



Politecnico di Torino

MsC. Mechatronic Engineering

A.a. 2025/2026

Graduation Session September 2026

Downwash-Based Leader Sensing and Relative-State Estimation for Implicit UAV Formation Control

Supervisor:

Prof. Alessandro Rizzo

Prof. Igor Paprotny

Candidate:

Francesco Pettinato

Acknowledgements

I would like to thank my supervisor, Prof. Alessandro Rizzo, for the constant guidance throughout this work, and for the support and availability at every stage of the project.
I am also grateful to Prof. Igor Paprotny for the constant guidance and for the many discussions that shaped the direction of this thesis, and for hosting this work at the University of Illinois Chicago.
My thanks also go to Tommaso Andina, with whom the bidirectional system was developed, for the shared work on the companion return channel.

Table of Contents

List of Figures VII

0 Introduction 1
- 0.1 Close-proximity flight and the need for implicit interaction in UAV formations . . . 1
- 0.2 Limitations of explicit communication in disturbed, short-range multi-UAV scenarios . . . 2
- 0.3 Research problem: estimating the leader state through aerodynamic wake effects . . . 3
- 0.4 Individual contribution of the thesis . . . 4
- 0.5 Structure of the thesis . . . 5

1 Leader-Follower Flight Through Aerodynamic Interaction 7
- 1.1 Overview of leader-follower architectures in close-proximity flight . . 7
- 1.2 Aerodynamic Wake and Downwash as Physical Information Carriers 9
- 1.3 From Disturbance to Measurement: Interpreting Wake-Induced Signals 11
- 1.4 The Baseline One-Way Leader-to-Follower System . . . 13
- 1.5 Open Issues in the Baseline System: Blind Zones, Range Ambiguity and Robustness . . . 15

2 State of the Art on Wake-Based Sensing and Relative Estimation 17
- 2.1 Aerodynamic interaction models for multirotor UAVs . . . 17
- 2.2 Turbulent jet models for rotor downwash characterization . . . 19
- 2.3 Downwash-aware control and compensation strategies . . . 21
- 2.4 Relative localization methods in GPS-denied and communication-limited settings . . . 23
- 2.5 Signal detection methods for weak aerodynamic signatures . . . 25
- 2.6 Positioning of this thesis with respect to existing work . . . 27

3 Modeling the Leader-to-Follower Aerodynamic Channel 29
- 3.1 Physical assumptions and operating conditions . . . 29

3.2 Downwash propagation model and spatial attenuation 31
3.3 Channel observability from onboard follower measurements 35
3.4 Noise sources and uncertainty in close-proximity flight 36
3.5 Signal-to-noise ratio as a function of relative position 38
3.6 Inversion of the wake model for range estimation 41
3.7 Limits of the model and expected failure regions 42

4 Detection and Estimation of the Leader Position 44
4.1 Vector-norm detector for rotationally invariant wake sensing 45
4.2 Adaptive thresholding based on online noise-floor estimation 47
4.3 Debounce logic for robust event detection 48
4.4 Approach-angle estimation from attitude residuals 50
4.5 Range reconstruction through downwash model inversion 51
4.6 Relative position estimation in polar coordinates 52
4.7 Filtering strategies: alpha filter, alpha-beta-gamma predictor and Kalman-based alternatives 53
4.8 Detection-rate and update-rate tradeoffs 54

5 Closed-Loop Pursuit Based on Implicit Sensing 56
5.1 Control objectives for follower alignment and distance regulation . . 57
5.2 Polar-coordinate pursuit controller 58
5.3 Coupling between estimator dynamics and control response 61
5.4 Stability considerations for the sensing-control loop 62
5.5 Robustness to intermittent detections and temporary loss of signal . 63
5.6 Practical implementation on the follower UAV 64

6 Interface with the Return Communication Channel 66
6.1 Role of follower-to-leader feedback in the shared bidirectional system 66
6.2 Minimal description of gesture-based acknowledgements 67
6.3 Interaction between leader sensing and return-channel activity . . . 68
6.4 Avoiding interference between forward sensing and feedback gestures 69
6.5 System-level timing constraints from the follower perspective 70

7 Experimental Validation of Leader Sensing and Pursuit 72
7.1 Experimental setup and instrumentation 72
7.2 Ground-level validation of wake detection 74
7.3 Hover tests for range and angle estimation 75
7.4 Diagonal approach and moving-leader scenarios 77
7.5 Spatial characterization of detection regions and blind zones 77
7.6 Closed-loop pursuit performance 80
7.7 Impact of follower-to-leader gestures on leader sensing 82
7.8 Discussion of experimental limitations 83

