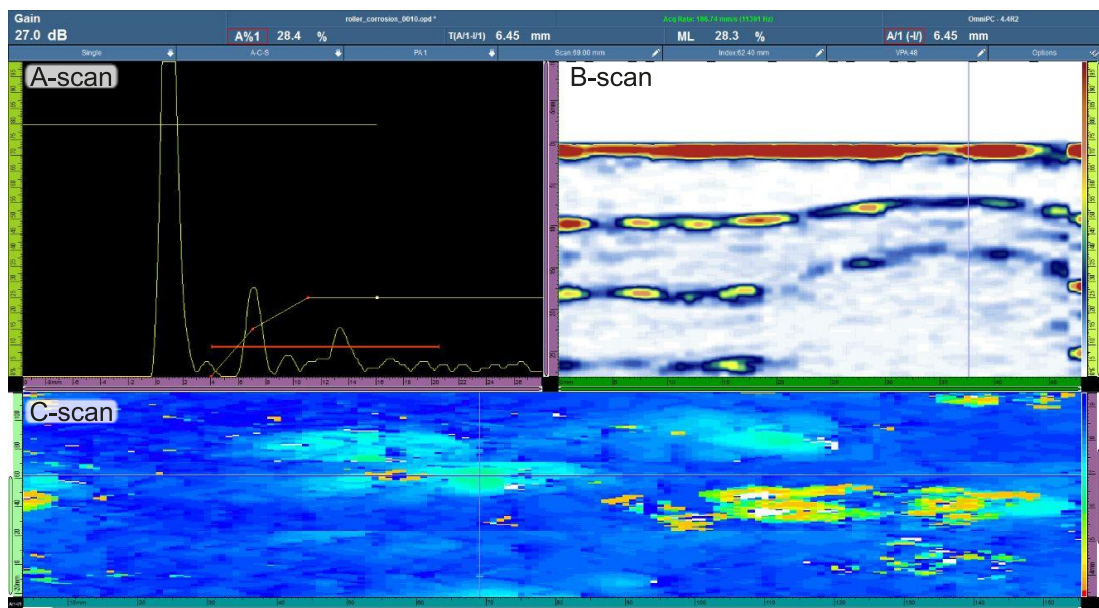


Figure B-6: Example NDE results from ultrasonic array inspection of a steel pipe section with 9 mm wall thickness and 508 mm diameter. From the top, left to right, this shows a demodulated A-scan, cross-sectional B-scan, and a C-scan area thickness image [420].

Figure B-6 also serves to highlight the flexibility of linear array transducers. By offering granular control of the voltage, trigger time offset, and waveform capture of each element it is possible for an array replicate the beam generation of a single transducer by stimulating select sub-groups of its elements [166]. Moreover, this beam can be dynamically refocussed to target specific regions of the sample and redirected through beam-steering to propagate in directions beyond the probe face normal without alterations to the ultrasonic hardware [166]. The B-scan and C-scan images of Figure B-6 can thereby be rapidly captured with minimal probe motion or distortion due to uncertainty in the transmitting element position, informing detailed assessment across large target surface areas. In light of the scale of many industrial structures, it is therefore highly desirable to replicate these outcomes aboard an airborne platform.

B.4 Full Matrix Capture Methods

Expanding on these beam forming techniques, ultrasonic array elements may also be excited individually in sequence, with the echo response to each transmission captured across all elements simultaneously in a process known as Full Matrix Capture (FMC). Post processing of these signals using the Total Focussing Method (TFM) [421] then enables the extraction and combination of signal amplitude corresponding to a given spatial location across all signals. This results in a cross-sectional image where each point appears as if targeted by a focussed transmission from all array elements, removing distortion artefacts due to reflector point spread and improving feature sizing consistency versus conventional beam forming methods. A similar strategy is taken in Phase Coherence Imaging (PCI) [422], in this case discarding amplitude and extracting the instantaneous phase from the FMC A-scan waveforms to produce a globally focussed cross-section image by identifying the consistent phase influence of small reflective and diffractive discontinuities amid the incoherent noise of the background material. This has the benefit of being amplitude independent, reducing sensitivity to material attenuation or instrumentation depth compensation gain settings, and improving detection of small defects near large reflectors, such as in high-temperature hydrogen attack and stress corrosion cracking [423]. As such, by permitting adaptive FMC data processing after the initial signal capture, both TFM and PCI thereby present

excellent applicability to detection and sizing of fine-scale defects in complex, multi-layered, materials [424], [425].

However, the instrumentation required to enable these features entails significant additional complexity relative to a single element transducer. With recent bench-top hardware weighing close to 5 kg [426], potential UAV applications would be restricted to exceptionally large, high-payload, platforms. FMC methods are therefore best revisited for UAV application following miniaturisation of the instrumentation by development activity in the wider field.

B.5 Alternate Methods for Ultrasound Transduction

In addition to the well-established range of piezoelectric transducers, ongoing development efforts have increased support for alternate physical mechanisms for ultrasound generation within NDE applications, offering unique benefits to enable new inspection modes and applications.

Examples such as Electro-Magnetic Acoustic Transducers (EMATs) present the opportunity for multiple frequency and wave mode excitation within consistent hardware [158] and so are well suited to the generation of complex surface and feature guided wave modes supporting large area inspection from a single probe location [151], [427]. These EMATs generate ultrasound directly within the test material using a mixture of the Lorentz force, where a static magnetic field interacts with eddy currents induced in an electrically conductive sample, and magnetostriction, where the volume of a material changes in response to an applied magnetic field [158]. As a result, EMATs also permit ultrasound generation at a small standoff, through non-conductive coatings and minor external surface fouling [308], a distinct improvement where manual preparation activity would be required by conventional piezoelectric transducers. However, certain aspects present obstacles to their UAV deployment such as the limited electromagnetic coupling efficiency, necessitating inclusion of large permanent magnets within already restricted UAV payload [99] or compensatory signal processing algorithms such as averaging or coded excitation [309].

Applicability is also constrained by inability of EMATs to produce ultrasound in non-conductive materials such as composite aerospace components or wind turbine blades.

Another means of ultrasound generation is presented by the use of laser light sources to transmit focussed optical energy directly into the target [157]. In this case, ultrasound is generated by the rapid heating and thermal expansion of a point on the target surface or, at higher energy levels, the reaction force associated with the expulsion of ablated material from the point of laser incidence [428]. Though this second excitation scheme cannot technically be termed non-destructive, an ablative coating may first be applied to the target piece to prevent surface damage. Propagation of the resulting ultrasonic vibrations across the target is then detectable using a conventional transducer or a laser interferometer. As such, ultrasonic inspection may be conducted in the absence of any physical coupling to the surface and with high transducer standoff, aspects advantageous to potential UAV applications. Utility of laser excitation is further shown in its generation of a broadband, high-resolution, pulse excitation with wave mode controllable by the impinging energy levels and adjustable without hardware changes. Coupled with a small laser cross-sectional area allowing coupling to tight radius curved geometry, laser ultrasonics have become well supported in select applications concerning inspection of aerospace composite structures [429] and additively manufactured structures [430].

However, the low sensitivity of interferometric signal capture necessitates signal averaging to eliminate noise, sometimes entailing upwards of 500 iterations [430]. This is problematic for a UAV deployment owing to the changes in the underlying signal with minor pose displacements exhibited during data capture timescales on the order of 100 ms. Furthermore, the comparative expense and impact sensitivity of optical equipment [99] poses a significant operation risk to airborne deployment. Considering these aspects, the generality and efficiency of ultrasonic transduction offered by inexpensive, lightweight, piezoelectric probes again proves advantageous.

APPENDIX C

Comparative Review of UAV Control Theory

This appendix provides a consolidated introduction to the many algorithms used in the control of multirotor UAVs. The basic operation of each method is presented from first principles. In each case, this is accompanied by a brief examination of the early works establishing its use for multirotor control and modern findings representative of its current state. Outcomes regarding the general attributes of each structure and its utility for airborne contact-based inspection are then derived by relative comparison.

As such, Section C.1 introduces the selection of typical control structures to be examined, while Section C.2 subsequently reviews their comparative benefits.

Further information can be obtained in the general UAV control reviews by Özbek *et al* [214], Shraim *et al* [215], Nascimento *et al* [216], and Nguyen *et al* [217].

C.1 Common UAV Control Algorithms

C.1.1 Proportional-Integral-Derivative Control

Having been established theoretically in 1922 [431], the Proportional, Integral, and Derivative (PID) controller remains one of the most widely deployed linear control strategies, with broad applicability to various robot archetypes and industrial machines. Its operation is based on the evaluation of an error function, $e(t) = x_{sp}(t) - x(t)$, expressing the difference between the current value of the parameter under control, x , and its desired value, x_{sp} . Mathematically, a PID control law may then be written as

$$u(t) = K_P e(t) + K_I \int e(t) dt + K_D \frac{d}{dt}(e(t)) \quad (\text{C-1})$$

Where $e(t)$ is the error term, $u(t)$ is the controller output (i.e. the input to the dynamic system), and K_p , K_I , and K_D are the proportional integral and derivative term gains, respectively. Within this structure, an increase in K_p will improve responsiveness and reduce steady state error but may lead to oscillations and instability. The integral component will remove steady state error over time, with larger K_I values speeding responsiveness but increasing oscillatory behaviour. Lastly, K_D gain acts to reduce oscillation, but may induce oscillation of opposing polarity if used excessively. Practical alterations to the computation of (C-1) such as anti-windup and derivative kick reduction, applied to the integral and derivative terms respectively, further aid performance [432]. The design objective is then to select gains such that the output is stable and tracks the reference to within response time and error specifications, either by iterative manual, algorithmic, or fully automated tuning processes [433], [434].

Owing to its simplicity and generality, PID control has been widely used in multirotor applications. Initially, its utilisation was most common for altitude and yaw control as a PD controller omitting the integral term [201], [202], before expanding to full PID regulation of the orientation control loop [180], [183], [200], [435]. Use of PID for translation control followed thereafter [183], [436], employing linearisation of the quadcopter dynamics for small angular deflections from the hover point.

However, these early PID implementations directly regulate Euler angle orientation [180], [195], [200]–[202], [435], with position control setting the desired pitch and roll angles based on goal position relative to the body. This introduces vulnerability to both non-linear effects and singularities during extreme manoeuvres, limiting trajectory tracking performance. More modern PID controllers thus decouple the position and attitude loops. These approaches compute the translation loop in the world frame, then separately identify a UAV pose best aligning thrust to meet trajectory demand signal using 3D geometry, before employing this as a setpoint in the attitude control loop. Attitude error versus the current pose can then be evaluated via rotation matrix, as first proposed by Lee *et al* [183], [184], [437], or using a quaternion representation, as detailed by Brescianini *et al* [350]. This results in near global stability, free from the non-linearities and singularities of earlier methods, allowing a simple PID control scheme to provide strong position regulation performance amid system uncertainty or

disturbances and during omnidirectional flight. PID control is then generally noted for its ease of implementation, low computational complexity, and smooth output function [214]–[216].

C.1.2 Linear Quadratic Regulators

Arising based on work by Kalman [438], a Linear Quadratic Regulator (LQR) is a form of full state feedback controller where a linear function is continuously applied to the state estimate in order to stabilise the closed loop system about the desired operating point. In contrast to manual pole placement methods, the LQR gain matrix is constructed algorithmically, so as to optimise a cost function in terms of both state tracking error and control effort [439].

An LQR is generically applicable to any stabilisable system [439] with linear input, state, and output that may be represented as follows

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \quad (\text{C-2})$$

Where \mathbf{x} is the system state vector with rate of change, $\dot{\mathbf{x}}$, expressed as a linear function of the current state and input vector, \mathbf{u} , via product with system matrices \mathbf{A} and \mathbf{B} . Both \mathbf{x} and \mathbf{u} vary as functions of time. The LQR control feedback is then a linear function of \mathbf{x} , found via static gain matrix \mathbf{K} .

$$\mathbf{u} = -\mathbf{K}\mathbf{x} \quad (\text{C-3})$$

Performance of the control law may then be evaluated by cost function, J .

$$J = \frac{1}{2} \int_0^{\infty} \mathbf{x}^T \mathbf{Q}\mathbf{x} + \mathbf{u}^T \mathbf{R}\mathbf{u} dt \quad (\text{C-4})$$

Where \mathbf{Q} and \mathbf{R} are tracking error and control effort diagonal weight matrices, respectively. Note that this is quadratic in terms of \mathbf{x} and \mathbf{u} , hence the LQR name. Weighting matrices may be selected to satisfy desired performance specification and emphasise either minimal control effort or tracking error in each system dimension.

The LQR design objective is thus to determine, for every initial state, an input \mathbf{u} such that $\mathbf{x} \rightarrow 0$ over time, $t \rightarrow \infty$, minimising (C-4) and satisfying the Hamilton-Jacobi-Bellman (HJB) equation [440]. This is approached by solving the Algebraic Riccati Equation (ARE), below, for matrix \mathbf{P} .

$$\mathbf{PA} + \mathbf{A}^T \mathbf{P} - \mathbf{PBR}^{-1} \mathbf{B}^T \mathbf{P} + \mathbf{Q} = \mathbf{0} \quad (\text{C-5})$$

It can then be shown [440] that the optimal state feedback control law is given by

$$\mathbf{K} = \mathbf{R}^{-1} \mathbf{B}^T \mathbf{P} \quad (\text{C-6})$$

LQR thereby represents a common approach to UAV control [215]. Among the earliest work regarding multirotor UAVs, Bouabdallah *et al* [200] propose an LQR method adapting the controller to repeatedly linearise around the current vehicle state for attitude control, demonstrating successful results in simulation. This was subsequently expanded by Cowling *et al* [441] who successfully simulate path following using LQR at the position control level amid disturbances due to wind and variation in actuator effectiveness. More recently, Foehn and Scaramuzza [442] develop a state-dependent LQR structure unifying position and attitude control, also accounting for time-varying system dynamics and disturbances. They experimentally demonstrate highly aggressive trajectory tracking amid such disturbances without loss of stability, or erroneous behaviour where the trajectory is infeasible, showcasing the high levels of acrobatic performance attainable within the LQR framework using state-based re-linearisation.

However, the direct dependence of LQR control on the system model can introduce vulnerability to parameter uncertainty. The early work of Bouabdallah *et al* [200] highlights this, showing steady-state attitude errors when controlling practical UAV hardware due to minor effectiveness variation between propellers. Foehn and Scaramuzza [442] find similar issues, showing non-zero steady-state position tracking error due to collective thrust uncertainty and lack of an integral component within the LQR controller. This presents an issue to NDE activity where precise sensor positioning is a key aspect of process feasibility.

C.1.3 Sliding Mode Control

As first proposed by Utkin [443], Sliding Mode Control (SMC) represents another popular control structure applied to non-linear systems. Its overarching principal is to drive the current system state-space representation onto a prescribed surface, or “manifold”, where its behaviour is well defined. Thereafter, the system state adheres to this surface, “sliding” along it to progresses towards a desired equilibrium. Design of this sliding manifold then confers the standard control properties of stabilisation, tracking, and regulation, while an attractor function based on Lyapunov stability ensures initial state progression towards the sliding surface and completes the controller. Design methods used to meet these criteria are detailed in the tutorial by DeCarlo *et al* [444].

SMC is applicable to any non-linear system of the general form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} \quad (\text{C-7})$$

Where \mathbf{x} is the system state vector, with change over time, $\dot{\mathbf{x}}$, described by a non-linear function of the current state, $\mathbf{f}(\mathbf{x})$, and the input vector, \mathbf{u} , via a second non-linear function, $\mathbf{g}(\mathbf{x})$. Each element of the SMC control function, u_i , then switches according to the position of the state above or below the sliding surface.

$$u_i(\mathbf{x}) = \begin{cases} u_i^+(\mathbf{x}), & \sigma_i(\mathbf{x}) > 0 \\ u_i^-(\mathbf{x}), & \sigma_i(\mathbf{x}) < 0 \end{cases} \quad (\text{C-8})$$

Where $\sigma_i(\mathbf{x})$ is the corresponding element of the sliding manifold, $\boldsymbol{\sigma}(\mathbf{x})$, itself a vector function, expressed as

$$\boldsymbol{\sigma}(\mathbf{x}) = [\sigma_1(\mathbf{x}), \sigma_2(\mathbf{x}), \dots]^T = 0 \quad (\text{C-9})$$

Within this system model, the sliding mode will exist provided that the state derivative, $\dot{\mathbf{x}}$, points towards the sliding surface, $\boldsymbol{\sigma}(\mathbf{x}) = 0$, at all points within the region of attraction where the control system is expected to operate [444].

In proposing SMC for multirotor platforms, Bouabdallah and Siegwart [445] demonstrate its capability to control the attitude loop of a multirotor UAV. In

comparative simulation versus pure PID and backstepping control schemes, Özbek *et al* [214] find SMC to produce the best compromise between tracking performance and simplicity, easily comprehended in terms of the physical UAV parameters. A number of works have additionally highlighted the strong robustness benefits of SMC due to its insensitivity to disturbances and model uncertainty [215]. Lee *et al* [446] show that SMC is able to retain strong tracking performance in the presence of significant position and orientation sensor noise, where other strategies become unstable. More recently, Mofid and Mobayen [447] have expanded this capability through an adaptive, self-tuning, SMC structure capable of online compensation amid unknown parameter uncertainty bounds.

However, SMC presents a fundamental challenge in the form of high-frequency vibrations, known as “chattering”, exhibited during repeated switching across either side of the sliding manifold [445], [446]. This behaviour is problematic owing to its ability to excite unmodelled harmonic vibrations within the UAV airframe that could result in instability [444]. Energy expended through chattering can also negatively impact UAV operational duration versus smoother control, limiting practical deployment duration. Moreover, such vibrations can induce visual blur or corrupt other onboard data capture, posing a significant obstacle to NDE deployment. Muñoz *et al* [448] compare methods to mitigate such effects within the controller at the expense of additional complexity, but find it cannot be fully eliminated.

C.1.4 Backstepping Control

Proposed circa 1990 by works including that of Kokotovic [449], Backstepping Control (BSC) is based on Lyapunov stability theory [450] and attempts to stabilise a non-linear system of the requisite form by deconstructing it into a number of nested sub-systems, with an irreducible sub-system at their core which may be easily stabilised by other methods. A feedback control law stabilising the full system is then recursively constructed by “back-stepping” the known stabilising function outwards from the core.