8 Scalability of Wake-Based Sensing in Multi-UAV Formations 86
8.1 Extension from one leader-follower pair to UAV chains 87
8.2 Accumulation of estimation error along the formation 88
8.3 String stability from the sensing and pursuit perspective 90
8.4 Graph representation of sensing dependencies 91
8.5 Constraints on formation geometry and inter-vehicle spacing 95
8.6 Scalability limits of wake-based leader tracking 96

9 Conclusions 98
9.1 Summary of the individual contribution 98
9.2 Main experimental findings 99
9.3 Limitations of the proposed sensing and estimation approach 101
9.4 Future work on robust wake-based formation control 102

Bibliography 105

List of Figures

1 Downwash effect in the air. 2

2 Comparison between explicit RF networks and the proposed implicit aerodynamic channel, with the physical and systemic limitations of each communication medium. 3

3 The shared bidirectional communication system. The forward path L→F, developed in this thesis, lets the follower sense the leader from its wake. The return path F→L, developed in the companion thesis, lets the follower reply by modulating its altitude, which the leader decodes with an I/Q matched filter and a receiving state machine. . 4

4 The noise-based detection threshold. The threshold is set at $\pm 4\sigma$ of the hovering noise, so that nominal hovering vibrations fall below it and only larger aerodynamic events are detected. 5

1.1 The implicit sensing pipeline. The aerodynamic interaction appears as uncommanded attitude perturbations. Subtracting the controller setpoints isolates the disturbance, whose peaks are processed into relative range and bearing and passed to the pursuit loop. 13

1.2 Non-monotonic relation between radial offset and induced moment magnitude. A given amplitude corresponds to two offsets, one on each side of the peak, so the inversion is not unique. To obtain a single-valued mapping, the numerical inversion is restricted to the monotone outer branch beyond r_{peak}. The peak location r_{peak} is determined empirically from the flight residuals. 14

3.1 Anemometric characterization bench. The leader is held at fixed hover RPM on a rig; a probe traverses the downwash jet, sampling the mean velocity at points (r, z) along the centreline and at increasing radial offsets. Radial profiles are repeated at azimuths 0° and 45° to check axisymmetry. 32

3.2 Longitudinal axial decay of the leader's downwash measured on the central wake axis ($r = 0$). Markers indicate the empirical mean velocity, with error bars representing the standard deviation across repeated measurements. The solid curve denotes the fitted hyperbolic centreline law $U_C(z) = 444/(z + 34)$ 33

3.3 Transverse radial profiles of the leader's downwash, discretized by axial depth. Solid markers indicate valid empirical measurements; open markers signify spatial coordinates where the local velocity falls below the sensor's absolute detection floor ($\approx 0.2\,\mathrm{m\,s^{-1}}$), denoting a complete structural collapse of the wake. The continuous blue curve illustrates the theoretical turbulent-jet distribution anchored to the respective on-axis peak velocity. Notably, at every depth, the flow sustains peak magnitude until roughly 10 cm, precipitating a severe spatial collapse between 10 cm and 15 cm. 34

3.4 Normalized autocorrelation of the pitch and roll residuals in quiet hover, averaged over the eleven leader-free flights. Both axes stay correlated for several hundred milliseconds, crossing the $1/e$ level at about 400 ms in pitch and 460 ms in roll. The residual is therefore strongly coloured by the attitude loop, and the number of independent noise samples per second is much smaller than the 200 Hz sample count. 39

3.5 Geometry of the wake signal. The leader downwash descends as a jet that reaches the follower below. The induced moment is weak on the axis, where the inflow is symmetric, and strongest at an intermediate offset inside the wake, where the radial gradient of the velocity field is largest. Beyond the edge of the jet the signal falls to the noise floor. The useful sensing region is the shaded band within the downwash, not the axis and not the far field. 40

4.1 Detection zone in the relative plane (r, θ) from simulation, for the OR logic of the baseline and for the vector norm. The OR logic detects (blue) across most angles but leaves a gap around 45°, where the wake splits between pitch and roll and each component falls below the fixed interval. The vector norm detects across the whole disc with no gap. The right panel shows the region the norm recovers (green), a band centred on 45°. 46

4.2 Three simulated approaches at $\theta = 0°$ (t=10 s), $\theta = 45°$ (t=17 s) and $\theta = 90°$ (t=24 s), with the wake injected through the Bauersfeld model. Top: the pitch and roll residuals. Middle: the residual norm against the adaptive and fixed thresholds. Bottom: the raw detection and the debounced output, which reduces the fragmented raw crossings to three clean events, one per approach. 49