Accordingly, backstepping control may be applied to the system of (C-7), provided it is also expressible in strict-feedback terms, as represented below [450].

$$\begin{aligned}
\dot{\mathbf{x}} &= f_0(\mathbf{x}) + g_0(\mathbf{x})z_1 \\
\dot{z}_1 &= f_1(\mathbf{x}_1, z_1) + g_1(\mathbf{x}_1, z_1)z_2 \\
&\dots \\
\dot{z}_{k-1} &= f_{k-1}(\mathbf{x}_1, z_1 \dots z_{k-2}) + g_{k-1}(\mathbf{x}_1, z_1 \dots z_{k-2})z_k \\
\dot{z}_k &= f_k(\mathbf{x}_1, z_1 \dots z_{i-2}, z_{k-1}) + g_i(\mathbf{x}_1, z_1 \dots z_{k-2}, z_{k-1})u
\end{aligned} \tag{C-10}$$

Where \mathbf{x} is the current system state, $f_0 \dots f_k$ and $g_0 \dots g_k$ are functions of \mathbf{x} and “feedback” virtual inputs, $z_1 \dots z_k$, such that $f(0, \dots 0) = 0$, and u is the overall control input to the system. The BSC design objective is then to derive a state feedback function such that \mathbf{x} and all z_i terms asymptotically tend to zero [450].

Application to multirotor UAVs is thereby well supported within the cascading position-attitude control scheme of Figure 2-12 and readily applied using the equations of motion derived from (2-32), as shown among the first papers to examine multirotor control by Altug *et al* [195]. The backstepping approach was thereafter investigated by numerous works for both multirotor attitude and position control [196], [445], [451]. Notably, Bouabdallah and Siegwart [452] propose an “integral backstepping” control structure combining the beneficial aspects of PID control with backstepping to reduce sensitivity to model uncertainty shown in pure backstepping, while also mitigating steady state errors. This is further advanced by Labbadi and Cherkaoui [453], by combination of adaptive backstepping control with a fast terminal sliding mode control to augment path tracking performance and extend robustness to time-varying uncertainties and external disturbances.

Generally, BSC shows good control performance where the system is well known, directly incorporating non-linear effects and yielding fast response to demand signals. However, backstepping also requires full state estimation, can show considerable sensitivity to disturbances and model uncertainty if implemented in the pure sense [452], [453], and can entail large control inputs, impacting operational flight time and requiring imposition of saturation limits to prevent instability [214], [215].

C.1.5 Model Predictive Control

First developed for control chemical processing plants by Richalet *et al* [454], Model Predictive Control (MPC) represents one of the more advanced control strategies used to govern multirotor flight. As detailed in the recent survey by Nguyen *et al* [217], MPC is applicable to both the linear and non-linear systems represented by (C-2) and (C-7), respectively. Like the LQR described in Section C.1.2, MPC functions by solving for the optimal input to the system minimising combined error and actuator effort per a relative weighting supplied in the design stage. However, whereas LQR determines a static, linear, state-feedback, gain matrix, MPC directly evaluates the optimal control signal at each timestep online and in real-time. Accordingly, the MPC forecasts system behaviour out to a finite horizon from the current time for a series of candidate control signals. That which grants the optimal response is applied to the physical system in the current timestep, before the calculation is repeated for the next timestep. This approach permits explicit inclusion of additional non-linear constraints on system performance due to actuation limits, system response dynamics, and other external conditions highly beneficial to control of complex UAV systems [217].

Mathematically, this core function of a MPC system is described as follows in the non-linear case.

$$\operatorname{argmin}_{\mathbf{u}} \int_t^{t+T} \|\mathbf{x} - \mathbf{x}_{sp}\|_{\mathbf{Q}}^2 + \|\mathbf{u} - \mathbf{u}_{sp}\|_{\mathbf{R}}^2 dt + \|\mathbf{x}_T - \mathbf{x}_{T,sp}\|_{\mathbf{S}}^2 \quad (\text{C-11})$$

$$\begin{aligned} \text{s.t.} \quad & \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} \\ & \mathbf{u} \in \mathbb{U} \end{aligned}$$

Where T is the horizon time offset from the current time, \mathbf{x}_{sp} is the desired state setpoint, \mathbf{u}_{sp} is the desired control input, often a zero vector, \mathbb{U} is the set of possible control inputs obeying actuation limits, and \mathbf{x}_T and $\mathbf{x}_{T,sp}$ are the system state at the end of the evaluated horizon time and its desired value, respectively. Matrices \mathbf{Q} , \mathbf{R} and \mathbf{S} provide relative weighting to the state tracking error, control effort, and final state error, respectively. Other symbols in common with (C-7) retain their previous meaning.

Owing to the significant embedded processing power requirements, MPC deployment aboard UAVs was established after many of the other control methods, and in initial research typically applied to the position control loop where slower dynamics permit less frequent control updates [217]. Among other early works, Alexis *et al* [455] successfully applied MPC to both position and orientation using a purpose built UAV with high-performance onboard computing, demonstrating position tracking and aggressive attitude control amid directional wind gusts. Subsequently, Kamel *et al* [456] have highlighted the adaptability of nonlinear model predictive attitude control to changes in system parameters, demonstrating aggressive trajectory tracking and maintaining stable flight in the event of rotor failure in a hexacopter UAV. In comparison of linear and non-linear MPC structures for operation in cluttered environments and amid disturbances, Kamel *et al* [220] also show improved responsiveness, aggressive trajectory tracking and disturbance rejection while reducing compute time by a factor of five using the non-linear MPC to encompass the rotor drag and roll-torque effects described in Section 2.4.1. More recently, Bicego *et al* [457] demonstrate a non-linear MPC strategy suitable for vehicles with arbitrary rotor layout, highlighting the flexibility offered by MPC when developing novel UAV systems. This generality also applies to new applications of existing systems and complex airborne tasks, e.g. control of a team of UAVs to transport large cargo slung beneath them using a distributed MPC structure by Wehbeh *et al* [458].

MPC is therefore identified as a means to provide excellent flight performance in scenarios where the system can be fully and accurately modelled, with the vehicle also supporting sufficient computer hardware to enable practical deployment [217].

C.2 Comparative Evaluation

In evaluating the relative capability of each method for airborne NDE, a number of key performance factors are identified. As a non-linear system, the multirotor platform may be expected to perform better under a non-linear control scheme such as SMC or BSC. This intuition is well supported by Bouabdallah and Siegwart [445], demonstrating the ability of BSC to provide robust attitude control under large perturbations from the stable horizontal pose where linear methods function best.

Likewise, the comparison by Özbek *et al* [214] identifies SMC among the top performers in a compromise between tracking performance and robustness.

However, sensors deployed aboard the UAV for NDE data collection often exhibit strong sensitivity to platform motion. Accordingly, the SMC chattering phenomenon will degrade data capture, severely limiting use within this application. BSC schemes show other non-ideal behaviour, generating large amplitude control signals that can lead to instability via actuator saturation and cause excessive energy usage [214], [452], presenting challenges to practical UAV deployment.

Model based control algorithms also commonly exhibit sensitivity to uncertainty in the system parameters, degrading their performance. Among the most sensitive, MPC requires likely disturbances to be included in the model and estimated in real time to adaptively account for discrepancies [217]. BSC is noted to require PID-like integral features to eliminate steady state error [215], [452]. Further divergence due to linearisation exaggerates the effect in LQR methods, with a comparison between PID and LQR quadcopter control by Bouabdallah *et al* [200] finding superior performance in the PID approach owing to its reduced sensitivity to model uncertainty. This issue continues to impact modern LQR performance, with recent work by Foehn and Scaramuzza [442] also exhibiting steady state error due to uncertainty in the collective thrust parameter. As with BSC, a compensatory integral term is recommended, further highlighting the robustness of PID to these effects.

Another factor influencing practical implementation is the computational complexity of the algorithm. As a real-time optimisation problem, MPC requires the highest compute resources of any method examined, with the more complex models requiring progressively higher processing power. This is noted to grant exceptional performance, directly accounting for minute effects down to the level of propeller motor response [457], but typically MPC cannot be computed by low-level flight control microcontrollers and so entails embedding a dedicated multi-thread processor with an order of magnitude higher performance aboard the UAV [220], [455], [456]. In extreme cases, calculations are offloaded to a ground station PC [457], introducing signal transport latency and necessitating an uninterrupted telemetry link which may be challenging to provide amid industrial structures. MPC is therefore best suited to

applications where high dynamic performance control can provide tangible practical benefit and the additional hardware can be fully accommodated aboard the UAV.

By contrast, PID entails minimal computational complexity, also providing the lowest overall pose tracking error in comparison to BSC and SMC in the study by Özbek *et al* [214]. While performance has been observed to degrade in the presence of external disturbances and large perturbations, the cause is attributed to linearisation of the system around the hover point in early PID formulations [200], [202], neglecting non-linear UAV behaviour and basing control on regulation of Euler angles. This introduces vulnerability to singularities and instability in the presence of large disturbances or perturbations from the setpoint, but is addressed by more recent geometric non-linear PID methods, such as those by Lee *et al* [184], [437]. As a result, PID is now able to provide near global asymptotic stability without the earlier vulnerabilities, allowing performance easily sufficient for most flight activity with minimal onboard compute requirements. In light of this and the reduced sensitivity to model uncertainty versus LQR, BSC and MPC methods, a strong case may be made for utilisation of the simpler PID-based control schemes within airborne NDE processes conducting low speed survey.

References

- [1] United Nations, “Net Zero Coalition: Climate Action.” <https://www.un.org/en/climatechange/net-zero-coalition> (accessed Dec. 22, 2022).
- [2] HM Government (UK), *Net Zero Strategy: Build Back Greener*. HH Associates Ltd, 2021. Accessed: Dec. 22, 2022. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1033990/net-zero-strategy-beis.pdf
- [3] BEIS, “Energy Trends, Electricity data.” Office for National Statistics, Apr. 27, 2023. [Microsoft Excel]. Available: <https://www.gov.uk/government/statistics/electricity-section-5-energy-trends>
- [4] BEIS, “Energy Trends: UK, July to September 2022.” Department for Business Energy & Industrial Strategy UK, Dec. 22, 2022. Accessed: Dec. 22, 2022. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1126161/Energy_Trends_December_2022.pdf
- [5] BEIS, “Electricity Generation Costs 2020.” Department for Business Energy & Industrial Strategy UK, Aug. 2020. Accessed: Dec. 28, 2022. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/911817/electricity-generation-cost-report-2020.pdf
- [6] “Why Britain is a world leader in offshore wind,” *The Economist*, Nov. 24, 2022. Accessed: Dec. 28, 2022. [Online]. Available: <https://www.economist.com/britain/2022/11/24/why-britain-is-a-world-leader-in-offshore-wind>
- [7] A. Glover and M. Gilbert, “Reimagining a Net Zero North Sea: An Integrated Energy Vision for 2050,” Net Zero Technology Centre, Aberdeen, UK, Nov. 2020. Accessed: Sep. 23, 2021. [Online]. Available: https://www.netzerotc.com/media/4523/integrated-energy-vision_nztc.pdf?vgo_ee=GWALVxCPsRO17aPo2%2BI6sp0Ia3UKpsppa4QgLSNmJZI%3D
- [8] G. Rannard, “Wind generated a record amount of electricity in 2022,” *BBC News*, Jan. 06, 2023. Accessed: Jan. 07, 2023. [Online]. Available: <https://www.bbc.com/news/science-environment-64179918>
- [9] HM Government (UK), “British Energy Security Strategy.” Apr. 07, 2022. Accessed: Dec. 28, 2022. [Online]. Available:

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1069973/british-energy-security-strategy-print-ready.pdf

- [10] The American Society for Nondestructive Testing, “Codes and Standards Overview,” 2019. <https://www.asnt.org/MajorSiteSections/Standards.aspx> (accessed Jan. 31, 2020).
- [11] UK Health and Safety Executive, “Inspection/Non-Destructive Testing: Regulatory Requirements.” <https://www.hse.gov.uk/comah/sragtech/techmeasndt.htm#RegulatoryRequirements> (accessed Jan. 16, 2020).
- [12] R. Sahal, J. G. Breslin, and M. I. Ali, “Big data and stream processing platforms for Industry 4.0 requirements mapping for a predictive maintenance use case,” *J. Manuf. Syst.*, vol. 54, pp. 138–151, Jan. 2020, doi: 10.1016/j.jmsy.2019.11.004.
- [13] J. Dalzochio *et al.*, “Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges,” *Comput. Ind.*, vol. 123, p. 103298, Dec. 2020, doi: 10.1016/j.compind.2020.103298.
- [14] The Institution of Engineering and Technology, “Engineering and tech giants join forces to STEM £1.5bn annual skills gap,” Nov. 17, 2021. <https://www.theiet.org/media/press-releases/press-releases-2021/press-releases-2021-october-december/17-november-2021-engineering-and-tech-giants-join-forces-to-stem-15bn-annual-skills-gap/> (accessed May 04, 2023).
- [15] The Institution of Engineering and Technology, “Addressing the STEM skills shortage challenge,” The IET, London, UK, E6D21002, May 2021.
- [16] D. Leaker, “Dataset A01: Summary of labour market statistics.” Office for National Statistics, Apr. 18, 2023. Accessed: May 04, 2023. [Microsoft Excel]. Available: <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/datasets/summaryoflabourmarketstatistics>
- [17] R. Young, K. Newton, T. Dunhill, and C. Huggins, “A Landscape for the Future of NDT in the UK Economy,” *Materials KTN*, Apr. 2014. Accessed: Apr. 19, 2022. [Online]. Available: <https://www.bindt.org/downloads/Materials-KTN-Future-of-NDT-in-UK-economy.pdf>
- [18] M. Benbouzid, T. Berghout, N. Sarma, S. Djurović, Y. Wu, and X. Ma, “Intelligent Condition Monitoring of Wind Power Systems: State of the Art Review,” *Energies*, vol. 14, no. 18, Art. no. 18, Jan. 2021, doi: 10.3390/en14185967.
- [19] D. A. Katsaprakakis, N. Papadakis, and I. Ntintakis, “A Comprehensive Analysis of Wind Turbine Blade Damage,” *Energies*, vol. 14, no. 18, Art. no. 18, Jan. 2021, doi: 10.3390/en14185974.

- [20] O. Adedipe, F. Brennan, and A. Kolios, "Review of corrosion fatigue in offshore structures: Present status and challenges in the offshore wind sector," *Renew. Sustain. Energy Rev.*, vol. 61, pp. 141–154, Aug. 2016, doi: 10.1016/j.rser.2016.02.017.
- [21] Weather Guard Lightning Tech, "Leading Edge Erosion: Why We Keep Failing To Solve It," Jan. 17, 2021. <https://weatherguardwind.com/leading-edge-erosion/> (accessed Dec. 30, 2022).
- [22] B. F. Sørensen *et al.*, "Improved design of large wind turbine blade of fibre composites based on studies of scale effects (Phase 1) - Summary Report," Risø National Laboratory, Roskilde, Denmark, Summary Report Risø-R-1390(EN), Sep. 2004.
- [23] A. Candela Garolera, S. F. Madsen, M. Nissim, J. D. Myers, and J. Holboell, "Lightning Damage to Wind Turbine Blades from Wind Farms in the U.S.," *IEEE Trans. Power Deliv.*, vol. 31, no. 3, pp. 1043–1049, Jun. 2016, doi: 10.1109/TPWRD.2014.2370682.
- [24] N. Gaudern, "A practical study of the aerodynamic impact of wind turbine blade leading edge erosion," *J. Phys. Conf. Ser.*, vol. 524, no. 1, p. 012031, Jun. 2014, doi: 10.1088/1742-6596/524/1/012031.
- [25] C. B. Hasager, F. Vejen, W. R. Skrzypiński, and A.-M. Tilg, "Rain Erosion Load and Its Effect on Leading-Edge Lifetime and Potential of Erosion-Safe Mode at Wind Turbines in the North Sea and Baltic Sea," *Energies*, vol. 14, no. 7, Art. no. 7, Jan. 2021, doi: 10.3390/en14071959.
- [26] K. B. Katnam, A. J. Comer, D. Roy, L. F. M. da Silva, and T. M. Young, "Composite Repair in Wind Turbine Blades: An Overview," *J. Adhes.*, vol. 91, no. 1–2, pp. 113–139, Jan. 2015, doi: 10.1080/00218464.2014.900449.
- [27] J. C. Marín, A. Barroso, F. París, and J. Cañas, "Study of fatigue damage in wind turbine blades," *Eng. Fail. Anal.*, vol. 16, no. 2, pp. 656–668, Mar. 2009, doi: 10.1016/j.engfailanal.2008.02.005.
- [28] UK Health and Safety Executive (HSE), "Work-related fatal injuries in Great Britain," Nov. 2022. <https://www.hse.gov.uk/statistics/fatals.htm> (accessed Jan. 06, 2023).
- [29] Industrial Rope Access Trade Association (IRATA), "International code of practice for industrial rope access - Part 3: Informative annexes - Annex O: Protecting rope access technicians against environmental conditions." IRATA, Sep. 2016. Accessed: Jan. 06, 2023. [Online]. Available: https://www.5thpoint.com.au/wp-content/uploads/2017/11/part-3-annex-o-2016-sep-01_1.pdf
- [30] L. Mishnaevsky and K. Thomsen, "Costs of repair of wind turbine blades: Influence of technology aspects," *Wind Energy*, vol. 23, no. 12, pp. 2247–2255, Dec. 2020, doi: 10.1002/we.2552.

- [31] B. Bedeschi, “Offshore wind operators use scale, analytics to cut vessel trips,” Mar. 07, 2019. <https://www.reutersevents.com/renewables/wind-energy-update/offshore-wind-operators-use-scale-analytics-cut-vessel-trips> (accessed Dec. 13, 2022).
- [32] I. B. Sperstad, E. E. Halvorsen-Weare, M. Hofmann, L. M. Nonås, M. Stålhane, and M. Wu, “A Comparison of Single- and Multi-parameter Wave Criteria for Accessing Wind Turbines in Strategic Maintenance and Logistics Models for Offshore Wind Farms,” *Energy Procedia*, vol. 53, pp. 221–230, Jan. 2014, doi: 10.1016/j.egypro.2014.07.231.
- [33] D. Foxwell, “Day rates for SOVs likely to continue upward course,” *Riviera*, May 28, 2019. <https://www.rivieramm.com/news-content-hub/news-content-hub/day-rates-for-sovs-likely-to-continue-upward-course-54572> (accessed Jan. 07, 2023).
- [34] P. H. Lauritsen *et al.*, “Strategy for Extending the Useful Lifetime of a Wind Turbine,” MegaVind, DK-1970 Frederiksberg C, Denmark, Aug. 2016. Accessed: May 18, 2020. [Online]. Available: <http://www.vindmoellegodkendelse.dk/media/1163/strategy-for-extending-the-useful-lifetime-of-a-wind-turbine.pdf>
- [35] A. Spyroudi, “End-of-Life Planning in Offshore Wind.” Offshore Renewable Energy Catapult (OREC), Apr. 2021. Accessed: Dec. 28, 2022. [Online]. Available: https://ore.catapult.org.uk/wp-content/uploads/2021/04/End-of-Life-decision-planning-in-offshore-wind_FINAL_AS-1.pdf
- [36] BEIS, “Renewable Energy Planning Database: quarterly extract.” Department for Business Energy & Industrial Strategy UK, Oct. 28, 2022. Accessed: Dec. 30, 2022. [Online]. Available: <https://www.gov.uk/government/publications/renewable-energy-planning-database-monthly-extract>
- [37] World Nuclear Association, “Nuclear Power in the United Kingdom,” Jul. 2022. <https://world-nuclear.org/information-library/country-profiles/countries-t-z/united-kingdom.aspx> (accessed Jan. 07, 2023).
- [38] P. Trampus, “Role and importance of NDE in nuclear power plant life extension,” *Procedia Struct. Integr.*, vol. 16, pp. 161–168, Jan. 2019, doi: 10.1016/j.prostr.2019.07.036.
- [39] International Atomic Energy Agency (IAEA), “Non-destructive testing for plant life assessment.” IAEA Austria, Aug. 2005. Accessed: Jan. 09, 2023. [Online]. Available: https://www-pub.iaea.org/mtcd/publications/pdf/tcs-26_web.pdf
- [40] UK Nuclear Decommissioning Authority (NDA), “Fact sheet: decommissioning of nuclear power facilities.” Aug. 2015. Accessed: Jan. 09, 2023. [Online]. Available: <https://ukinventory.nda.gov.uk/wp-content/uploads/2014/01/Fact-sheet-decommissioning-of-nuclear-power-facilities.pdf>

- [41] E10 Committee, "E1281 Standard Guide for Nuclear Facility Decommissioning Plans," ASTM International, Washington DC, USA, Standard E1281, Feb. 2021. doi: 10.1520/E1281-15R21.
- [42] UK Office for Nuclear Regulation (ONR), "Nuclear Decommissioning," Aug. 09, 2021. <https://www.onr.org.uk/decomissioning.htm> (accessed Jan. 09, 2023).
- [43] Curtiss-Wright, "Aerial Inspection with Drone Technology and Visual Asset Management." Aug. 30, 2021. Accessed: Jan. 04, 2023. [Online]. Available: https://s21.q4cdn.com/920789611/files/download_library/scientech/Aerial-Inspection-with-Drone-Technology-Case-Study.pdf
- [44] "ASME BPVC Section XI - Rules for Inservice Inspection of Nuclear Power Plant Components," American Society of Mechanical Engineers (ASME), New York, NY, USA, Standard BPVC-XI-1, 2021. Accessed: Jan. 04, 2023. [Online]. Available: <https://www.asme.org/codes-standards/find-codes-standards/bpvc-xi-bpvc-section-xi-rules-inservice-inspection-nuclear-power-plant-components/2021/print-book>
- [45] A. Smith, "Guest takeover by Sellafield Ltd. Robotics and Artificial Intelligence Integrated Research Team (RAI-IRT) - UAV Inspections," presented at the RAIN Webinar Series, Jul. 09, 2020. Accessed: Jan. 04, 2023. [Online]. Available: <https://rainhub.org.uk/events/guest-takeover-by-sellafield-ltd-robotics-and-artificial-intelligence-integrated-research-team-rai-irt/>
- [46] Flyability, "Elios 2 tested at the Chernobyl Nuclear Power Plant," Nov. 2020. <https://www.flyability.com/news/chernobyl-mission> (accessed Nov. 20, 2020).
- [47] UK Statutory Instruments, "The Ionising Radiations Regulations 2017 - Schedule 3: Dose Limits." 2017. Accessed: Jan. 10, 2023. [Online]. Available: <https://www.legislation.gov.uk/ukxi/2017/1075/schedule/3/made/data.pdf>
- [48] UK Health and Safety Executive (HSE), "Risk management: ALARP at a glance." <https://www.hse.gov.uk/managing/theory/alarplance.htm> (accessed Jan. 10, 2023).
- [49] I. Tsitsimpelis, C. J. Taylor, B. Lennox, and M. J. Joyce, "A review of ground-based robotic systems for the characterization of nuclear environments," *Prog. Nucl. Energy*, vol. 111, pp. 109–124, Mar. 2019, doi: 10.1016/j.pnucene.2018.10.023.
- [50] H. M. Crockett and J. S. Horowitz, "Erosion in Nuclear Piping Systems," *J. Press. Vessel Technol.*, vol. 132, no. 2, p. 024501, Apr. 2010, doi: 10.1115/1.4000509.
- [51] V. Kain, "Flow Accelerated Corrosion: Forms, Mechanisms and Case Studies," *Procedia Eng.*, vol. 86, pp. 576–588, Jan. 2014, doi: 10.1016/j.proeng.2014.11.083.

- [52] Pumps & Systems, “Flow-Accelerated Corrosion Damages Plant Systems,” *Pumps and Systems Magazine*, Apr. 11, 2014. <https://www.pumpsandsystems.com/flow-accelerated-corrosion-damages-plant-systems> (accessed Jan. 07, 2023).
- [53] Intertek, “Flow Accelerated Corrosion in Nuclear Power Plants,” Apr. 28, 2020. <https://www.intertek.com/blog/2020-04-28-fac/> (accessed Jan. 11, 2023).
- [54] P. C. Wu, “Erosion/Corrosion-Induced Pipe Wall Thinning in U.S. Nuclear Power Plants,” US Nuclear Regulatory Commission, Washington DC, USA, NUREG 1344, TI89 011833, Apr. 1989. doi: 10.2172/6152848.
- [55] E. Bismut, M. D. Pandey, and D. Straub, “Reliability-based inspection and maintenance planning of a nuclear feeder piping system,” *Reliab. Eng. Syst. Saf.*, vol. 224, p. 108521, Aug. 2022, doi: 10.1016/j.res.2022.108521.
- [56] “ASME BPVC Section V - Nondestructive Examination,” American Society of Mechanical Engineers (ASME), New York, NY, USA, Standard ASME BPVC-V, 2017.
- [57] Trading Economics, “United Kingdom Electricity Price - Historical Data,” Jan. 11, 2023. <https://tradingeconomics.com/united-kingdom/electricity-price> (accessed Jan. 11, 2023).
- [58] Offshore Energies UK (OEUK), “Economic Report 2021,” London, UK, Aug. 2021. Accessed: Jan. 12, 2023. [Online]. Available: <https://oeuk.org.uk/wp-content/uploads/2021/08/OGUK-Economic-Report-2021.pdf>
- [59] E. Bowman *et al.*, “International Measures of Prevention, Application, and Economics of Corrosion Technologies Study.” NACE International, Mar. 01, 2016. Accessed: Apr. 07, 2022. [Online]. Available: <http://impact.nace.org/documents/Nace-International-Report.pdf>
- [60] Gecko Robotics, “The Cost of Unplanned Downtime for Refineries,” Jul. 28, 2021. <https://blog.geckorobotics.com/the-cost-of-unplanned-downtime-for-refineries> (accessed Jan. 12, 2023).
- [61] American Petroleum Institute (API), “API 571 - Damage Mechanisms Affecting Fixed Equipment in the Refining Industry,” American Petroleum Institute, Washington DC, USA, Standard API 571, Mar. 2020. Accessed: Jun. 20, 2022. [Online]. Available: https://www.techstreet.com/api/standards/api-rp-571?product_id=2108271
- [62] C. Nwaoha, “5 Key Questions About Corrosion in Carbon Capture Processing Plants,” *Corrosionpedia*, Apr. 04, 2019. <https://www.corrosionpedia.com/corrosion-in-carbon-capture-plants/2/6623> (accessed Jan. 08, 2023).

- [63] Kaefer, “Corrosion under Insulation,” 2022. <https://kaeferltd.co.uk/Surface-Protection/Corrosion-under-Insulation.html> (accessed Jan. 17, 2023).
- [64] Q. Cao *et al.*, “A Review of Corrosion under Insulation: A Critical Issue in the Oil and Gas Industry,” *Metals*, vol. 12, no. 4, Art. no. 4, Apr. 2022, doi: 10.3390/met12040561.
- [65] “Solving the corrosion under insulation challenge,” *NetZeroTC*. <https://www.netzerotc.com/news-events/newsroom/news/2018/solving-the-corrosion-under-insulation-challenge/> (accessed Oct. 25, 2021).
- [66] American Petroleum Institute (API), “API 570 - Piping Inspection Code: In-service Inspection, Rating, Repair, and Alteration of Piping Systems,” American Petroleum Institute, Washington DC, USA, Standard API 570, Feb. 2016. Accessed: Jun. 20, 2022. [Online]. Available: https://www.techstreet.com/api/standards/api-570?product_id=1910713
- [67] M. Z. Ab Rashid, M. F. Mohd Yakub, S. A. Zaki bin Shaikh Salim, N. Mamat, S. M. Syed Mohd Putra, and S. A. Roslan, “Modeling of the in-pipe inspection robot: A comprehensive review,” *Ocean Eng.*, vol. 203, p. 107206, May 2020, doi: 10.1016/j.oceaneng.2020.107206.
- [68] A. Shukla and H. Karki, “Application of robotics in onshore oil and gas industry—A review Part I,” *Robot. Auton. Syst.*, vol. 75, pp. 490–507, Jan. 2016, doi: 10.1016/j.robot.2015.09.012.
- [69] “ASME BPVC Section VIII - Rules for Construction of Pressure Vessels,” American Society of Mechanical Engineers (ASME), New York, NY, USA, Standard ASME BPVC-VIII, 2017.
- [70] American Petroleum Institute (API), “API 653 - Tank Inspection, Repair, Alteration, and Reconstruction,” American Petroleum Institute, Washington DC, USA, Standard API 653, 2003.
- [71] L. Hanna, “Digital Twin Offers Permanent Record of Inspection,” *Eddyfi*, Jul. 21, 2020. <https://blog.eddyfi.com/en/digital-twin-offers-permanent-record-of-inspection> (accessed Jan. 08, 2023).
- [72] Bernardo Ecenarro, “Types of Corrosion and Corrosiveness Categories,” *BESA*, 2022. <https://www.bernardoecenarro.com/en/besa-lab/types-of-corrosion-and-corrosiveness-categories/> (accessed Jan. 08, 2023).
- [73] Society of Petroleum Engineers (SPE) International, “Oil Storage,” *PetroWiki*, Nov. 23, 2021. https://petrowiki.spe.org/Oil_storage (accessed Jan. 15, 2023).
- [74] UK Health and Safety Executive (HSE), “ATEX and explosive atmospheres - Fire and explosion,” 2022. <https://www.hse.gov.uk/fireandexplosion/atex.htm> (accessed Jan. 16, 2023).

- [75] UK Health and Safety Executive (HSE), *Dangerous Substances and Explosive Atmospheres Regulations (DSEAR) 2002*, 2nd ed. Norwich, UK: HSE Books, 2013.
- [76] R. E. Melchers, “The effect of corrosion on the structural reliability of steel offshore structures,” *Corros. Sci.*, vol. 47, no. 10, pp. 2391–2410, Oct. 2005, doi: 10.1016/j.corsci.2005.04.004.
- [77] G. P. Drumond, I. P. Pasqualino, B. C. Pinheiro, and S. F. Estefen, “Pipelines, risers and umbilicals failures: A literature review,” *Ocean Eng.*, vol. 148, pp. 412–425, Jan. 2018, doi: 10.1016/j.oceaneng.2017.11.035.
- [78] T. Klein, “Splash Zone Painting Campaign,” *Sub C*, Feb. 25, 2020. <https://www.subcpartner.com/news/splash-zone-painting-campaign/> (accessed Jan. 08, 2023).
- [79] N. Meyendorf *et al.*, “NDE 4.0—NDE for the 21st Century—The Internet of Things and Cyber Physical Systems will Revolutionize NDE,” in *Center for Nondestructive Evaluation Conference Papers, Posters and Presentations*, Singapore, Jan. 2017. [Online]. Available: https://lib.dr.iastate.edu/cnde_conf/117
- [80] UK Health and Safety Executive, “Human factors/ergonomics, health and safety in the workplace.” <http://www.hse.gov.uk/humanfactors/index.htm> (accessed Oct. 31, 2019).
- [81] P. Nooralishahi *et al.*, “Drone-Based Non-Destructive Inspection of Industrial Sites: A Review and Case Studies,” *Drones*, vol. 5, no. 4, Art. no. 4, Dec. 2021, doi: 10.3390/drones5040106.
- [82] L. Wang, Z. Zhang, and X. Luo, “A Two-Stage Data-Driven Approach for Image-Based Wind Turbine Blade Crack Inspections,” *IEEEASME Trans. Mechatron.*, vol. 24, no. 3, pp. 1271–1281, Jun. 2019, doi: 10.1109/TMECH.2019.2908233.
- [83] P. Buchan, “How Drones Are Being Used for Inspection In Industrial Environments,” presented at the BINDT Multicopters Workshop 2018, Manchester, UK, Jul. 2018. Accessed: May 13, 2021. [Online]. Available: <https://www.bindt.org/admin/Downloads/Multicopters%20for%20Inspection%20-%2005%20Buchan%20-%20July%2018%20V1.1.pdf>
- [84] X. Li, Q. Yang, Z. Chen, X. Luo, and W. Yan, “Visible defects detection based on UAV-based inspection in large-scale photovoltaic systems,” *IET Renew. Power Gener.*, vol. 11, no. 10, pp. 1234–1244, 2017, doi: 10.1049/iet-rpg.2017.0001.
- [85] O. Menéndez, M. Pérez, and F. Auat Cheein, “Visual-Based Positioning of Aerial Maintenance Platforms on Overhead Transmission Lines,” *Appl. Sci.*, vol. 9, no. 1, p. 165, Jan. 2019, doi: 10.3390/app9010165.

- [86] J. Seo, L. Duque, and J. Wacker, “Drone-enabled bridge inspection methodology and application,” *Autom. Constr.*, vol. 94, pp. 112–126, Oct. 2018, doi: 10.1016/j.autcon.2018.06.006.
- [87] P. Rossi, F. Mancini, M. Dubbini, F. Mazzone, and A. Capra, “Combining nadir and oblique UAV imagery to reconstruct quarry topography: methodology and feasibility analysis,” *Eur. J. Remote Sens.*, vol. 50, no. 1, pp. 211–221, Jan. 2017, doi: 10.1080/22797254.2017.1313097.
- [88] J. Deign, “Fully automated drones could double wind turbine inspection rates | Reuters Events | Renewables,” Sep. 27, 2016. <https://www.reutersevents.com/renewables/wind-energy-update/fully-automated-drones-could-double-wind-turbine-inspection-rates> (accessed Aug. 16, 2022).
- [89] C. Stout and D. Thompson, “UAV Approaches to Wind Turbine Inspection.” *Offshore Renewable Energy Catapult (OREC)*, Mar. 2019. Accessed: Jan. 04, 2023. [Online]. Available: <https://www.arpas.uk/wp-content/uploads/2019/04/Cyberhawks-Approach-to-UAV-Inspection-Craig-Stout-ORE-Catapult.pdf>
- [90] Nuclear Engineering International, “Eye in the sky,” May 02, 2018. <https://www.neimagazine.com/features/featureeye-in-the-sky-6136284/> (accessed Jan. 04, 2023).
- [91] Cyberhawk, “World First ASME Code Inspection at a US Nuclear Plant,” *Cyberhawk*, 2018. https://thecyberhawk.com/case_studies/world-first-asme-code-inspection-at-a-us-nuclear-plant/ (accessed Jan. 04, 2023).
- [92] Flyability, “Elios: Inspection of a Nuclear Power Plant,” Jun. 2017. <https://blog.flyability.com/casestudies/inspection-of-a-nuclear-power-plant> (accessed Jan. 04, 2023).
- [93] TEXO DSI, “Pioneering Unmanned Inspection of North Sea FPSO Cargo Tank,” *TEXO*. <https://texo.co.uk/case-studies/pioneering-unmanned-inspection-of-north-sea-fps-o-cargo-tank/> (accessed Jan. 05, 2023).
- [94] Cyberhawk, “North Sea Underdeck Inspection for Oil and Gas Major,” *Cyberhawk*, Apr. 2014. http://thecyberhawk.com/case_studies/north-sea-underdeck-inspection-for-oil-and-gas-major/ (accessed Jan. 05, 2023).
- [95] N. Ida and N. Meyendorf, *Handbook of modern non-destructive testing*, 1st ed. Switzerland: Springer Nature, 2017. [Online]. Available: <https://doi.org/10.1007/978-3-319-26553-7>
- [96] H. Shakhathreh *et al.*, “Unmanned Aerial Vehicles (UAVs): A Survey on Civil Applications and Key Research Challenges,” *IEEE Access*, vol. 7, pp. 48572–48634, 2019, doi: 10.1109/ACCESS.2019.2909530.

- [97] M. Abbas and M. Shafiee, "An overview of maintenance management strategies for corroded steel structures in extreme marine environments," *Mar. Struct.*, vol. 71, p. 102718, May 2020, doi: 10.1016/j.marstruc.2020.102718.
- [98] J. Fausing Olesen and H. R. Shaker, "Predictive Maintenance for Pump Systems and Thermal Power Plants: State-of-the-Art Review, Trends and Challenges," *Sensors*, vol. 20, no. 8, Art. no. 8, Jan. 2020, doi: 10.3390/s20082425.
- [99] C. Hellier, *Handbook of nondestructive evaluation*. New York: McGraw-Hill, 2003. Accessed: Nov. 05, 2020. [Online]. Available: <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=90795>
- [100] BINDT, "An Introduction to NDT Common Methods," British Institute of Non-Destructive Testing, Midsummer House, Riverside Way, Bedford Road, Northampton NN1 5NX, UK, Guidance Document AA050, Nov. 2015. Accessed: Aug. 16, 2022. [Online]. Available: <https://www.bindt.org/admin/Downloads/Apprenticeship-Guidance-Document.pdf>
- [101] S. Gholizadeh, "A review of non-destructive testing methods of composite materials," *Procedia Struct. Integr.*, vol. 1, no. Supplement C, pp. 50–57, Jan. 2016, doi: 10.1016/j.prostr.2016.02.008.
- [102] C. F. Britton, "4.36 Corrosion Monitoring and Inspection," in *Shreir's Corrosion*, Bob Cottis, Michael Graham, Robert Lindsay, Stuart Lyon, Tony Richardson, David Scantlebury, and Howard Stott, Eds., Elsevier, 2010, pp. 3117–3166. Accessed: Dec. 08, 2020. [Online]. Available: <https://doi.org/10.1016/B978-044452787-5.00130-X>
- [103] Pexraytech, "PXR NDT systems | Product Catalogue." Feb. 02, 2023. Accessed: May 29, 2023. [Online]. Available: https://pexraytech.com/wp-content/uploads/2023/02/Pexraytech-Product-Catalogue-PXR-NDT-Systems_2023-01-11.pdf
- [104] Teledyne ICM, "Portable X-Ray Solutions | Product Catalogue." Nov. 02, 2021. Accessed: May 30, 2023. [Online]. Available: https://www.teledyneicm.com/wp-content/uploads/2021/09/DOC_NDT_Digital_Radiography.pdf
- [105] A. N. AbdAlla, M. A. Faraj, F. Samsuri, D. Rifai, K. Ali, and Y. Al-Douri, "Challenges in improving the performance of eddy current testing: Review," *Meas. Control*, vol. 52, no. 1–2, pp. 46–64, Jan. 2019, doi: 10.1177/0020294018801382.
- [106] J. Janesick and G. Putnam, "Developments and Applications of High-Performance CCD and CMOS Imaging Arrays," *Annu. Rev. Nucl. Part. Sci.*, vol. 53, no. 1, pp. 263–300, 2003, doi: 10.1146/annurev.nucl.53.041002.110431.