4.3 Estimated range against the true range in simulation, as the virtual leader approaches the follower. The virtual leader follows a prescribed trajectory, so the true range is known at every instant. The estimated range, recovered by inverting the wake profile, follows the true range with an error below one millimetre on the monotone side of the peak. 52

5.1 Geometry of the polar pursuit law. The follower commands a velocity $\mathbf{v}_{\text{cmd}}$ along the line of sight to the leader, enforcing the fundamental condition of pure pursuit. The scalar magnitude of that command is dictated by the differential between the estimated range r_{est} and the target setpoint r_{target}, ensuring the follower approaches when beyond the target distance and recedes when within it, ultimately stabilizing on the radial boundary defined by r_{target}. 59

5.2 The radial velocity command parameterized as a function of the range error. A defined dead zone in the vicinity of zero error suppresses the command output, preventing the follower from erroneously reacting to stochastic estimation noise at the equilibrium setpoint. Outside this dead zone, the command scales proportionally with the error via the gain coefficient K_P, generating restorative effort toward the set distance. The command is strictly bounded at $\pm V_{\max}$, restricting the follower's kinematics to velocities safely within the tracking capabilities of the intermittent measurement updates. 60

5.3 Block diagram of the sensing-control loop architecture. The aerodynamic wake induces attitude residuals on the follower airframe; the estimator maps these residuals into a relative spatial position; the pursuit controller subsequently generates a regulated velocity command; and the resulting kinematic motion alters the relative spatial geometry, thereby feeding back into the wake perturbation. Positioned in series within the closed loop, the phase lag and hold characteristics of the estimator directly modulate the resulting control response. 61

6.1 Geometry of the vertical role-swap between the two channels. In sensing mode (left) the follower stays below the leader and reads the wake-induced attitude residuals, while the leader encodes information through its horizontal motion. In reply mode (right) the follower climbs above the leader and superimposes a vertical oscillation, which the leader decodes as a vertical-velocity perturbation. The two channels act on orthogonal axes and never occupy the same vertical configuration. 70

7.1 Experimental setup. The follower nano-quadrotor relies only on its onboard IMU; an identical leader hovers above it over a floor grid used as position reference. For the static tests the leader is placed at fixed lateral offsets r from 0 cm to 40 cm; for the dynamic tests it is moved by hand along a low-speed circular path around the follower (top-view inset). 73

7.2 Angular coverage of the bearing estimate over a pursuit flight. Each wedge is one 30° sector, and its radius is the number of detections with the leader in that sector. The leader is localized in all twelve sectors around the follower, including the diagonal sectors near 45° where the single-axis baseline was blind. The bearing is therefore observable across the full circle. 76

7.3 Estimated leader position in the follower body frame during one extended pursuit flight, coloured by time. The follower sits at the origin. Points populate the plane around it while respecting the two operational bounds: the inner dashed circle is the clamp at the offset of maximum moment (0.13 m), the outer dotted circle is the far-field edge (0.50 m). No estimate falls inside the clamp, confirming the follower stays in the observable annulus while tracking the leader around the full circle. 78

7.4 Wake signal magnitude parameterized against nominal radial offset, derived from the controlled-offset flight telemetry. The signal amplitude is defined as the ninetieth percentile of the residual norm; error bars denote the statistical spread across independent flight trials at each discrete offset. The signal is attenuated along the central axis, escalates to a spatial maximum at an intermediate offset near 10 cm to 20 cm, and subsequently decays at larger distances, while remaining robustly above the ambient noise floor throughout the evaluated domain. Because the prescribed offsets are nominal, this distribution visualizes the macroscopic topological shape of the signal rather than a metrologically exact calibration curve. 79

7.5 Estimated radial range spanning a continuous closed-loop pursuit flight. The dynamic range smoothly fluctuates in response to the leader's relative kinematics, asymptotically increasing toward the far-field boundary during signal dropouts and actively decreasing during proximity encounters. The estimate strictly respects the defined operational bounds, never violating the numerical clamp located at the peak moment offset (dashed line). This confirms that the trailing vehicle consistently maintains a spatial separation where the radial range remains fully observable. 81

8.1 Topological representation of the formation as a sequential chain. Each UAV derives its pursuit trajectory by sensing the downwash of its immediate predecessor, thereby rendering each link a discrete, unidirectional wake-sensing edge. The terminal lead vehicle, V_1, strictly regulates its state toward a mathematically invariant virtual node, ν, which deterministically anchors the entire kinematic chain to a prescribed spatial coordinate. 87