- [107] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge: Cambridge University Press, 2004. doi: 10.1017/CBO9780511811685.
- [108] P. Corke, *Robotics, vision and control*, 2nd ed. New York, NY: Springer Berlin Heidelberg, 2017.
- [109] K. Whitehead *et al.*, “Remote sensing of the environment with small unmanned aircraft systems (UASs), part 2: scientific and commercial applications,” *J. Unmanned Veh. Syst.*, vol. 02, no. 03, pp. 86–102, Sep. 2014, doi: 10.1139/juvs-2014-0007.
- [110] M. S. Nielsen, I. Nikolov, E. K. Kruse, J. Garnæs, and C. B. Madsen, “High-Resolution Structure-from-Motion for Quantitative Measurement of Leading-Edge Roughness,” *Energies*, vol. 13, no. 15, Art. no. 15, Jan. 2020, doi: 10.3390/en13153916.
- [111] C. Biscarini, I. Catapano, N. Cavalagli, G. Ludeno, F. A. Pepe, and F. Ubertini, “UAV photogrammetry, infrared thermography and GPR for enhancing structural and material degradation evaluation of the Roman masonry bridge of Ponte Lucano in Italy,” *NDT E Int.*, vol. 115, p. 102287, Oct. 2020, doi: 10.1016/j.ndteint.2020.102287.
- [112] U. Bacher, “Hybrid Aerial Sensor Data as Basis for a Geospatial Digital Twin,” in *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Copernicus GmbH, Jun. 2022, pp. 653–659. doi: 10.5194/isprs-archives-XLIII-B4-2022-653-2022.
- [113] S. Hosseinzadeh *et al.*, “A Novel Centralization Method for Pipe Image Stitching,” *IEEE Sens. J.*, vol. 21, no. 10, pp. 11889–11898, May 2021, doi: 10.1109/JSEN.2020.3031637.
- [114] D. Zhang, W. Jackson, G. Dobie, G. West, and C. MacLeod, “Structure-from-motion based image unwrapping and stitching for small bore pipe inspections,” *Comput. Ind.*, vol. 139, p. 103664, Aug. 2022, doi: 10.1016/j.compind.2022.103664.
- [115] M. W. Smith, J. L. Carrivick, and D. J. Quincey, “Structure from motion photogrammetry in physical geography,” *Prog. Phys. Geogr. Earth Environ.*, vol. 40, no. 2, pp. 247–275, Apr. 2016, doi: 10.1177/0309133315615805.
- [116] J. Huyan, W. Li, S. Tighe, R. Deng, and S. Yan, “Illumination Compensation Model with k-Means Algorithm for Detection of Pavement Surface Cracks with Shadow,” *J. Comput. Civ. Eng.*, vol. 34, no. 1, p. 04019049, Jan. 2020, doi: 10.1061/(ASCE)CP.1943-5487.0000869.
- [117] M. Johnson-Roberson, O. Pizarro, S. B. Williams, and I. Mahon, “Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys,” *J. Field Robot.*, vol. 27, no. 1, pp. 21–51, 2010, doi: 10.1002/rob.20324.

- [118] S. Giancola, M. Valenti, and R. Sala, *A Survey on 3D Cameras: Metrological Comparison of Time-of-Flight, Structured-Light and Active Stereoscopy Technologies*. in SpringerBriefs in Computer Science. Cham: Springer International Publishing, 2018. doi: 10.1007/978-3-319-91761-0.
- [119] “ZED 2i - Industrial AI Stereo Camera.” <https://www.stereolabs.com/zed-2i/> (accessed Nov. 16, 2022).
- [120] “Depth Camera D455 – Intel® RealSense™ Depth and Tracking Cameras.” <https://www.intelrealsense.com/depth-camera-d455/> (accessed Nov. 16, 2022).
- [121] M. von Übel, “The Best Photogrammetry Software of 2022,” *All3DP*, Sep. 09, 2022. <https://all3dp.com/1/best-photogrammetry-software/> (accessed Nov. 08, 2022).
- [122] J. L. Schönberger, E. Zheng, J.-M. Frahm, and M. Pollefeys, “Pixelwise View Selection for Unstructured Multi-View Stereo,” in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., in Lecture Notes in Computer Science, vol. 9907. Cham: Springer International Publishing, 2016, pp. 501–518. doi: 10.1007/978-3-319-46487-9_31.
- [123] C. Griwodz *et al.*, “AliceVision Meshroom: An open-source 3D reconstruction pipeline,” in *Proceedings of the 12th ACM Multimedia Systems Conference*, in MMSys ’21. New York, NY, USA: Association for Computing Machinery, Sep. 2021, pp. 241–247. doi: 10.1145/3458305.3478443.
- [124] P. Moulon, P. Monasse, R. Perrot, and R. Marlet, “OpenMVG: Open Multiple View Geometry,” in *Reproducible Research in Pattern Recognition*, B. Kerautret, M. Colom, and P. Monasse, Eds., in Lecture Notes in Computer Science. Cham: Springer International Publishing, 2017, pp. 60–74. doi: 10.1007/978-3-319-56414-2_5.
- [125] “Agisoft Metashape.” <https://www.agisoft.com/> (accessed Oct. 27, 2021).
- [126] “RealityCapture,” *CapturingReality*, Oct. 27, 2021. <https://www.capturingreality.com/> (accessed Oct. 27, 2021).
- [127] “Generate 2D and 3D information, purely from images with Pix4D,” *Pix4D*. <https://pix4d.com/> (accessed Sep. 03, 2018).
- [128] K. Kingsland, “Comparative analysis of digital photogrammetry software for cultural heritage,” *Digit. Appl. Archaeol. Cult. Herit.*, vol. 18, p. e00157, Sep. 2020, doi: 10.1016/j.daach.2020.e00157.

- [129] E.-K. Stathopoulou, M. Welponer, and F. Remondino, “Open-Source Image-Based 3D Reconstruction Pipelines: Review, Comparison and Evaluation,” in *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Strasbourg, France, Nov. 2019, pp. 331–338. doi: 10.5194/isprs-archives-XLII-2-W17-331-2019.
- [130] S. Bianco, G. Ciocca, and D. Marelli, “Evaluating the Performance of Structure from Motion Pipelines,” *J. Imaging*, vol. 4, no. 8, Art. no. 8, Aug. 2018, doi: 10.3390/jimaging4080098.
- [131] “What Is An Orthomosaic? Orthomosaic Maps & Orthophotos Explained,” *dronegenuity*, Nov. 27, 2019. <https://www.dronegenuity.com/orthomosaic-maps-explained/> (accessed Oct. 26, 2021).
- [132] J. Zhang, S. Xu, Y. Zhao, J. Sun, S. Xu, and X. Zhang, “Aerial orthoimage generation for UAV remote sensing: Review,” *Inf. Fusion*, vol. 89, pp. 91–120, Jan. 2023, doi: 10.1016/j.inffus.2022.08.007.
- [133] C. V. Lopez, H. Anglberger, and U. Stilla, “Fusion of very high resolution SAR and optical images for the monitoring of urban areas,” in *2017 Joint Urban Remote Sensing Event (JURSE)*, Mar. 2017, pp. 1–4. doi: 10.1109/JURSE.2017.7924551.
- [134] D. Lee and J. Park, “Development of Solar-Panel Monitoring Method Using Unmanned Aerial Vehicle and Thermal Infrared Sensor,” *IOP Conf. Ser. Mater. Sci. Eng.*, vol. 611, no. 1, p. 012085, Oct. 2019, doi: 10.1088/1757-899X/611/1/012085.
- [135] A. Chang, J. Jung, M. M. Maeda, and J. Landivar, “Crop height monitoring with digital imagery from Unmanned Aerial System (UAS),” *Comput. Electron. Agric.*, vol. 141, pp. 232–237, Sep. 2017, doi: 10.1016/j.compag.2017.07.008.
- [136] W. Jackson, G. Dobie, C. MacLeod, G. West, C. Mineo, and L. McDonald, “Error Analysis and Calibration for a Novel Pipe Profiling Tool,” *IEEE Sens. J.*, vol. 20, no. 7, pp. 3545–3555, Apr. 2020, doi: 10.1109/JSEN.2019.2960939.
- [137] T. Jayakumar and C. K. Mukhopadhyay, “NDT Techniques: Acoustic Emission,” in *Reference Module in Materials Science and Materials Engineering*, Elsevier, 2016, p. B9780128035818027000. doi: 10.1016/B978-0-12-803581-8.02682-5.
- [138] D. G. Honegger and D. Wijewickreme, “25 - Seismic risk assessment for oil and gas pipelines,” in *Handbook of Seismic Risk Analysis and Management of Civil Infrastructure Systems*, S. Tesfamariam and K. Goda, Eds., in Woodhead Publishing Series in Civil and Structural Engineering. Woodhead Publishing, 2013, pp. 682–715. doi: 10.1533/9780857098986.4.682.
- [139] Encyclopedia Britanica, “Ultrasonics,” Oct. 06, 2017. <https://www.britannica.com/science/ultrasonics> (accessed Oct. 26, 2022).

- [140] Olympus, “Ultrasonic Transducers: Technical Notes.” Olympus Industrial Resources, Aug. 06, 2019. Accessed: Oct. 20, 2020. [Online]. Available: <https://www.olympus-ims.com/en/resources/white-papers/ultrasonic-transducer-technical-notes/>
- [141] P. Anugonda, J. S. Wiehn, and J. A. Turner, “Diffusion of ultrasound in concrete,” *Ultrasonics*, vol. 39, no. 6, pp. 429–435, Oct. 2001, doi: 10.1016/S0041-624X(01)00077-4.
- [142] C. Fei *et al.*, “Ultrahigh Frequency (100 MHz–300 MHz) Ultrasonic Transducers for Optical Resolution Medical Imaging,” *Sci. Rep.*, vol. 6, no. 1, Art. no. 1, Jun. 2016, doi: 10.1038/srep28360.
- [143] P. Cawley, “Practical Long Range Guided Wave Inspection — Managing Complexity,” in *AIP Conference Proceedings*, Bellingham, Washington (USA): AIP, 2003, pp. 22–40. doi: 10.1063/1.1570116.
- [144] M. J. S. Lowe, “WAVE PROPAGATION | Guided Waves in Structures,” in *Encyclopedia of Vibration*, Elsevier, 2001, pp. 1551–1559. doi: 10.1006/rwvb.2001.0173.
- [145] M. S. Greenwood and J. A. Bamberger, “Measurement of viscosity and shear wave velocity of a liquid or slurry for on-line process control,” *Ultrasonics*, vol. 39, no. 9, pp. 623–630, Aug. 2002, doi: 10.1016/S0041-624X(02)00372-4.
- [146] Encyclopedia Britanica, “Bulk Modulus,” Oct. 12, 2022. <https://www.britannica.com/science/bulk-modulus> (accessed Oct. 25, 2022).
- [147] M. J. S. Lowe, “WAVE PROPAGATION | Waves in an Unbounded Medium,” in *Encyclopedia of Vibration*, Elsevier, 2001, pp. 1565–1570. doi: 10.1006/rwvb.2001.0140.
- [148] Encyclopedia Britanica, “Young’s Modulus,” Aug. 30, 2022. <https://www.britannica.com/science/Youngs-modulus> (accessed Oct. 25, 2022).
- [149] The Engineering Toolbox, “Poisson’s Ratio,” 2008. https://www.engineeringtoolbox.com/poissons-ratio-d_1224.html (accessed Oct. 26, 2022).
- [150] M. H. Rosli, R. S. Edwards, and Y. Fan, “In-plane and out-of-plane measurements of Rayleigh waves using EMATs for characterising surface cracks,” *NDT E Int.*, vol. 49, pp. 1–9, Jul. 2012, doi: 10.1016/j.ndteint.2012.03.002.
- [151] P. Wilcox, M. Lowe, and P. Cawley, “Omnidirectional guided wave inspection of large metallic plate structures using an EMAT array,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 52, no. 4, pp. 653–665, Apr. 2005, doi: 10.1109/TUFFC.2005.1428048.

- [152] S. C. Olisa, M. A. Khan, and A. Starr, “Review of Current Guided Wave Ultrasonic Testing (GWUT) Limitations and Future Directions,” *Sensors*, vol. 21, no. 3, Art. no. 3, Jan. 2021, doi: 10.3390/s21030811.
- [153] M. Vasilev *et al.*, “Non-contact in-process ultrasonic screening of thin fusion welded joints,” *J. Manuf. Process.*, vol. 64, pp. 445–454, Apr. 2021, doi: 10.1016/j.jmapro.2021.01.033.
- [154] M. Tabatabaeipour, J. Hettler, S. Delrue, and K. Van Den Abeele, “Visualization of Delaminations in Composite Structures Using a Baseline-Free, Sparse Array Imaging Technique Based on Nonlinear Lamb Wave Propagation,” *Acta Acust. United Acust.*, vol. 103, no. 6, pp. 987–997, Nov. 2017, doi: 10.3813/AAA.919128.
- [155] E. A. Foster *et al.*, “Inspection of nuclear assets with limited access using Feature Guided Waves,” *NDT E Int.*, vol. 131, p. 102695, Oct. 2022, doi: 10.1016/j.ndteint.2022.102695.
- [156] M. J. S. Lowe, “WAVE PROPAGATION | Interaction of Waves with Boundaries,” in *Encyclopedia of Vibration*, Elsevier, 2001, pp. 1559–1564. doi: 10.1006/rwvb.2001.0172.
- [157] M. Johansmann and G. Wirth, “Laser Doppler vibrometry for measuring vibration in ultrasonic transducers,” in *Ultrasonic Transducers*, K. Nakamura, Ed., in Woodhead Publishing Series in Electronic and Optical Materials. Woodhead Publishing, 2012, pp. 277–313. doi: 10.1533/9780857096302.2.277.
- [158] G. Hübschen, “Electromagnetic acoustic transducers,” in *Ultrasonic Transducers*, K. Nakamura, Ed., in Woodhead Publishing Series in Electronic and Optical Materials. Woodhead Publishing, 2012, pp. 36–69. doi: 10.1533/9780857096302.1.36.
- [159] D. A. Hutchins and A. Neild, “Airborne ultrasound transducers,” in *Ultrasonic Transducers*, K. Nakamura, Ed., in Woodhead Publishing Series in Electronic and Optical Materials. Woodhead Publishing, 2012, pp. 374–407. doi: 10.1533/9780857096302.3.374.
- [160] S. Cochran, “Piezoelectricity and basic configurations for piezoelectric ultrasonic transducers,” in *Ultrasonic Transducers*, K. Nakamura, Ed., in Woodhead Publishing Series in Electronic and Optical Materials. Woodhead Publishing, 2012, pp. 3–35. doi: 10.1533/9780857096302.1.3.
- [161] J. A. Brown and G. R. Lockwood, “Low-cost, high-performance pulse generator for ultrasound imaging,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 49, no. 6, pp. 848–851, Jun. 2002, doi: 10.1109/TUFFC.2002.1009345.
- [162] W. Qiu, Y. Yu, F. K. Tsang, and L. Sun, “A multifunctional, reconfigurable pulse generator for high-frequency ultrasound imaging,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 59, no. 7, pp. 1558–1567, Jul. 2012, doi: 10.1109/TUFFC.2012.2355.

- [163] J. E. Michaels, S. J. Lee, A. J. Croxford, and P. D. Wilcox, "Chirp excitation of ultrasonic guided waves," *Ultrasonics*, vol. 53, no. 1, pp. 265–270, Jan. 2013, doi: 10.1016/j.ultras.2012.06.010.
- [164] T. Misaridis and J. A. Jensen, "Use of modulated excitation signals in medical ultrasound. Part I: basic concepts and expected benefits," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 52, no. 2, pp. 177–191, Feb. 2005, doi: 10.1109/TUFFC.2005.1406545.
- [165] R. Y. Chiao and X. Hao, "Coded excitation for diagnostic ultrasound: a system developer's perspective," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 52, no. 2, pp. 160–170, Feb. 2005, doi: 10.1109/TUFFC.2005.1406543.
- [166] B. W. Drinkwater and P. D. Wilcox, "Ultrasonic arrays for non-destructive evaluation: A review," *NDT E Int.*, vol. 39, no. 7, pp. 525–541, Oct. 2006, doi: 10.1016/j.ndteint.2006.03.006.
- [167] M. J. S. Lowe, "ULTRASONICS," in *Encyclopedia of Vibration*, Elsevier, 2001, pp. 1437–1441. doi: 10.1006/rwvb.2001.0143.
- [168] Olympus, "High-Temperature Ultrasonic Testing." <https://www.olympus-ims.com/en/high-temperature-ultrasonic-testing/> (accessed Oct. 27, 2022).
- [169] E. Brunner, "How Ultrasound System Considerations Influence Front-End Component Choice." Analog Devices, Inc, 2002. Accessed: May 11, 2023. [Online]. Available: <https://www.analog.com/en/analog-dialogue/articles/ultrasound-considerations-influence-front-end.html>
- [170] K. A. Fowler, G. M. Elfbaum, K. A. Smith, and T. J. Nelligan, "Theory and Application of Precision Ultrasonic Thickness Gaging," *Insight - Non-Destr. Test. Cond. Monit.*, vol. 2, no. 10, pp. 1–10, Oct. 1997.
- [171] Technical Committee ISO/TC 135, "ISO 16809:2017 Non-destructive testing – Ultrasonic thickness measurement," International Standards Organisation, Geneva, Switzerland, International Standard ISO 16809:2017, Nov. 2017. Accessed: Nov. 01, 2019. [Online]. Available: <https://www.iso.org/standard/72430.html>
- [172] E07 Committee, "E797: Standard Practice for Measuring Thickness by Manual Ultrasonic Pulse-Echo Contact Method," ASTM International, Dec. 2015. doi: 10.1520/E0797_E0797M-15.
- [173] C. Mineo *et al.*, "Robotic Geometric and Volumetric Inspection of High Value and Large Scale Aircraft Wings," in *2019 IEEE 5th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, Torino, Italy: IEEE, Jun. 2019, pp. 82–86. doi: 10.1109/MetroAeroSpace.2019.8869667.

- [174] E. Duernberger, C. MacLeod, D. Lines, C. Loukas, and M. Vasilev, “Adaptive optimisation of multi-aperture ultrasonic phased array imaging for increased inspection speeds of wind turbine blade composite panels,” *NDT E Int.*, vol. 132, p. 102725, Dec. 2022, doi: 10.1016/j.ndteint.2022.102725.
- [175] R. McMillan *et al.*, “Crawler-based automated non-contact ultrasonic inspection of large structural assets,” in *Proceedings of the ASME 2022 49th Annual Review of Progress in Quantitative Nondestructive Evaluation*, San Diego, CA, USA, Jul. 2022, p. 6.
- [176] W. Z. Stepniewski and C. N. Keys, *Rotary-wing aerodynamics*. New York: Dover Publications, 1984.
- [177] J. G. Leishman, *Principles of helicopter aerodynamics*, 2nd ed. in Cambridge aerospace series, no. 18. Cambridge: Cambridge university press, 2006.
- [178] R. W. Prouty, *Helicopter performance, stability, and control*. Malabar, FL: R.E. Krieger Pub. Co, 1995.
- [179] P. Pounds, R. Mahony, and P. Corke, “Modelling and Control of a Quad-Rotor Robot,” in *Proceedings of the Australasian Conference on Robotics and Automation*, Auckland, New Zealand, Dec. 2006, p. 10.
- [180] P. Pounds, R. Mahony, and P. Corke, “Modelling and control of a large quadrotor robot,” *Control Eng. Pract.*, vol. 18, no. 7, pp. 691–699, Jul. 2010, doi: 10.1016/j.conengprac.2010.02.008.
- [181] F. Furrer, M. Burri, M. Achtelik, and R. Siegwart, “RotorS—A Modular Gazebo MAV Simulator Framework,” in *Robot Operating System (ROS): The Complete Reference (Volume 1)*, A. Koubaa, Ed., in Studies in Computational Intelligence. Cham: Springer International Publishing, 2016, pp. 595–625. doi: 10.1007/978-3-319-26054-9_23.
- [182] P. Martin and E. Salaün, “The true role of accelerometer feedback in quadrotor control,” in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 1623–1629. doi: 10.1109/ROBOT.2010.5509980.
- [183] R. Mahony, V. Kumar, and P. Corke, “Multirotor Aerial Vehicles: Modeling, Estimation, and Control of Quadrotor,” *IEEE Robot. Autom. Mag.*, vol. 19, no. 3, pp. 20–32, Sep. 2012, doi: 10.1109/MRA.2012.2206474.
- [184] T. Lee, M. Leok, and N. H. McClamroch, “Geometric tracking control of a quadrotor UAV on SE(3),” in *49th IEEE Conference on Decision and Control (CDC)*, Dec. 2010, pp. 5420–5425. doi: 10.1109/CDC.2010.5717652.

- [185] M. Kamel *et al.*, “The Voliro Omniorientational Hexacopter: An Agile and Maneuverable Tilttable-Rotor Aerial Vehicle,” *IEEE Robot. Autom. Mag.*, vol. 25, no. 4, pp. 34–44, Dec. 2018, doi: 10.1109/MRA.2018.2866758.
- [186] Technical Committee ISO/TC 20, “ISO 1151-2:1985 Flight dynamics - Concepts, quantities and Symbols -Part 2: Motions of the aircraft and the atmosphere relative to the Earth,” International Organization for Standardization, Geneva, Switzerland, International Standard ISO 1151-2:1985, 1985.
- [187] E. W. Weisstein, “Euler’s Rotation Theorem.” <https://mathworld.wolfram.com/EulersRotationTheorem.html> (accessed Nov. 22, 2022).
- [188] J. Diebel, “Representing attitude: Euler angles, unit quaternions, and rotation vectors,” *Matrix*, vol. 58, no. 15–16, pp. 1–35, 2006.
- [189] D. Baraff, “An Introduction to Physically Based Modeling:” Accessed: Aug. 18, 2022. [Online]. Available: <http://www.cs.cmu.edu/~baraff/sigcourse/notesd1.pdf>
- [190] W. R. Hamilton, *Elements of Quaternions*, 1st ed. London, UK: Longmans, Green & Co., 1866. Accessed: Nov. 22, 2022. [Online]. Available: <https://ia600700.us.archive.org/5/items/elementsofquater00hamirich/elementsofquater00hamirich.pdf>
- [191] B. Siciliano and O. Khatib, Eds., *Springer handbook of robotics*. Berlin: Springer, 2008.
- [192] T. Bajd, M. Mihelj, and M. Munih, “Rotation and Orientation,” in *Introduction to Robotics*, Dordrecht: Springer Netherlands, 2013, pp. 9–36. doi: 10.1007/978-94-007-6101-8_2.
- [193] O. M. O’Reilly, *Intermediate Dynamics for Engineers: Newton-Euler and Lagrangian Mechanics*, 2nd ed. Cambridge University Press, 2020. doi: 10.1017/9781108644297.
- [194] K. M. Lynch and F. C. Park, *Modern robotics: mechanics, planning, and control*, Reprinted with additions and Corrections. Cambridge New York, NY Port Melbourne: Cambridge University Press, 2017.
- [195] E. Altug, J. P. Ostrowski, and R. Mahony, “Control of a quadrotor helicopter using visual feedback,” in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, May 2002, pp. 72–77 vol.1. doi: 10.1109/ROBOT.2002.1013341.
- [196] T. Hamel, R. Mahony, R. Lozano, and J. Ostrowski, “Dynamic Modelling and Configuration Stabilization for an X4-Flyer,” in *IFAC Proceedings Volumes*, in 15th IFAC World Congress, vol. 35. Barcelona, Spain, Jul. 2002, pp. 217–222. doi: 10.3182/20020721-6-ES-1901.00848.

- [197] P. Pounds, R. Mahony, P. Hynes, and J. Roberts, "Design of a Four-Rotor Aerial Robot," in *Proc. 2002 Australasian Conference on Robotics and Automation*, Auckland, New Zealand, Nov. 2002, pp. 145–150.
- [198] P. E. I. Pounds, "Design, Construction and Control of a Large Quadrotor Micro Air Vehicle," Australian National University, Canberra, Australia, 2007. [Online]. Available: <https://core.ac.uk/download/pdf/160609737.pdf>
- [199] S. Bouabdallah, P. Murrieri, and R. Siegwart, "Design and control of an indoor micro quadrotor," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*, Apr. 2004, pp. 4393–4398 Vol.5. doi: 10.1109/ROBOT.2004.1302409.
- [200] S. Bouabdallah, A. Noth, and R. Siegwart, "PID vs LQ control techniques applied to an indoor micro quadrotor," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, Sep. 2004, pp. 2451–2456 vol.3. doi: 10.1109/IROS.2004.1389776.
- [201] P. Castillo, A. Dzul, and R. Lozano, "Real-time stabilization and tracking of a four rotor mini-rotorcraft," in *2003 European Control Conference (ECC)*, Cambridge UK, Sep. 2003, pp. 3123–3128. doi: 10.23919/ECC.2003.7086519.
- [202] P. Castillo, A. Dzul, and R. Lozano, "Real-time stabilization and tracking of a four-rotor mini rotorcraft," *IEEE Trans. Control Syst. Technol.*, vol. 12, no. 4, pp. 510–516, Jul. 2004, doi: 10.1109/TCST.2004.825052.
- [203] W. M. Silver, "On the Equivalence of Lagrangian and Newton-Euler Dynamics for Manipulators," *Int. J. Robot. Res.*, vol. 1, no. 2, pp. 60–70, Jun. 1982, doi: 10.1177/027836498200100204.
- [204] M. Ryll *et al.*, "6D interaction control with aerial robots: The flying end-effector paradigm," *Int. J. Robot. Res.*, vol. 38, no. 9, pp. 1045–1062, Aug. 2019, doi: 10.1177/0278364919856694.
- [205] E. Ebeid, M. Skriver, K. H. Terkildsen, K. Jensen, and U. P. Schultz, "A survey of Open-Source UAV flight controllers and flight simulators," *Microprocess. Microsyst.*, vol. 61, pp. 11–20, Sep. 2018, doi: 10.1016/j.micpro.2018.05.002.
- [206] Open Robotics, "Gazebo." <https://gazebo.org/home> (accessed Nov. 24, 2022).
- [207] Px4, "PX4/jMAVSim." PX4 Autopilot for Drones, Jul. 30, 2022. Accessed: Nov. 24, 2022. [Online]. Available: <https://github.com/PX4/jMAVSim>

- [208] Microsoft Research, “Home - AirSim,” 2021. <https://microsoft.github.io/AirSim/> (accessed Nov. 24, 2022).
- [209] Mathworks, “MATLAB and Simulink for Unmanned Aerial Vehicles,” 2022. <https://uk.mathworks.com/solutions/robotics/uav.html> (accessed Nov. 24, 2022).
- [210] Coppelia Robotics, “CoppeliaSim: Robot Simulator,” Sep. 2022. <https://www.coppeliarobotics.com/> (accessed Nov. 24, 2022).
- [211] NVIDIA, “Isaac Sim,” Dec. 11, 2019. <https://developer.nvidia.com/isaac-sim> (accessed Nov. 24, 2022).
- [212] S. Aggarwal and N. Kumar, “Path planning techniques for unmanned aerial vehicles: A review, solutions, and challenges,” *Comput. Commun.*, vol. 149, pp. 270–299, Jan. 2020, doi: 10.1016/j.comcom.2019.10.014.
- [213] N. Cao and A. F. Lynch, “Inner–Outer Loop Control for Quadrotor UAVs With Input and State Constraints,” *IEEE Trans. Control Syst. Technol.*, vol. 24, no. 5, pp. 1797–1804, Sep. 2016, doi: 10.1109/TCST.2015.2505642.
- [214] N. S. Özbek, M. Önkol, and M. Ö. Efe, “Feedback control strategies for quadrotor-type aerial robots: a survey,” *Trans. Inst. Meas. Control*, vol. 38, no. 5, pp. 529–554, May 2016, doi: 10.1177/0142331215608427.
- [215] H. Shraim, A. Awada, and R. Youness, “A survey on quadrotors: Configurations, modeling and identification, control, collision avoidance, fault diagnosis and tolerant control,” *IEEE Aerosp. Electron. Syst. Mag.*, vol. 33, no. 7, pp. 14–33, Jul. 2018, doi: 10.1109/MAES.2018.160246.
- [216] T. P. Nascimento and M. Saska, “Position and attitude control of multi-rotor aerial vehicles: A survey,” *Annu. Rev. Control*, vol. 48, pp. 129–146, Jan. 2019, doi: 10.1016/j.arcontrol.2019.08.004.
- [217] H. Nguyen, M. Kamel, K. Alexis, and R. Siegwart, “Model Predictive Control for Micro Aerial Vehicles: A Survey,” in *2021 European Control Conference (ECC)*, Virtual, Jun. 2021, pp. 1556–1563. doi: 10.23919/ECC54610.2021.9654841.
- [218] J. C. A. Barata and M. S. Hussein, “The Moore–Penrose Pseudoinverse: A Tutorial Review of the Theory,” *Braz. J. Phys.*, vol. 42, no. 1, pp. 146–165, Apr. 2012, doi: 10.1007/s13538-011-0052-z.
- [219] T. A. Johansen and T. I. Fossen, “Control allocation—A survey,” *Automatica*, vol. 49, no. 5, pp. 1087–1103, May 2013, doi: 10.1016/j.automatica.2013.01.035.

- [220] M. Kamel, M. Burri, and R. Siegwart, “Linear vs Nonlinear MPC for Trajectory Tracking Applied to Rotary Wing Micro Aerial Vehicles,” *IFAC-Pap.*, vol. 50, no. 1, pp. 3463–3469, Jul. 2017, doi: 10.1016/j.ifacol.2017.08.849.
- [221] “ArduPilot Open Source Autopilot.” <http://ardupilot.org/> (accessed Sep. 03, 2018).
- [222] “PX4: The Professional Autopilot,” *PX4 Open Source Autopilot*. <http://px4.io/> (accessed Aug. 25, 2018).
- [223] “Betaflight: Open Source Flight Controller,” Sep. 03, 2018. <https://github.com/betaflight/betaflight> (accessed Sep. 03, 2018).
- [224] Dronecode Project, “MAVLink | Introduction & Developer Guide,” Dec. 09, 2022. <https://mavlink.io/en/> (accessed Dec. 14, 2022).
- [225] Open Robotics, “Robotic Operating System (ROS).” <https://www.ros.org/about-ros/> (accessed Nov. 20, 2020).
- [226] R. Duffy, “Mapping the UK drone industry,” *nesta*, Sep. 06, 2018. <https://www.nesta.org.uk/blog/uk-drone-industry-map/> (accessed Dec. 05, 2022).
- [227] Cyberhawk, “Drone Inspection | Asset Maintenance,” 2022. <https://thecyberhawk.com/> (accessed Dec. 06, 2022).
- [228] SkySpecs, “Renewable energy asset management software and services,” 2022. <https://skyspecs.com/> (accessed Dec. 06, 2022).
- [229] Field Group, “Serving the world with insights beyond data,” 2022. <https://field.group/> (accessed Dec. 06, 2022).
- [230] T. Czimmermann *et al.*, “Visual-Based Defect Detection and Classification Approaches for Industrial Applications—A SURVEY,” *Sensors*, vol. 20, no. 5, p. 1459, Mar. 2020, doi: 10.3390/s20051459.
- [231] S. Dorafshan, R. J. Thomas, and M. Maguire, “Comparison of deep convolutional neural networks and edge detectors for image-based crack detection in concrete,” *Constr. Build. Mater.*, vol. 186, pp. 1031–1045, Oct. 2018, doi: 10.1016/j.conbuildmat.2018.08.011.
- [232] A. Khaloo, D. Lattanzi, K. Cunningham, R. Dell’Andrea, and M. Riley, “Unmanned aerial vehicle inspection of the Placer River Trail Bridge through image-based 3D modelling,” *Struct. Infrastruct. Eng.*, vol. 14, no. 1, pp. 124–136, Jan. 2018, doi: 10.1080/15732479.2017.1330891.

- [233] S. Zhao, F. Kang, J. Li, and C. Ma, “Structural health monitoring and inspection of dams based on UAV photogrammetry with image 3D reconstruction,” *Autom. Constr.*, vol. 130, p. 103832, Oct. 2021, doi: 10.1016/j.autcon.2021.103832.
- [234] D. Zhang, R. Watson, G. Dobie, C. MacLeod, A. Khan, and G. Pierce, “Quantifying impacts on remote photogrammetric inspection using unmanned aerial vehicles,” *Eng. Struct.*, vol. 209, p. 109940, Apr. 2020, doi: 10.1016/j.engstruct.2019.109940.
- [235] S. Omari, P. Gohl, M. Burri, M. Achtelik, and R. Siegwart, “Visual industrial inspection using aerial robots,” in *Proceedings of the 2014 3rd International Conference on Applied Robotics for the Power Industry*, Oct. 2014, pp. 1–5. doi: 10.1109/CARPI.2014.7030056.
- [236] J. Quenzel, M. Nieuwenhuisen, D. Droschel, M. Beul, S. Houben, and S. Behnke, “Autonomous MAV-based Indoor Chimney Inspection with 3D Laser Localization and Textured Surface Reconstruction,” *J. Intell. Robot. Syst.*, May 2018, doi: 10.1007/s10846-018-0791-y.
- [237] A. Bircher, M. S. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, “Receding horizon path planning for 3D exploration and surface inspection,” *Auton. Robots*, vol. 42, Feb. 2018, doi: 10.1007/s10514-016-9610-0.
- [238] D. T. Connor *et al.*, “Radiological comparison of a FDNPP waste storage site during and after construction,” *Environ. Pollut.*, vol. 243, pp. 582–590, Dec. 2018, doi: 10.1016/j.envpol.2018.08.099.
- [239] T. Sugimoto, K. Sugimoto, I. Uechi, N. Utagawa, and C. Kuroda, “Outer Wall Inspection Using Acoustic Irradiation Induced Vibration from UAV for Noncontact Acoustic Inspection Method,” in *2018 IEEE International Ultrasonics Symposium (IUS)*, Oct. 2018, pp. 1–9. doi: 10.1109/ULTSYM.2018.8579757.
- [240] K. Dinh, N. Gucunski, K. Tran, A. Novo, and T. Nguyen, “Full-resolution 3D imaging for concrete structures with dual-polarization GPR,” *Autom. Constr.*, vol. 125, p. 103652, May 2021, doi: 10.1016/j.autcon.2021.103652.
- [241] A. Massaro, N. Savino, S. Selicato, A. Panarese, A. Galiano, and G. Dipierro, “Thermal IR and GPR UAV and Vehicle Embedded Sensor Non-Invasive Systems for Road and Bridge Inspections,” in *2021 IEEE International Workshop on Metrology for Industry 4.0 IoT (MetroInd4.0 IoT)*, Jun. 2021, pp. 248–253. doi: 10.1109/MetroInd4.0IoT51437.2021.9488483.
- [242] E. Valence, M. Baraer, E. Rosa, F. Barbecot, and C. Monty, “Drone-based ground-penetrating radar (GPR) application to snow hydrology,” *The Cryosphere*, vol. 16, no. 9, pp. 3843–3860, Sep. 2022, doi: 10.5194/tc-16-3843-2022.