8.2 Theoretical bounding of accumulated state estimation error along the formation chain, normalized against the unitary per-link standard deviation σ_1. Assuming strict statistical independence, the terminal error profile expands proportionally to $\sqrt{n}$. Conversely, total statistical correlation induces linear expansion scaling with n. Given the stochastic heterogeneity of the wake channel, the localized errors fundamentally tend toward independence; thus, the lower $\sqrt{n}$ trajectory characterizes the expected scaling behavior, while the upper linear trajectory defines the absolute conservative limit. . . . 89

8.3 The augmented graph in its two modes. In sensing mode the wake carries from the leader to the follower below it, and the leader regulates toward the reference ν. In reply mode the leader descends below the follower and the sensing edge reverses, while the reference still drives the leader. No edge enters ν in either mode, so the reference is never influenced by the vehicles, and the path from ν to the follower is restored after each reply. 92

8.4 Transient exponential decay profile of the aggregate formation error converging toward the terminal boundary band, mathematically modeled for the minimal pair coupled with the virtual reference node. The error magnitude diminishes exponentially at a rate governed exclusively by the non-zero Laplacian eigenvalue k. The timeout parameter $t^{\star}$ defines the precise temporal coordinate at which the decaying error intersects the specified tracking tolerance ε. Employing the established parameters $k = 0.08\,\mathrm{s}^{-1}$, initial error $e_0 = 0.30\,\mathrm{m}$, and an aggressive tolerance bound $\varepsilon = 0.02\,\mathrm{m}$, the resultant analytical timeout is bounded at approximately 34 s. . . . 94

9.1 Blade-passing signature against blade count, measured on the propeller bench at near-equal shaft speed. The three propellers turn at about the same rate, so the shaft line sits near 69 Hz in all three, but the blade-passing frequency moves with the blade count: 139 Hz for two blades, 208 Hz for three, and 272 Hz for four. Different propellers therefore carry distinct signatures at the same rotation speed, which is the basis for identifying a drone by its wake. All three frequencies lie above the 100 Hz limit of the 200 Hz inertial unit, so none is accessible to the current sensor. 103

Chapter 0

Introduction

0.1 Close-proximity flight and the need for implicit interaction in UAV formations

Over the last few decades, Unmanned Aerial Vehicles (UAVs) have become building blocks of multi-agent systems, in which several vehicles operate in a coordinated way within the same airspace. This growth was driven by the falling cost and size of multi-rotor platforms and, above all, by planning and control algorithms that handle nonlinear dynamics in real time on resource-constrained hardware. Many applications followed, such as formation flight for surveillance and reconnaissance, cooperative inspection of extended infrastructures, and search and rescue in complex environments.

A common element across these applications is the need to operate under close-proximity conditions, referred to in the literature as close-proximity flight. This regime brings out phenomena that become prominent as vehicles approach one another. In particular, the downwash, the high-velocity downward airflow generated by the rotors of a multirotor drone, is the source of a disturbance field that propagates through the surrounding space and follows a turbulent-jet structure from about 2.5 motor-to-motor distances l below the vehicle [1].

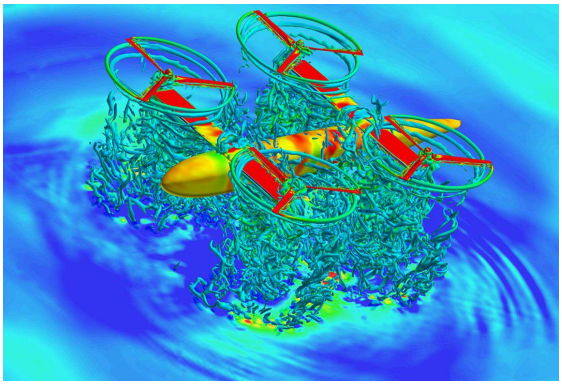

Figure 1: Downwash effect in the air.

This disturbance field can impair a vehicle operating within it if it is neglected or left unmodeled, both for a two-vehicle configuration [2] and for a larger heterogeneous swarm [3]. Most of the literature treats the downwash as a disturbance to be compensated through robust controllers. This work takes a complementary view: the downwash can also carry information between vehicles, which avoids relying only on radio-transmitted signals. This dual nature, disturbance on one side and signal on the other, is the thread of the whole work and is developed in the following chapters, from the prior work to the contributions of this project.