- [243] R. Bähnemann *et al.*, “Under the Sand: Navigation and Localization of a Micro Aerial Vehicle for Landmine Detection with Ground-Penetrating Synthetic Aperture Radar,” *Field Robot.*, vol. 2, no. 1, pp. 1028–1067, Mar. 2022, doi: 10.55417/fr.2022034.
- [244] “TEXO DSI - Advanced UAV Integrated UT Gauge Unveiled.” <https://texodsi.co.uk/news-post/advanced-uav-integrated-ut-gauge-unveiled/> (accessed Oct. 30, 2019).
- [245] “Multigauge 6500 Drone with Thickness Gauge,” *Tritex NDT Ltd.* <https://www.tritexndt.com/product/multigauge-6500-drone-including-multigauge-6000> (accessed Oct. 24, 2021).
- [246] Air Control Entech, “Ultrasonic Inspection UAV,” *Air Control Entech.* <https://aircontrolentech.com/> (accessed Dec. 02, 2022).
- [247] “Thickness Measurement by Ultrasonic NDT,” *Terra-Inspectioneering.* <https://www.terra-inspectioneering.com/service-en/thickness-measurement-by-ultrasonic-ndt> (accessed Dec. 02, 2022).
- [248] D. Bizley, “Texo Drone Survey and Inspection deploys world’s first UT integrated UAV system,” *Oilfield Technology*, Sep. 29, 2017. <https://www.oilfieldtechnology.com/hse/29092017/texo-drone-survey-and-inspection-deploys-worlds-first-ut-integrated-uav-system/> (accessed Dec. 02, 2022).
- [249] L. M. González-deSantos, J. Martínez-Sánchez, H. González-Jorge, F. Navarro-Medina, and P. Arias, “UAV payload with collision mitigation for contact inspection,” *Autom. Constr.*, vol. 115, p. 103200, Mar. 2020, doi: 10.1016/j.autcon.2020.103200.
- [250] B. B. Kocer, T. Tjahjowidodo, M. Pratama, and G. G. L. Seet, “Inspection-while-flying: An autonomous contact-based nondestructive test using UAV-tools,” *Autom. Constr.*, vol. 106, p. 102895, Oct. 2019, doi: 10.1016/j.autcon.2019.102895.
- [251] D. Zhang, R. Watson, C. MacLeod, G. Dobie, W. Galbraith, and G. Pierce, “Implementation and evaluation of an autonomous airborne ultrasound inspection system,” *Nondestruct. Test. Eval.*, vol. 37, no. 1, pp. 1–21, Jan. 2022, doi: 10.1080/10589759.2021.1889546.
- [252] S. Hamaza *et al.*, “Sensor Installation and Retrieval Operations Using an Unmanned Aerial Manipulator,” *IEEE Robot. Autom. Lett.*, vol. 4, no. 3, pp. 2793–2800, Jul. 2019, doi: 10.1109/LRA.2019.2918448.
- [253] L. M. González-deSantos, J. Martínez-Sánchez, H. González-Jorge, M. Ribeiro, J. B. de Sousa, and P. Arias, “Payload for Contact Inspection Tasks with UAV Systems,” *Sensors*, vol. 19, no. 17, p. 3752, Jan. 2019, doi: 10.3390/s19173752.

- [254] T. Ikeda *et al.*, “Wall contact by octo-rotor UAV with one DoF manipulator for bridge inspection,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 5122–5127. doi: 10.1109/IROS.2017.8206398.
- [255] R. Voyles and G. Jiang, “Hexrotor UAV platform enabling dextrous interaction with structures — Preliminary work,” in *2012 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, Nov. 2012, pp. 1–7. doi: 10.1109/SSRR.2012.6523891.
- [256] M. Ryll, H. H. Bühlhoff, and P. R. Giordano, “A Novel Overactuated Quadrotor Unmanned Aerial Vehicle: Modeling, Control, and Experimental Validation,” *IEEE Trans. Control Syst. Technol.*, vol. 23, no. 2, pp. 540–556, Mar. 2015, doi: 10.1109/TCST.2014.2330999.
- [257] A. Ollero, M. Tognon, A. Suarez, D. Lee, and A. Franchi, “Past, Present, and Future of Aerial Robotic Manipulators,” *IEEE Trans. Robot.*, vol. 38, no. 1, pp. 626–645, Feb. 2022, doi: 10.1109/TRO.2021.3084395.
- [258] M. Á. Trujillo, J. R. Martínez-de Dios, C. Martín, A. Viguria, and A. Ollero, “Novel Aerial Manipulator for Accurate and Robust Industrial NDT Contact Inspection: A New Tool for the Oil and Gas Inspection Industry,” *Sensors*, vol. 19, no. 6, Mar. 2019, doi: 10.3390/s19061305.
- [259] K. Bodie *et al.*, “Active Interaction Force Control for Contact-Based Inspection with a Fully Actuated Aerial Vehicle,” *IEEE Trans. Robot.*, pp. 1–14, 2020, doi: 10.1109/TRO.2020.3036623.
- [260] M. Tognon *et al.*, “A Truly-Redundant Aerial Manipulator System With Application to Push-and-Slide Inspection in Industrial Plants,” *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 1846–1851, Apr. 2019, doi: 10.1109/LRA.2019.2895880.
- [261] P. J. Sanchez-Cuevas *et al.*, “Fully-Actuated Aerial Manipulator for Infrastructure Contact Inspection: Design, Modeling, Localization, and Control,” *Sensors*, vol. 20, no. 17, p. 4708, Aug. 2020, doi: 10.3390/s20174708.
- [262] K. Bodie *et al.*, “An Omnidirectional Aerial Manipulation Platform for Contact-Based Inspection,” presented at the Robotics: Science and Systems XV, Freiburg im Breisgau, Germany, Jun. 2019. Accessed: Oct. 07, 2019. [Online]. Available: <http://www.roboticsproceedings.org/rss15/p19.html>
- [263] R. Watson *et al.*, “Dry Coupled Ultrasonic Non-Destructive Evaluation Using an Over-Actuated Unmanned Aerial Vehicle,” *IEEE Trans. Autom. Sci. Eng.*, pp. 1–16, 2021, doi: 10.1109/TASE.2021.3094966.

- [264] D. Schmidt and K. Berns, “Climbing robots for maintenance and inspections of vertical structures—A survey of design aspects and technologies,” *Robot. Auton. Syst.*, vol. 61, no. 12, pp. 1288–1305, Dec. 2013, doi: 10.1016/j.robot.2013.09.002.
- [265] Eddyfi, “Scorpion 2 Ultrasonic Tank Shell Inspection | UT Thickness Readings,” 2019. <https://www.eddyfi.com/en/product/scorpion-2> (accessed Dec. 03, 2022).
- [266] MISTRAS Group, “Pipeline Crawler Inspection Services,” 2021. <https://www.mistrasgroup.com/how-we-help/field-inspections/advanced-ndt/digital-radiography/pipe-crawlers/> (accessed Dec. 03, 2022).
- [267] P. Chatzacos, Y. P. Markopoulos, K. Hrissagis, and A. Khalid, “On the Development of a Modular External-pipe Crawling Omni-directional Mobile Robot,” in *Climbing and Walking Robots*, M. O. Tokhi, G. S. Virk, and M. A. Hossain, Eds., Berlin, Heidelberg: Springer, 2006, pp. 693–700. doi: 10.1007/3-540-26415-9_83.
- [268] C. Choi, B. Park, and S. Jung, “The Design and Analysis of a Feeder Pipe Inspection Robot With an Automatic Pipe Tracking System,” *IEEEASME Trans. Mechatron.*, vol. 15, no. 5, pp. 736–745, Oct. 2010, doi: 10.1109/TMECH.2009.2032541.
- [269] M. B. Khan *et al.*, “iCrawl: An Inchworm-Inspired Crawling Robot,” *IEEE Access*, vol. 8, pp. 200655–200668, 2020, doi: 10.1109/ACCESS.2020.3035871.
- [270] N. B. David and D. Zarrouk, “Design and Analysis of FCSTAR, a Hybrid Flying and Climbing Sprawl Tuned Robot,” *IEEE Robot. Autom. Lett.*, vol. 6, no. 4, pp. 6188–6195, Oct. 2021, doi: 10.1109/LRA.2021.3077851.
- [271] W. Myeong and H. Myung, “Development of a Wall-Climbing Drone Capable of Vertical Soft Landing Using a Tilt-Rotor Mechanism,” *IEEE Access*, vol. 7, pp. 4868–4879, 2019, doi: 10.1109/ACCESS.2018.2889686.
- [272] J. Meng, J. Buzzatto, Y. Liu, and M. Liarakapis, “On Aerial Robots with Grasping and Perching Capabilities: A Comprehensive Review,” *Front. Robot. AI*, vol. 8, 2022, Accessed: Apr. 15, 2022. [Online]. Available: <https://www.frontiersin.org/article/10.3389/frobt.2021.739173>
- [273] R. Jarvis, A. Farinha, M. Kovac, and F. Cegla, “NDE sensor delivery using unmanned aerial vehicles,” *Insight - Non-Destr. Test. Cond. Monit.*, vol. 60, no. 8, pp. 463–467, Aug. 2018, doi: 10.1784/insi.2018.60.8.463.
- [274] P. Yu, G. Chamitoff, and K. C. Wong, “Perching Upside Down with Bi-directional Thrust Quadrotor,” in *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, Sep. 2020, pp. 1697–1703. doi: 10.1109/ICUAS48674.2020.9213946.

- [275] H. Zhang, E. Lerner, B. Cheng, and J. Zhao, “Compliant Bistable Grippers Enable Passive Perching for Micro Aerial Vehicles,” *IEEEASME Trans. Mechatron.*, vol. 26, no. 5, pp. 2316–2326, Oct. 2021, doi: 10.1109/TMECH.2020.3037303.
- [276] F. J. Garcia Rubiales, P. Ramon Soria, B. C. Arrue, and A. Ollero, “Soft-Tentacle Gripper for Pipe Crawling to Inspect Industrial Facilities Using UAVs,” *Sensors*, vol. 21, no. 12, Art. no. 12, Jan. 2021, doi: 10.3390/s21124142.
- [277] A. Lopez-Lora, P. J. Sanchez-Cuevas, A. Suarez, A. Garofano-Soldado, A. Ollero, and G. Heredia, “MHYRO: Modular HYbrid RObot for contact inspection and maintenance in oil gas plants,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 1268–1275. doi: 10.1109/IROS45743.2020.9341639.
- [278] J. Cacace, G. A. Fontanelli, and V. Lippiello, “A Novel Hybrid Aerial-Ground Manipulator for Pipeline Inspection tasks,” in *2021 Aerial Robotic Systems Physically Interacting with the Environment (AIRPHARO)*, Oct. 2021, pp. 1–6. doi: 10.1109/AIRPHARO52252.2021.9571034.
- [279] S. Zhao, F. Ruggiero, G. A. Fontanelli, V. Lippiello, Z. Zhu, and B. Siciliano, “Nonlinear Model Predictive Control for the Stabilization of a Wheeled Unmanned Aerial Vehicle on a Pipe,” *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 4314–4321, Oct. 2019, doi: 10.1109/LRA.2019.2931821.
- [280] S. R. Nekoo, J. Á. Acosta, G. Heredia, and A. Ollero, “A benchmark mechatronics platform to assess the inspection around pipes with variable pitch quadrotor for industrial sites,” *Mechatronics*, vol. 79, p. 102641, Nov. 2021, doi: 10.1016/j.mechatronics.2021.102641.
- [281] Flyability, “Elios 3.” <https://www.flyability.com/elios-3> (accessed Dec. 01, 2022).
- [282] P. De Petris *et al.*, “RMF-Owl: A Collision-Tolerant Flying Robot for Autonomous Subterranean Exploration,” in *2022 International Conference on Unmanned Aircraft Systems (ICUAS)*, Jun. 2022, pp. 536–543. doi: 10.1109/ICUAS54217.2022.9836115.
- [283] Cleo Robotics, “Dronut X1,” 2021. <https://cleorobotics.com/> (accessed Jan. 26, 2022).
- [284] H. Herrmann and J. Pagan, “Showcase of use of Multicopters for Inspection – Offshore and Onshore Wind,” presented at the BINDT Multicopters Workshop, Manchester, UK, Jul. 09, 2018. [Online]. Available: <http://www.bindt.org/admin/Downloads/Multicopters%20for%20Inspection%20-%2004%20Herrmann%20-%20EDF%20RD%20UKC%20BINDT%202018.pdf>
- [285] M. Tranzatto *et al.*, “CERBERUS: Autonomous Legged and Aerial Robotic Exploration in the Tunnel and Urban Circuits of the DARPA Subterranean Challenge,” *Field Robot.*, vol. 2, no. 1, pp. 274–324, Mar. 2022, doi: 10.55417/fr.2022011.

- [286] Royal Mail Group Ltd, “Drones: connecting remote communities across the UK,” 2022. <https://www.royalmail.com/sustainability/environment/drones-connecting-remote-communities-across-the-uk> (accessed Dec. 13, 2022).
- [287] CAA, “CAP722: Unmanned Aircraft System Operations in UK Airspace – Policy and Guidance, 9th ed.” Dec. 2022. Accessed: Dec. 13, 2022. [Online]. Available: https://publicapps.caa.co.uk/docs/33/CAP722_Edition_9.pdf
- [288] CAA, “CAP1861: BVLOS in Non-Segregated Airspace.” 2020. Accessed: Oct. 07, 2021. [Online]. Available: <https://publicapps.caa.co.uk/docs/33/CAP%201861%20-%20BVLOS%20Fundamentals%20v2.pdf>
- [289] “Sees.ai | Sees.ai.” <https://www.sees.ai/> (accessed Oct. 07, 2021).
- [290] Connected Places Catapult, “Implementing an Open-Access UTM Framework for the UK.” UK Department for Transport, Jun. 2021. Accessed: Dec. 13, 2022. [Online]. Available: <https://cp.catapult.org.uk/wp-content/uploads/2021/06/Implementing-an-Open-Access-UTM-Framework.pdf>
- [291] Skydio, “Skydio Dock™,” 2022. <https://www.skydio.com/skydio-dock/> (accessed Dec. 13, 2022).
- [292] DJI, “DJI Dock,” 2022. <https://www.dji.com/uk/photo> (accessed Dec. 13, 2022).
- [293] Percepto, “Percepto Base: Autonomous drone-in-a-box docking station,” 2022. <https://percepto.co/drone-in-a-box/percepto-base/> (accessed Dec. 13, 2022).
- [294] Z. Dukowitz, “One Month of Work Saved on Complex Nuclear Inspection Using the Elios 2,” Oct. 2021. <https://www.flyability.com/casestudies/asme-xi-nuclear-inspection> (accessed Oct. 21, 2021).
- [295] F. P. García Márquez and A. M. Peco Chacón, “A review of non-destructive testing on wind turbines blades,” *Renew. Energy*, vol. 161, pp. 998–1010, Dec. 2020, doi: 10.1016/j.renene.2020.07.145.
- [296] N. Meyendorf, N. Ida, R. Singh, and J. Vrana, Eds., *Handbook of Nondestructive Evaluation 4.0*. Cham: Springer International Publishing, 2022. doi: 10.1007/978-3-030-73206-6.
- [297] Y. Abdulla and Y. Narasimha, “Successful inspection of two large ammonia storage tanks,” *Process Saf. Prog.*, vol. 28, no. 1, pp. 45–59, 2009, doi: 10.1002/prs.10278.

- [298] International Association of Classification Societies (IACS), “Safe Use of Rafts or Boats for Survey,” International Association of Classification Societies (IACS), Guidance No. 39, 2009.
- [299] J. Harfield, “Client guide for the inspection of steel chimneys,” Association of Technical Lightning & Access Specialists (ATLAS), London, UK, ASG 003, Dec. 2018. Accessed: May 14, 2020. [Online]. Available: <https://atlas.org.uk/?mdocs-file=1364>
- [300] L. T. Popoola, A. S. Grema, G. K. Latinwo, B. Gutti, and A. S. Balogun, “Corrosion problems during oil and gas production and its mitigation,” *Int. J. Ind. Chem.*, vol. 4, no. 1, p. 35, Sep. 2013, doi: 10.1186/2228-5547-4-35.
- [301] L. R. Hilbert, A. R. Black, F. Andersen, and T. Mathiesen, “Inspection and monitoring of corrosion inside monopile foundations for offshore wind turbines,” presented at the EUROCORR, Stockholm, Sweden, Sep. 2011, p. 14.
- [302] D. Mitchell *et al.*, “A review: Challenges and opportunities for artificial intelligence and robotics in the offshore wind sector,” *Energy AI*, vol. 8, p. 100146, May 2022, doi: 10.1016/j.egyai.2022.100146.
- [303] K. Kabbabe Poleo, W. J. Crowther, and M. Barnes, “Estimating the impact of drone-based inspection on the Levelised Cost of electricity for offshore wind farms,” *Results Eng.*, vol. 9, p. 100201, Mar. 2021, doi: 10.1016/j.rineng.2021.100201.
- [304] G. Splitt, “New Applications for Dry Coupling Probes,” in *Non-Destructive Testing*, J. M. Farley and R. W. Nichols, Eds., London, UK: Pergamon, Sep. 1987, pp. 2300–2304. doi: 10.1016/B978-0-08-036221-2.50015-2.
- [305] J. M. Allin, P. Cawley, and M. J. S. Lowe, “Adhesive disbond detection of automotive components using first mode ultrasonic resonance,” *NDT E Int.*, vol. 36, no. 7, pp. 503–514, Oct. 2003, doi: 10.1016/S0963-8695(03)00045-8.
- [306] B. Drinkwater and P. Cawley, “An Ultrasonic Wheel Probe Alternative to Liquid Coupling,” in *Review of Progress in Quantitative Nondestructive Evaluation*, D. O. Thompson and D. E. Chimenti, Eds., Boston, MA: Springer US, 1995, pp. 983–989. doi: 10.1007/978-1-4615-1987-4_124.
- [307] P. Khalili and P. Cawley, “The choice of ultrasonic inspection method for the detection of corrosion at inaccessible locations,” *NDT E Int.*, vol. 99, pp. 80–92, Oct. 2018, doi: 10.1016/j.ndteint.2018.06.003.