0.2 Limitations of explicit communication in disturbed, short-range multi-UAV scenarios

Coordinating a UAV swarm requires a continuous exchange of state information between vehicles, such as position, velocity, attitude and maneuvering intentions. This information is usually sent over dedicated radio-frequency channels that carry data packets between the vehicles or to a ground station. Radio-frequency channels, however, have several drawbacks, such as limited bandwidth and signal fading.
As studied by Gielis et al. [4], radio-frequency networks are affected by environmental clutter, routing complexity and architectural incompatibility, which undermine communication performance. Radio-frequency channels were not designed to serve multi-agent robot systems, so they are not well suited to this task [4]. Several solutions have been proposed over the years to give multi-robot systems a reliable direct channel, for example the use of machine learning. Reinforcement-learning methods, in particular, can address many classic robotic tasks, such as multi-robot motion planning in the continuous domain and path finding, through Graph Neural Network (GNN) approaches [4]. More recent work extends GNNs to message-aware strategies, so that robots can separate important from less important information

within a flock [4]. Despite these advances, a scalable explicit communication method for multi-robot systems is still lacking, in part because of the shortage of real-world deployments.
Since the limits of explicit communication are not yet overcome, there is a clear case for alternative coordination strategies, such as using the downwash as a source of implicit communication between the drones of a swarm.

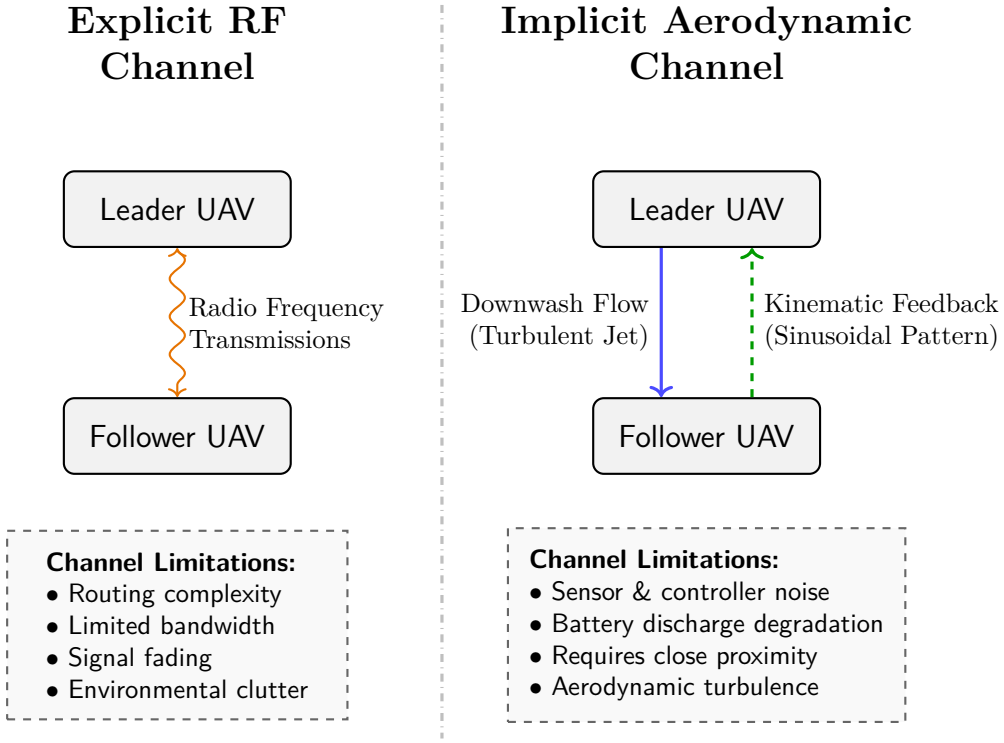

Figure 2: Comparison between explicit RF networks and the proposed implicit aerodynamic channel, with the physical and systemic limitations of each communication medium.

0.3 Research problem: estimating the leader state through aerodynamic wake effects

The downwash carries information about the relative position of two vehicles, and explicit communication is limited at short range. Together these two facts define the research problem of this work: can a follower estimate the relative state of a leader, that is the range and bearing to it, using only its own dynamic response to the leader wake, and can it use that estimate to pursue the leader in closed loop. The follower must do this without a radio link, without markers and without any

added sensor, relying only on the inertial unit it already carries for flight control. The full system that frames this problem is bidirectional, as shown in Figure 3. The leader talks to the follower through its wake, and the follower answers by modulating its own altitude. This thesis develops the forward path, from leader to follower, that is the sensing, estimation and pursuit side. The return path, in which the follower modulates its altitude and the leader decodes the gesture, is the subject of the companion thesis and appears here only as the interface of Chapter 6.