- [308] O. Trushkevych *et al.*, “Miniaturised SH EMATs for Fast Robotic Screening of Wall Thinning in Steel Plates,” *IEEE Sens. J.*, vol. 21, no. 2, pp. 1386–1394, Jan. 2021, doi: 10.1109/JSEN.2020.3021526.
- [309] J. Isla and F. Cegla, “Coded Excitation for Pulse-Echo Systems,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 64, no. 4, pp. 736–748, Apr. 2017, doi: 10.1109/TUFFC.2017.2661383.
- [310] “Voliro Airborne Robotics.” <https://www.voliro.com/> (accessed Jun. 29, 2020).
- [311] Vicon Motion Systems Ltd, “Vicon Applications: Engineering.” <https://www.vicon.com/applications/engineering/> (accessed Mar. 12, 2020).
- [312] “Total Stations | Leica Geosystems.” <https://leica-geosystems.com/en-GB/products/total-stations> (accessed Nov. 01, 2019).
- [313] D. Scaramuzza and Z. Zhang, “Aerial Robots, Visual-Inertial Odometry of,” in *Encyclopedia of Robotics*, M. H. Ang, O. Khatib, and B. Siciliano, Eds., Berlin, Heidelberg: Springer, 2020, pp. 1–9. doi: 10.1007/978-3-642-41610-1_71-1.
- [314] C. G. Mayhew, R. G. Sanfelice, and A. R. Teel, “Quaternion-Based Hybrid Control for Robust Global Attitude Tracking,” *IEEE Trans. Autom. Control*, vol. 56, no. 11, pp. 2555–2566, Nov. 2011, doi: 10.1109/TAC.2011.2108490.
- [315] “Home | Eddyfi.” <https://www.eddyfi.com> (accessed Oct. 09, 2020).
- [316] B. Drinkwater, R. Dwyer-Joyce, and P. Cawley, “A study of the transmission of ultrasound across solid–rubber interfaces,” *J. Acoust. Soc. Am.*, vol. 101, no. 2, pp. 970–981, 1997, doi: 10.1121/1.418055.
- [317] D. Benstock, F. Cegla, and M. Stone, “The influence of surface roughness on ultrasonic thickness measurements,” *J. Acoust. Soc. Am.*, vol. 136, no. 6, pp. 3028–3039, Dec. 2014, doi: 10.1121/1.4900565.
- [318] G. Dobie, R. Summan, C. MacLeod, G. Pierce, and W. Galbraith, “An automated miniature robotic vehicle inspection system,” in *40th Annual Review of Progress in Quantitative Non-Destructive Evaluation*, Baltimore, Maryland, USA, 2014, pp. 1881–1888. doi: 10.1063/1.4865053.
- [319] Cygnus instruments, “Cygnus 2+ Hands-Free.” Accessed: Dec. 02, 2020. [Online]. Available: <https://www.cygnusinstruments.co.uk/product/cygnus-2-hands-free-2/>

- [320] Waygate Technologies, “DM5E corrosion thickness gauges,” *Baker Hughes Digital Solutions*. <https://www.bakerhughesds.com/ndt-ultrasonic-testing/thickness-gauges/dm5e-corrosion-thickness-gauges> (accessed Dec. 01, 2020).
- [321] C. N. MacLeod, G. Dobie, S. G. Pierce, R. Summan, and M. Morozov, “Machining-Based Coverage Path Planning for Automated Structural Inspection,” *IEEE Trans. Autom. Sci. Eng.*, vol. 15, no. 1, pp. 202–213, Jan. 2018, doi: 10.1109/TASE.2016.2601880.
- [322] “Kuka KR 6 R900 sixx | Datasheet.” https://www.kuka.com/-/media/kuka-downloads/imported/6b77eecacfe542d3b736af377562ecaa/0000205456_en.pdf (accessed Aug. 11, 2022).
- [323] A. V. Oppenheim and R. W. Schaffer, *Discrete-time signal processing*, 3rd ed. Upper Saddle River: Pearson, 2010.
- [324] G. E. Farin and D. Hansford, *Mathematical principles for scientific computing and visualization*. Wellesley, Mass: AK Peters, 2008.
- [325] E. W. Weisstein, “Line-Plane Intersection,” *Mathworld-- A Wolfram Web Resource*. <http://mathworld.wolfram.com/Line-PlaneIntersection.html> (accessed Oct. 24, 2019).
- [326] M. Lamari, “Corrosion Gaging with Dual Element Transducers.” Evident - Olympus, Apr. 25, 2003. Accessed: Jan. 04, 2020. [Online]. Available: <https://www.olympus-ims.com/en/applications/corrosion-gaging-dual-element-transducers/>
- [327] Dwight E. Gray, *American Institute of Physics Handbook*, 3rd ed. New York, NY, USA: McGraw-Hill, 1972.
- [328] M. Parrilla, J. J. Anaya, and C. Fritsch, “Digital signal processing techniques for high accuracy ultrasonic range measurements,” *IEEE Trans. Instrum. Meas.*, vol. 40, no. 4, pp. 759–763, Aug. 1991, doi: 10.1109/19.85348.
- [329] D. Marioli, C. Narduzzi, C. Offelli, D. Petri, E. Sardini, and A. Taroni, “Digital time-of-flight measurement for ultrasonic sensors,” *IEEE Trans. Instrum. Meas.*, vol. 41, no. 1, pp. 93–97, Feb. 1992, doi: 10.1109/19.126639.
- [330] R. Summan *et al.*, “Spatial calibration of large volume photogrammetry based metrology systems,” *Measurement*, vol. 68, pp. 189–200, May 2015, doi: 10.1016/j.measurement.2015.02.054.
- [331] P. Riseborough, “Using the ECL EKF,” *PX4 User Guide*, May 03, 2023. https://docs.px4.io/main/en/advanced_config/tuning_the_ecl_ekf.html (accessed May 07, 2023).

- [332] K. Shoemake, “Animating Rotation with Quaternion Curves,” in *Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques*, in SIGGRAPH ’85. New York, NY, USA: ACM, 1985, pp. 245–254. doi: 10.1145/325334.325242.
- [333] K. S. Arun, T. S. Huang, and S. D. Blostein, “Least-Squares Fitting of Two 3-D Point Sets,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-9, no. 5, pp. 698–700, Sep. 1987, doi: 10.1109/TPAMI.1987.4767965.
- [334] “27MG Ultrasonic Thickness Gage.” [https://www.olympus-ims.com/en/27mg/#!cms\[tab\]=%2F27mg%2Fspecifications](https://www.olympus-ims.com/en/27mg/#!cms[tab]=%2F27mg%2Fspecifications) (accessed Jun. 29, 2020).
- [335] “Multigauge 5600 Hand Held Ultrasonic Thickness Meter | Tritex NDT,” *Tritex NDT Ltd.* <https://www.tritexndt.com/product/mg5600-ultrasonic-thickness-meter> (accessed Nov. 04, 2019).
- [336] J. P. Pauwelussen, “Chapter Two - Fundamentals of Tire Behavior,” in *Essentials of Vehicle Dynamics*, J. P. Pauwelussen, Ed., Oxford: Butterworth-Heinemann, 2015, pp. 7–74. doi: 10.1016/B978-0-08-100036-6.00002-9.
- [337] M. J. Lettich, “Insulation Management and Its Value to Industry,” *Insulation Outlook Magazine*, Sep. 01, 2003. <https://insulation.org/io/articles/insulation-management-and-its-value-to-industry/> (accessed Oct. 11, 2022).
- [338] S. F. Burch, “HOIS/OGTC Guidelines for in-situ inspection of corrosion under insulation (CUI),” ESR Technology Ltd., Milton Park, Abingdon, HOIS-G-023, Jan. 2021. Accessed: Apr. 06, 2022. [Online]. Available: <https://www.bindt.org/shopbindt/hois-documents/HOIS-OGTC-Guidelines-for-In-Situ-Inspection-of-Corrosion-Under-Insulation-CUI-HOIS-G-023-Issue-2.html>
- [339] C. Paz, E. Suárez, C. Gil, and C. Baker, “CFD analysis of the aerodynamic effects on the stability of the flight of a quadcopter UAV in the proximity of walls and ground,” *J. Wind Eng. Ind. Aerodyn.*, vol. 206, p. 104378, Nov. 2020, doi: 10.1016/j.jweia.2020.104378.
- [340] “BLHeli_32,” *GitHub*. <https://github.com/bitdump/BLHeli> (accessed Apr. 21, 2022).
- [341] “ASME B36.10M - Welded and Seamless Wrought Steel Pipe,” American Society of Mechanical Engineers (ASME), Standard ASME B36.10M, 2004.
- [342] R. M. Murray, Z. Li, and S. Sastry, *A mathematical introduction to robotic manipulation*. Boca Raton: CRC Press, 1994.
- [343] “Common Friction Coefficients.” https://www.engineeringtoolbox.com/friction-coefficients-d_778.html (accessed Aug. 17, 2022).

- [344] Tyto Robotics, “Series 1580 Thrust Tester,” *Tyto Robotics*. <https://www.tytorobotics.com/pages/series-1580-1585> (accessed Oct. 07, 2022).
- [345] K. Bodie, Z. Taylor, M. Kamel, and R. Siegwart, “Towards Efficient Full Pose Omnidirectionality with Overactuated MAVs,” in *Proceedings of the 2018 International Symposium on Experimental Robotics*, J. Xiao, T. Kröger, and O. Khatib, Eds., in Springer Proceedings in Advanced Robotics. Buenos Aires, Argentina: Springer International Publishing, Nov. 2018, pp. 85–95. doi: 10.1007/978-3-030-33950-0_8.
- [346] L. Meier, D. Honegger, and M. Pollefeys, “PX4: A node-based multithreaded open source robotics framework for deeply embedded platforms,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 6235–6240. doi: 10.1109/ICRA.2015.7140074.
- [347] Px4, “Safety Configuration (Failsafes) | User Guide,” Mar. 26, 2023. <https://docs.px4.io/main/en/config/safety.html> (accessed May 31, 2023).
- [348] S. R. Nekoo, J. Á. Acosta, G. Heredia, and A. Ollero, “Soft-Landing of Multi-Rotor Drones using a Robust Nonlinear Control and Wind Modeling,” in *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*, Jun. 2021, pp. 1070–1079. doi: 10.1109/ICUAS51884.2021.9476763.
- [349] S. W. Shepperd, “Quaternion from Rotation Matrix,” *J. Guid. Control*, vol. 1, no. 3, pp. 223–224, May 1978, doi: 10.2514/3.55767b.
- [350] D. Brescianini, M. Hehn, and R. D’Andrea, “Nonlinear Quadcopter Attitude Control: Technical Report,” ETH Zurich, Zurich, Switzerland, 2013. doi: 10.3929/ETHZ-A-009970340.
- [351] Vicon, “Tracker Motion Capture Software for VR and Object Tracking.” <http://www.vicon.com/products/software/tracker> (accessed Jun. 30, 2020).
- [352] “FAQs: How Accurate / Precise Are Your Systems?,” *Vicon*. <https://www.vicon.com/support/faqs/> (accessed Aug. 31, 2022).
- [353] United States Air Force (USAF), “MIL-STD-150A: Photographic Lenses,” US Airforce, Standard MIL-STD-150A, May 1959.
- [354] S. F. Burch, “HOIS Guidance on Image Quality for UAV/UAS based external remote visual inspection in the oil & gas industry,” HOIS, Oxfordshire, UK, Guidance HOIS-G-005, Jun. 2018. Accessed: Nov. 16, 2019. [Online]. Available: <https://www.esrtechnology.com/images/hoispages/publicRPguides/HOIS-G-005-HOIS-Guidance-for-UAV-based-external-RVI---Issue-1-Final-Open.pdf>

- [355] “Agisoft Metashape: Standard Edition.” <https://www.agisoft.com/features/standard-edition/> (accessed Sep. 08, 2022).
- [356] T. Rabbani, “Efficient Hough Transform for Automatic Detection of Cylinders in Point Clouds,” presented at the International Society for Photogrammetry and Remote Sensing Workshop, Enschede, Netherlands: ISPRS, Sep. 2005, pp. 60–65.
- [357] T. Rabbani, “PROTO Toolbox,” 2005. <http://viztronics.com/tahir/data/PROTO/doc/menu.html> (accessed Mar. 26, 2019).
- [358] E. W. Weisstein, “Cylindrical Coordinates.” <https://mathworld.wolfram.com/> (accessed Sep. 07, 2022).
- [359] Blender Foundation, “Blender 3.3,” *blender.org*. <https://www.blender.org/> (accessed Sep. 09, 2022).
- [360] A. Shukla, H. Xiaoqian, and H. Karki, “Autonomous tracking and navigation controller for an unmanned aerial vehicle based on visual data for inspection of oil and gas pipelines,” in *2016 16th International Conference on Control, Automation and Systems (ICCAS)*, Oct. 2016, pp. 194–200. doi: 10.1109/ICCAS.2016.7832320.
- [361] M. Salvago, F. J. Pérez-Gran, J. Parra, M. A. Trujillo, and A. Viguria, “Robust and Efficient Pose Estimation of Pipes for Contact Inspection using Aerial Robots,” in *2021 Aerial Robotic Systems Physically Interacting with the Environment (AIRPHARO)*, Oct. 2021, pp. 1–6. doi: 10.1109/AIRPHARO52252.2021.9571059.
- [362] E. Guerra, R. Munguía, and A. Grau, “UAV Visual and Laser Sensors Fusion for Detection and Positioning in Industrial Applications,” *Sensors*, vol. 18, no. 7, Art. no. 7, Jul. 2018, doi: 10.3390/s18072071.
- [363] “Agisoft Metashape User Manual - Standard Edition, Version 1.8.” Aug. 17, 2022. Accessed: Sep. 15, 2022. [Online]. Available: https://www.agisoft.com/pdf/metashape_1_8_en.pdf
- [364] A01 Committee, “Specification for Pipe, Steel, Black and Hot-Dipped, Zinc-Coated, Welded and Seamless,” ASTM International, Standard ASTM A53/A53M-22, Aug. 2022. doi: 10.1520/A0053_A0053M-22.
- [365] “ISO 4628-1:2016 Evaluation of degradation of coatings: General introduction and designation system,” International Standards Organisation, Geneva, Switzerland, International Standard ISO 4628-1:2016, 2016.

- [366] S. Agnisarman, S. Lopes, K. Chalil Madathil, K. Piratla, and A. Gramopadhye, “A survey of automation-enabled human-in-the-loop systems for infrastructure visual inspection,” *Autom. Constr.*, vol. 97, pp. 52–76, Jan. 2019, doi: 10.1016/j.autcon.2018.10.019.
- [367] F.-C. Chen and M. R. Jahanshahi, “NB-CNN: Deep Learning-Based Crack Detection Using Convolutional Neural Network and Naïve Bayes Data Fusion,” *IEEE Trans. Ind. Electron.*, vol. 65, no. 5, pp. 4392–4400, May 2018, doi: 10.1109/TIE.2017.2764844.
- [368] Q. Luo, X. Fang, L. Liu, C. Yang, and Y. Sun, “Automated Visual Defect Detection for Flat Steel Surface: A Survey,” *IEEE Trans. Instrum. Meas.*, vol. 69, no. 3, pp. 626–644, Mar. 2020, doi: 10.1109/TIM.2019.2963555.
- [369] Z. Shang and Z. Shen, “Single-pass inline pipeline 3D reconstruction using depth camera array,” *Autom. Constr.*, vol. 138, p. 104231, Jun. 2022, doi: 10.1016/j.autcon.2022.104231.
- [370] R. Jarvis, “Current Deflection NDE for Pipe Inspection and Monitoring,” EngD, Imperial College London, London, UK, 2017. Accessed: Mar. 09, 2021. [Online]. Available: <https://spiral.imperial.ac.uk/handle/10044/1/56865>
- [371] Y. Furukawa and C. Hernández, *Multi-view stereo: a tutorial*. in *Foundation and trends in computer graphics and vision*, no. 9,1/2. Boston Delft: Now, 2015.
- [372] D. C. Brown, “Decentering Distortion of Lenses,” *Photogramm. Eng.*, vol. 32, no. 3, pp. 444–462, 1966.
- [373] A. W. Fitzgibbon, “Simultaneous linear estimation of multiple view geometry and lens distortion,” in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, Kauai, HI, USA: IEEE Comput. Soc, 2001, p. I-125–I-132. doi: 10.1109/CVPR.2001.990465.
- [374] MATLAB, “Camera Calibrator.” <https://uk.mathworks.com/help/vision/ref/cameracalibrator-app.html> (accessed Nov. 02, 2022).
- [375] OpenCV, “Camera Calibration.” https://docs.opencv.org/4.x/dc/dbb/tutorial_py_calibration.html (accessed Nov. 02, 2022).
- [376] J. Salvi, X. Armangué, and J. Batlle, “A comparative review of camera calibrating methods with accuracy evaluation,” *Pattern Recognit.*, vol. 35, no. 7, pp. 1617–1635, Jul. 2002, doi: 10.1016/S0031-3203(01)00126-1.

- [377] D. Griffiths and H. Burningham, “Comparison of pre- and self-calibrated camera calibration models for UAS-derived nadir imagery for a SfM application,” *Prog. Phys. Geogr. Earth Environ.*, vol. 43, no. 2, pp. 215–235, Apr. 2019, doi: 10.1177/0309133318788964.
- [378] D. Scaramuzza and F. Fraundorfer, “Visual Odometry [Tutorial],” *IEEE Robot. Autom. Mag.*, vol. 18, no. 4, pp. 80–92, Dec. 2011, doi: 10.1109/MRA.2011.943233.
- [379] C. Cadena *et al.*, “Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age,” *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, Dec. 2016, doi: 10.1109/TRO.2016.2624754.
- [380] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual–Inertial, and Multimap SLAM,” *IEEE Trans. Robot.*, vol. 37, no. 6, pp. 1874–1890, Dec. 2021, doi: 10.1109/TRO.2021.3075644.
- [381] D. G. Lowe, “Distinctive Image Features from Scale-Invariant Keypoints,” *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, Nov. 2004, doi: 10.1023/B:VISI.0000029664.99615.94.
- [382] H. Bay, T. Tuytelaars, and L. Van Gool, “SURF: Speeded Up Robust Features,” in *Computer Vision – ECCV 2006*, A. Leonardis, H. Bischof, and A. Pinz, Eds., in Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2006, pp. 404–417. doi: 10.1007/11744023_32.
- [383] S. Leutenegger, M. Chli, and R. Y. Siegwart, “BRISK: Binary Robust invariant scalable keypoints,” in *2011 International Conference on Computer Vision*, Barcelona, Spain: IEEE, Nov. 2011, pp. 2548–2555. doi: 10.1109/ICCV.2011.6126542.
- [384] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “ORB: An efficient alternative to SIFT or SURF,” in *2011 International Conference on Computer Vision*, Barcelona, Spain: IEEE, Nov. 2011, pp. 2564–2571. doi: 10.1109/ICCV.2011.6126544.
- [385] E. Rosten and T. Drummond, “Machine Learning for High-Speed Corner Detection,” in *Computer Vision – ECCV 2006*, A. Leonardis, H. Bischof, and A. Pinz, Eds., in Lecture Notes in Computer Science, vol. 3951. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 430–443. doi: 10.1007/11744023_34.
- [386] D. Hutchison *et al.*, “BRIEF: Binary Robust Independent Elementary Features,” in *Computer Vision – ECCV 2010*, K. Daniilidis, P. Maragos, and N. Paragios, Eds., in Lecture Notes in Computer Science, vol. 6314. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 778–792. doi: 10.1007/978-3-642-15561-1_56.