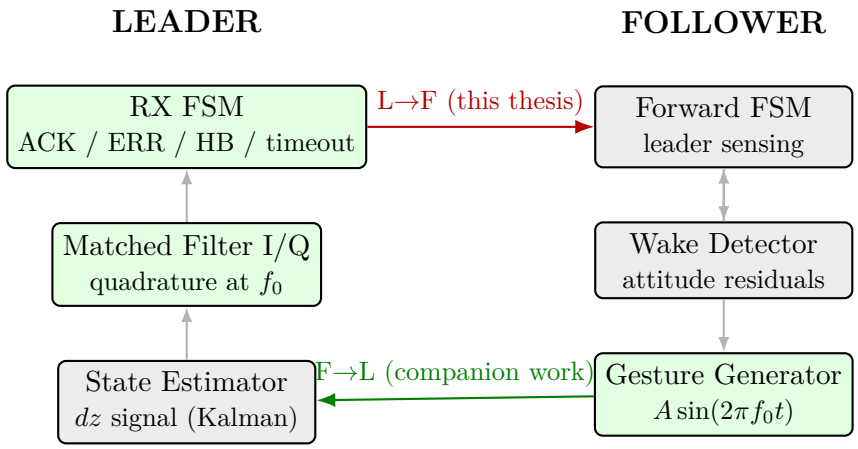

Figure 3: The shared bidirectional communication system. The forward path L→F, developed in this thesis, lets the follower sense the leader from its wake. The return path F→L, developed in the companion thesis, lets the follower reply by modulating its altitude, which the leader decodes with an I/Q matched filter and a receiving state machine.

0.4 Individual contribution of the thesis

The system studied here is a bidirectional implicit communication architecture for a leader-follower pair, built on the baseline of Cerruti [5], which showed through preliminary tests that an implicit leader-to-follower channel is feasible by sensing and interpreting downwash-induced disturbances. Within the forward channel, the individual contribution of this thesis is threefold.

- **Rotation-invariant detection and bearing recovery.** The baseline relied only on the follower pitch angle, which carries information when the leader sits almost directly above the follower. This work combines the pitch and roll residuals into a single quantity that responds to the wake from any horizontal direction, and recovers the bearing to the leader from the two residuals. This extends sensing from the on-axis case to non-negligible horizontal offsets,

and is supported by the fluid-dynamic characterization of the downwash in Bauersfeld et al. [1].

- **Full relative-state reconstruction and polar pursuit.** The baseline returned only a scalar leader velocity. This work reconstructs both the relative range and the relative bearing and uses them to close a pursuit loop written in polar coordinates, which keeps the follower in the region where the wake signal is strongest, away from the on-axis null.
- **Noise-based detection threshold.** The baseline used a fixed threshold calibrated once under specific conditions, which ignored the growth of sensor and controller noise as the battery discharges. This work sets the threshold from the measured noise instead, estimated online from the follower signal, so that the threshold tracks the noise floor rather than sitting at a fixed value. The final form is an ordered-statistic detector, adopted after a simpler variance-based threshold was found to freeze on the real flight logs.

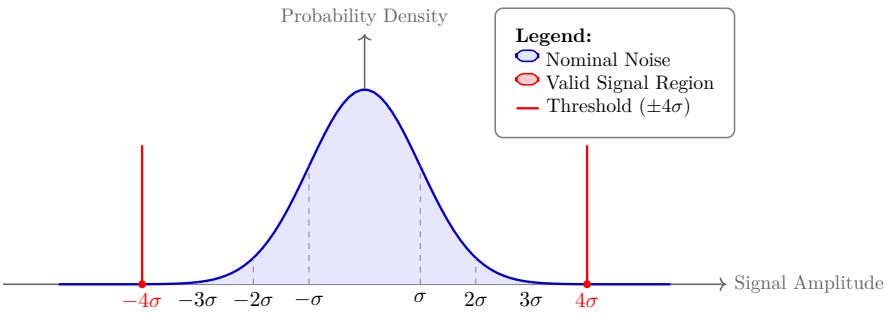

Figure 4: The noise-based detection threshold. The threshold is set at $\pm 4\sigma$ of the hovering noise, so that nominal hovering vibrations fall below it and only larger aerodynamic events are detected.

The methodology is two-fold. The prior work is first reproduced and checked to establish a solid basis, and the elements above are then added to extend it. The forward channel is placed within the shared bidirectional system, and its coexistence with the return channel is analyzed so that forward sensing and the return gestures do not interfere.

0.5 Structure of the thesis

- Chapter 1 describes leader-follower flight through aerodynamic interaction, the wake as a carrier of information, the one-way baseline and its open issues, and states the requirements for autonomous leader sensing.

- Chapter 2 reviews the state of the art on wake-based sensing, turbulent-jet models of rotor downwash, downwash-aware control and relative localization in GPS-denied settings.
- Chapter 3 models the leader-to-follower aerodynamic channel, its spatial attenuation, its observability from onboard measurements and its noise.
- Chapter 4 presents the detection and estimation pipeline, from the vector-norm detector to the polar range and bearing estimate.
- Chapter 5 builds the closed-loop pursuit controller and studies its coupling with the estimator.
- Chapter 6 describes the interface with the follower-to-leader return channel.
- Chapter 7 reports the experimental work on detection, estimation and pursuit.
- Chapter 8 examines how wake-based sensing scales to leader-follower chains.
- Chapter 9 summarizes the contribution and outlines future work.

Chapter 1

Leader-Follower Flight Through Aerodynamic Interaction

During a flight, a quadrotor generates a pronounced aerodynamic wake that coalesces into a highly structured, high-velocity stream beneath the vehicle, thereby producing an aerodynamic perturbation that influences the other drones. While a substantial body of literature on close-proximity flight models this interaction as disturbances to be rejected, this work aims to investigate another paradigm: exploiting the downwash to model an implicit communication channel, since the aerodynamic forces and moments acting on the Follower due to the Leader are intrinsically coupled with the relative spatial configuration between the two vehicles. Therefore, the Follower can infer the Leader's position solely by quantifying its own dynamic response to the downwash, thus eliminating the necessity for dedicated exteroceptive sensors. This chapter thus establishes the theoretical and practical foundations for this approach by reviewing the fluid-dynamic principles that render the passive aerodynamic sensing viable for modelling the communication channel. Finally, this chapter formalizes a rigorous set of requirements that must be satisfied to achieve robust, autonomous leader sensing and closed-loop pursuit.

1.1 Overview of leader-follower architectures in close-proximity flight

A formation of unmanned aerial vehicles (UAVs) is defined as a multi-agent system designed to maintain a prescribed geometric configuration during coordinated movement. In particular, formation methodologies seek to drive each agent to

an assigned state within this collective geometry; as categorized by Oh et al. [6], these strategies are classified by the state variables designated for regulation. In particular, the Leader-Follower framework could be classified as a displacement-based architecture, meaning that the communication channel relies strongly on the relative scalar ranges and directional vectors to adjacent agents. Specifically, the Leader usually navigates along a predetermined trajectory while the Followers are tasked with maintaining a fixed relative pose, either with respect to the Leader or to a preceding follower. Since the inverse communication channel relies solely on the local state information between the drones, this topology can be expanded without the computational overhead of centralized, global planning. This type of communication presents some advantages, since, for example, Onboard Global Navigation Satellite Systems (GNSS), even with Real-Time Kinematic Corrections, provide a relative accuracy that is low with respect to the scale of a close formation, thereby forcing a large inter-vehicle separation to guarantee safe flight and thus making the GNSS unsuitable for the tight spacing required in close-proximity UAV formations [7, 8]. Nonetheless, GNSS is also subject to degradation in urban multipath and fails indoors. Moreover, alternative technologies such as Ultra-Wideband (UWB) or Vision Systems that track a drone through structural models or fiducial markers heavily depend on uninterrupted line-of-sight, adequate illumination, and high frame rates, thus making them unreliable in cluttered or fast-changing scenes [9]. Furthermore, it has been widely studied how onboard cameras and the processing needed for these pipelines burden the computational power on micro and nano-scale UAVs, thus resulting in a perception update rate that is far below the control-loop rate; hence, adding latency to the loop [10, 9]. Such hardware penalties are particularly prohibitive for micro-aerial vehicles (MAVs). In particular, for the purposes of this work, a nano-quadrotor such as the Parrot Mambo possesses a total mass of 63 g and operates under stringent energetic and computational constraints.

Table 1.1: Key parameters of the Parrot Mambo platform utilized in this research. The motor spacing l serves as the characteristic length scale for the aerodynamic wake model, while the half-wheelbase dictates the moment arm through which asymmetric rotor loading acts upon the airframe.