- [387] S. A. K. Tareen and Z. Saleem, "A comparative analysis of SIFT, SURF, KAZE, AKAZE, ORB, and BRISK," in *2018 International Conference on Computing, Mathematics and Engineering Technologies (iCoMET)*, Mar. 2018, pp. 1–10. doi: 10.1109/ICOMET.2018.8346440.
- [388] M. Muja and D. G. Lowe, "Fast Approximate Nearest Neighbors with Automatic Algorithm Configuration," in *Proceedings of the Fourth International Conference on Computer Vision Theory and Applications*, Lisboa, Portugal: SciTePress - Science and Technology Publications, Feb. 2009, pp. 331–340. doi: 10.5220/0001787803310340.
- [389] K. Yousif, A. Bab-Hadiashar, and R. Hoseinnezhad, "An Overview to Visual Odometry and Visual SLAM: Applications to Mobile Robotics," *Intell. Ind. Syst.*, vol. 1, no. 4, pp. 289–311, Dec. 2015, doi: 10.1007/s40903-015-0032-7.
- [390] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981, doi: 10.1145/358669.358692.
- [391] M. I. A. Lourakis and A. A. Argyros, "SBA: A software package for generic sparse bundle adjustment," *ACM Trans. Math. Softw.*, vol. 36, no. 1, p. 2:1-2:30, Mar. 2009, doi: 10.1145/1486525.1486527.
- [392] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle Adjustment — A Modern Synthesis," in *Vision Algorithms: Theory and Practice*, B. Triggs, A. Zisserman, and R. Szeliski, Eds., in Lecture Notes in Computer Science, vol. 1883. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 298–372. doi: 10.1007/3-540-44480-7_21.
- [393] N. Snavely, I. Simon, M. Goesele, R. Szeliski, and S. M. Seitz, "Scene Reconstruction and Visualization From Community Photo Collections," *Proc. IEEE*, vol. 98, no. 8, pp. 1370–1390, Aug. 2010, doi: 10.1109/JPROC.2010.2049330.
- [394] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski, "A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms," in *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, Jun. 2006, pp. 519–528. doi: 10.1109/CVPR.2006.19.
- [395] M. Goesele, N. Snavely, B. Curless, H. Hoppe, and S. M. Seitz, "Multi-View Stereo for Community Photo Collections," in *2007 IEEE 11th International Conference on Computer Vision*, Rio de Janeiro, Brazil: IEEE, 2007, pp. 1–8. doi: 10.1109/ICCV.2007.4408933.
- [396] Y. Furukawa and J. Ponce, "Accurate, Dense, and Robust Multiview Stereopsis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 8, pp. 1362–1376, Aug. 2010, doi: 10.1109/TPAMI.2009.161.

- [397] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski, “Towards Internet-scale multi-view stereo,” in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, San Francisco, CA, USA: IEEE, Jun. 2010, pp. 1434–1441. doi: 10.1109/CVPR.2010.5539802.
- [398] F. Langguth, K. Sunkavalli, S. Hadap, and M. Goesele, “Shading-Aware Multi-view Stereo,” in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., in Lecture Notes in Computer Science. Cham: Springer International Publishing, 2016, pp. 469–485. doi: 10.1007/978-3-319-46487-9_29.
- [399] “CloudCompare - Open Source project.” <http://www.cloudcompare.org/> (accessed Dec. 19, 2022).
- [400] B. Curless and M. Levoy, “A volumetric method for building complex models from range images,” in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques - SIGGRAPH '96*, New Orleans, Louisiana, USA: ACM Press, 1996, pp. 303–312. doi: 10.1145/237170.237269.
- [401] M. Kazhdan, M. Bolitho, and H. Hoppe, “Poisson surface reconstruction,” in *Eurographics Symposium on Geometry Processing*, Cagliari, Sardinia, Italy, Jun. 2006, p. 10. doi: 10.2312/SGP/SGP06/061-070.
- [402] M. Kazhdan and H. Hoppe, “Screened poisson surface reconstruction,” *ACM Trans. Graph.*, vol. 32, no. 3, p. 29:1-29:13, Jul. 2013, doi: 10.1145/2487228.2487237.
- [403] S. Fuhrmann and M. Goesele, “Floating scale surface reconstruction,” *ACM Trans. Graph.*, vol. 33, no. 4, p. 46:1-46:11, Jul. 2014, doi: 10.1145/2601097.2601163.
- [404] B. Ummerhofer and T. Brox, “Global, Dense Multiscale Reconstruction for a Billion Points,” in *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec. 2015, pp. 1341–1349. doi: 10.1109/ICCV.2015.158.
- [405] G. M. Treece, R. W. Prager, and A. H. Gee, “Regularised marching tetrahedra: improved iso-surface extraction,” *Comput. Graph.*, vol. 23, no. 4, pp. 583–598, Aug. 1999, doi: 10.1016/S0097-8493(99)00076-X.
- [406] W. E. Lorensen and H. E. Cline, “Marching cubes: A high resolution 3D surface construction algorithm,” *ACM SIGGRAPH Comput. Graph.*, vol. 21, no. 4, pp. 163–169, Aug. 1987, doi: 10.1145/37402.37422.
- [407] Y. Yan and J. F. Hajjar, “Automated extraction of structural elements in steel girder bridges from laser point clouds,” *Autom. Constr.*, vol. 125, p. 103582, May 2021, doi: 10.1016/j.autcon.2021.103582.

- [408] H. Cui, J. Li, Q. Hu, and Q. Mao, “Real-Time Inspection System for Ballast Railway Fasteners Based on Point Cloud Deep Learning,” *IEEE Access*, vol. 8, pp. 61604–61614, 2020, doi: 10.1109/ACCESS.2019.2961686.
- [409] P. E. Debevec, C. J. Taylor, and J. Malik, “Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach,” in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, in SIGGRAPH ’96. New York, NY, USA: Association for Computing Machinery, Aug. 1996, pp. 11–20. doi: 10.1145/237170.237191.
- [410] L. Wang, S. B. Kang, R. Szeliski, and H.-Y. Shum, “Optimal texture map reconstruction from multiple views,” in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, Kauai, HI, USA: IEEE Comput. Soc, 2001, p. I-347-I-354. doi: 10.1109/CVPR.2001.990496.
- [411] V. Lempitsky and D. Ivanov, “Seamless Mosaicing of Image-Based Texture Maps,” in *2007 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2007, pp. 1–6. doi: 10.1109/CVPR.2007.383078.
- [412] P. Pérez, M. Gangnet, and A. Blake, “Poisson image editing,” in *ACM SIGGRAPH 2003 Papers*, in SIGGRAPH ’03. New York, NY, USA: Association for Computing Machinery, Jul. 2003, pp. 313–318. doi: 10.1145/1201775.882269.
- [413] Z. Chen, J. Zhou, Y. Chen, and G. Wang, “3D Texture Mapping in Multi-view Reconstruction,” in *Advances in Visual Computing*, G. Bebis, R. Boyle, B. Parvin, D. Koracin, C. Fowlkes, S. Wang, M.-H. Choi, S. Mantler, J. Schulze, D. Acevedo, K. Mueller, and M. Papka, Eds., in *Lecture Notes in Computer Science*, vol. 7431. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 359–371. doi: 10.1007/978-3-642-33179-4_35.
- [414] W. Xu and J. Mulligan, “Performance evaluation of color correction approaches for automatic multi-view image and video stitching,” in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Jun. 2010, pp. 263–270. doi: 10.1109/CVPR.2010.5540202.
- [415] K. Hormann, B. Lévy, and A. Sheffer, “Mesh Parameterization: Theory and Practice.” *Siggraph Course Notes*, Aug. 2007. Accessed: Nov. 10, 2022. [Online]. Available: <https://www.inf.usi.ch/hormann/papers/Hormann.2007.MPT.pdf>
- [416] B. Lévy, S. Petitjean, N. Ray, and J. Maillot, “Least squares conformal maps for automatic texture atlas generation,” *ACM Trans. Graph.*, vol. 21, no. 3, pp. 362–371, Jul. 2002, doi: 10.1145/566654.566590.
- [417] V. N. Whittaker, “A review of non-destructive measurement of flaw size,” *Non-Destr. Test.*, vol. 5, no. 2, pp. 92–100, Apr. 1972, doi: 10.1016/0029-1021(72)90101-6.

- [418] R. Halmshaw, *Non-Destructive Testing*, 2nd ed. Northampton, UK: British Institute of NDT, 2004.
- [419] N. Thrane, J. Wismer, H. Konstantin-Hansen, and S. Gade, “Practical use of the ‘Hilbert transform.’” Brüel&Kjær, May 1997. Accessed: Aug. 12, 2022. [Online]. Available: <https://www.bksv.com/media/doc/bo0437.pdf>
- [420] Olympus, “Corrosion Mapping on Carbon Steel Using the RollerFORM Scanner.” <https://www.olympus-ims.com/en/corrosion-mapping-on-carbon-steel-using-the-rollerform-scanner/> (accessed Oct. 28, 2022).
- [421] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, “Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation,” *NDT E Int.*, vol. 38, no. 8, pp. 701–711, Dec. 2005, doi: 10.1016/j.ndteint.2005.04.002.
- [422] J. Camacho, M. Parrilla, and C. Fritsch, “Phase Coherence Imaging,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 56, no. 5, pp. 958–974, May 2009, doi: 10.1109/TUFFC.2009.1128.
- [423] Evident, “Phase Coherence Imaging (PCI) Best Practices: Getting Started with the PCI Technique.” Jul. 2022. Accessed: Oct. 31, 2022. [Online]. Available: https://www.olympus-ims.com/en/.downloads/download/?file=285219336&fl=en_US&inline
- [424] C. Mineo, D. Lines, and D. Cerniglia, “Generalised bisection method for optimum ultrasonic ray tracing and focusing in multi-layered structures,” *Ultrasonics*, vol. 111, p. 106330, Mar. 2021, doi: 10.1016/j.ultras.2020.106330.
- [425] R. Zimmermann *et al.*, “Increasing the speed of automated ultrasonic inspection of as-built additive manufacturing components by the adoption of virtual source aperture,” *Mater. Des.*, vol. 220, p. 110822, Aug. 2022, doi: 10.1016/j.matdes.2022.110822.
- [426] “LTPA 64/128 Product | Peak NDT,” *PeakNDT*. <https://www.peakndt.com/products/ltpa-64-128/> (accessed Oct. 28, 2022).
- [427] M. Clough, M. Fleming, and S. Dixon, “Circumferential guided wave EMAT system for pipeline screening using shear horizontal ultrasound,” *NDT E Int.*, vol. 86, pp. 20–27, Mar. 2017, doi: 10.1016/j.ndteint.2016.11.010.
- [428] T. Kundu, *Ultrasonic nondestructive evaluation: engineering and biological material characterization*. Boca Raton, Florida, USA: CRC Press, 2004.

- [429] M. Kalms, O. Focke, and C. v. Kopylow, “Applications of laser ultrasound NDT methods on composite structures in aerospace industry,” presented at the Ninth International Symposium on Laser Metrology, C. Quan and A. Asundi, Eds., Unknown, Singapore, Sep. 2008, p. 71550E. doi: 10.1117/12.814512.
- [430] D. Pieris *et al.*, “Laser Induced Phased Arrays (LIPA) to detect nested features in additively manufactured components,” *Mater. Des.*, vol. 187, p. 108412, Feb. 2020, doi: 10.1016/j.matdes.2019.108412.
- [431] S. Bennett, “A brief history of automatic control,” *IEEE Control Syst. Mag.*, vol. 16, no. 3, pp. 17–25, Jun. 1996, doi: 10.1109/37.506394.
- [432] S. Dormido and A. Visioli, “PID Control,” in *Encyclopedia of Systems and Control*, J. Baillieul and T. Samad, Eds., Cham: Springer International Publishing, 2021, pp. 1724–1733. doi: 10.1007/978-3-030-44184-5_245.
- [433] S. Skogestad, “Simple analytic rules for model reduction and PID controller tuning,” *J. Process Control*, vol. 13, no. 4, pp. 291–309, Jun. 2003, doi: 10.1016/S0959-1524(02)00062-8.
- [434] K. J. Åström and T. Hägglund, “Automatic tuning of simple regulators with specifications on phase and amplitude margins,” *Automatica*, vol. 20, no. 5, pp. 645–651, Sep. 1984, doi: 10.1016/0005-1098(84)90014-1.
- [435] H. Huang, G. M. Hoffmann, S. L. Waslander, and C. J. Tomlin, “Aerodynamics and control of autonomous quadrotor helicopters in aggressive maneuvering,” in *2009 IEEE International Conference on Robotics and Automation*, Kobe, Japan: IEEE, May 2009, pp. 3277–3282. doi: 10.1109/ROBOT.2009.5152561.
- [436] M. Alpen, K. Frick, and J. Horn, “Nonlinear modeling and position control of an industrial quadrotor with on-board attitude control,” in *2009 IEEE International Conference on Control and Automation*, Dec. 2009, pp. 2329–2334. doi: 10.1109/ICCA.2009.5410536.
- [437] F. Goodarzi, D. Lee, and T. Lee, “Geometric nonlinear PID control of a quadrotor UAV on SE(3),” in *2013 European Control Conference (ECC)*, Jul. 2013, pp. 3845–3850. doi: 10.23919/ECC.2013.6669644.
- [438] R. E. Kalman, “Contributions to the Theory of Optimal Control,” *Boletín Soc. Mat. Mex.*, vol. 5, pp. 102–119, 1960.
- [439] H. L. Trentelman, “Linear Quadratic Optimal Control,” in *Encyclopedia of Systems and Control*, J. Baillieul and T. Samad, Eds., Cham: Springer International Publishing, 2021, pp. 1121–1126. doi: 10.1007/978-3-030-44184-5_201.

- [440] R. Tedrake, “Underactuated Robotics: Algorithms for Walking, Running, Swimming, Flying, and Manipulation (Course Notes for MIT 6.832).” Accessed: Dec. 01, 2022. [Online]. Available: <http://underactuated.mit.edu/>
- [441] I. D. Cowling, O. A. Yakimenko, J. F. Whidborne, and A. K. Cooke, “A prototype of an autonomous controller for a quadrotor UAV,” in *2007 European Control Conference (ECC)*, Jul. 2007, pp. 4001–4008. doi: 10.23919/ECC.2007.7068316.
- [442] P. Foehn and D. Scaramuzza, “Onboard State Dependent LQR for Agile Quadrotors,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 6566–6572. doi: 10.1109/ICRA.2018.8460885.
- [443] V. Utkin, “Variable structure systems with sliding modes,” *IEEE Trans. Autom. Control*, vol. 22, no. 2, pp. 212–222, Apr. 1977, doi: 10.1109/TAC.1977.1101446.
- [444] R. A. DeCarlo, S. H. Zak, and G. P. Matthews, “Variable structure control of nonlinear multivariable systems: a tutorial,” *Proc. IEEE*, vol. 76, no. 3, pp. 212–232, Mar. 1988, doi: 10.1109/5.4400.
- [445] S. Bouabdallah and R. Siegwart, “Backstepping and Sliding-mode Techniques Applied to an Indoor Micro Quadrotor,” in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, Apr. 2005, pp. 2247–2252. doi: 10.1109/ROBOT.2005.1570447.
- [446] D. Lee, H. Jin Kim, and S. Sastry, “Feedback linearization vs. adaptive sliding mode control for a quadrotor helicopter,” *Int. J. Control Autom. Syst.*, vol. 7, no. 3, pp. 419–428, Jun. 2009, doi: 10.1007/s12555-009-0311-8.
- [447] O. Mofid and S. Mobayen, “Adaptive sliding mode control for finite-time stability of quad-rotor UAVs with parametric uncertainties,” *ISA Trans.*, vol. 72, pp. 1–14, Jan. 2018, doi: 10.1016/j.isatra.2017.11.010.
- [448] F. Muñoz, I. González-Hernández, S. Salazar, E. S. Espinoza, and R. Lozano, “Second order sliding mode controllers for altitude control of a quadrotor UAS: Real-time implementation in outdoor environments,” *Neurocomputing*, vol. 233, pp. 61–71, Apr. 2017, doi: 10.1016/j.neucom.2016.08.111.
- [449] P. V. Kokotovic, “The joy of feedback: nonlinear and adaptive,” *IEEE Control Syst. Mag.*, vol. 12, no. 3, pp. 7–17, Jun. 1992, doi: 10.1109/37.165507.
- [450] H. K. Khalil, *Nonlinear systems*, 3rd ed. Upper Saddle River, N.J: Prentice Hall, 2002.

- [451] T. Madani and A. Benallegue, “Backstepping Control for a Quadrotor Helicopter,” in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China: IEEE, Oct. 2006, pp. 3255–3260. doi: 10.1109/IROS.2006.282433.
- [452] S. Bouabdallah and R. Siegwart, “Full control of a quadrotor,” in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2007, pp. 153–158. doi: 10.1109/IROS.2007.4399042.
- [453] M. Labbadi and M. Cherkaoui, “Robust adaptive backstepping fast terminal sliding mode controller for uncertain quadrotor UAV,” *Aerosp. Sci. Technol.*, vol. 93, p. 105306, Oct. 2019, doi: 10.1016/j.ast.2019.105306.
- [454] J. Richalet, A. Rault, J. L. Testud, and J. Papon, “Model predictive heuristic control: Applications to industrial processes,” *Automatica*, vol. 14, no. 5, pp. 413–428, Sep. 1978, doi: 10.1016/0005-1098(78)90001-8.
- [455] K. Alexis, G. Nikolakopoulos, and A. Tzes, “Model predictive quadrotor control: attitude, altitude and position experimental studies,” *IET Control Theory Appl.*, vol. 6, no. 12, pp. 1812–1827, Aug. 2012, doi: 10.1049/iet-cta.2011.0348.
- [456] M. Kamel, K. Alexis, M. Achtelik, and R. Siegwart, “Fast nonlinear model predictive control for multicopter attitude tracking on $SO(3)$,” in *2015 IEEE Conference on Control Applications (CCA)*, Sep. 2015, pp. 1160–1166. doi: 10.1109/CCA.2015.7320769.
- [457] D. Bicego, J. Mazzetto, R. Carli, M. Farina, and A. Franchi, “Nonlinear Model Predictive Control with Enhanced Actuator Model for Multi-Rotor Aerial Vehicles with Generic Designs,” *J. Intell. Robot. Syst.*, vol. 100, no. 3, pp. 1213–1247, Dec. 2020, doi: 10.1007/s10846-020-01250-9.
- [458] J. Wehbeh, S. Rahman, and I. Sharf, “Distributed Model Predictive Control for UAVs Collaborative Payload Transport,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 11666–11672. doi: 10.1109/IROS45743.2020.9341541.