Quantity	Value
Mass	$\approx 63\,\mathrm{g}$
Configuration	Quadrotor, X-frame geometry
Motor spacing l	$\approx 65\,\mathrm{mm}$
Half-wheelbase	$\approx 0.06\,\mathrm{m}$
Onboard sensing	Inertial Measurement Unit (IMU), downward optical flow, ultrasonic altimeter
Control loop rate	$200\,\mathrm{Hz}$

Any mass addition, such as LIDAR lasers or other sensing systems, might disproportionately degrade flight endurance and alter the platform inertia tensors, thus compromising the efficacy of controllers. Therefore, to tackle these difficulties, an implicit communication channel provides a way to transmit state information continuously through a physical medium, the downwash, which exists independently of intentional data transmission and cannot be deactivated and, most importantly, does not depend on other sensing systems. Furthermore, implicit sensing is highly advantageous because it operates without dedicated transreceiver hardware and is not jeopardized by environments subject to severe electromagnetic interference, thus allowing the formation topology and detection to be maintained solely based on the physical coupling induced by the downwash.

1.2 Aerodynamic Wake and Downwash as Physical Information Carriers

In order to exploit the aerodynamic wake as a reliable signal, a rigorous characterization of its fluid dynamics is indeed requisite. Therefore, being considered an isolated rotor in hover, momentum theory [11] determines that the mean induced velocity v_i of the airflow traversing a rotor plane of area A and operating in an atmosphere of density ρ is:

$$v_i = \sqrt{\frac{T}{2\,\rho\,A}}, \tag{1.1}$$

where T is the generated thrust. Equation 1.1 follows from actuator-disk momentum theory, which models the rotor as an infinitesimally thin disk with a discontinuous pressure rise, and neglects viscosity, compressibility, and wake rotation [11].

Conservation of mass and momentum gives a fully developed slipstream velocity of exactly $2v_i$ in the ideal case [11]. Since v_i squared is proportional to the disk loading T/A, the downwash velocity grows with thrust. In hover, the thrust balances the vehicle weight, so a heavier vehicle, or a manoeuvre that demands thrust above the hover value, produces a faster wake. On the other hand, a standard quadrotor is comprised of four rotors distributed symmetrically; thereby, in the near-field region immediately below the vehicle, each rotor generates a distinct slipstream, and as the vertical distance increases, these four distinct streams diffuse and merge as demonstrated by Bauersfeld et al. [1]. In particular, it has been empirically evaluated that at vertical distances exceedingly approximately 2.5 times the vehicle diameter, the merged flow is accurately modelled as a single, fully developed turbulent jet that exhibits a mean axial velocity inversely proportional to the vertical distance from the source and a radial velocity profile at any constant depth that approximates a Gaussian distribution, peaking at the central axis and decaying laterally; the variance of this profile increases with depth due to jet expansion. By normalizing the flow field against a characteristic exit velocity and the motor spacing l, Bauersfeld et al. [1] show that the mean velocity profiles of differently sized quadrotors collapse onto a universal self-similar curve in the far field, once the rotor jets have merged. This scaling provides a principled way to transfer their velocity model to the scale of the Parrot Mambo, although two caveats apply: the collapse holds in the merged far field, whereas close-formation flight also samples the near field, and the model describes the jet velocity rather than the aerodynamic moment sensed by the follower IMU. Both points are addressed in the following sections. The turbulent-jet description of the downwash gives the flow field a definite spatial structure: the velocity is largest on the vertical axis aligned with the leader and decreases with radial offset, and it also weakens as the vertical separation grows. Each relative position therefore maps to a distinct value of the flow, so the field varies in a repeatable way with the offset (r, z) between follower and leader. A follower placed inside this field is subjected to an airflow whose magnitude and direction depend on where it sits relative to the leader. In this sense the wake encodes the relative position variables, and the aerodynamic disturbance carries information about (r, z). The exploitation of aerodynamic interactions is a well-documented phenomenon in biological systems: for example, avian species frequently leverage it since a flapping wing generates trailing tip vortices that induce a localized downwash immediately aft of the wing and a compensatory upwash outboard of the wingtips. Moreover, Lissaman and Shollenberger [12] theorized that birds flying in a V-formation shape exploit this upwash to reduce the required power output. Subsequent empirical studies validated this hypothesis: Weimerskirch et al. [13], for example, recorded substantial reductions in the heart rates of pelicans flying in formation, while Portugal et al. [14] demonstrated that northern bald ibises actively regulate their spatial positioning and wingbeat phasing

to maximize upwash interception and avoid downwash zones. Consequently, this thesis aims to replicate this adaptive behaviour mathematically by substituting biological mechanoreceptors and vision with inertial telemetry. Two boundary conditions frame the fluid model: first, the turbulent-jet approximation holds in the far field, from about 2.5 motor-to-motor distances below the rotor plane, where the four rotor slipstreams have merged into a single jet. Closer to the leader, in the near field, the individual slipstreams are still separated, and the flow is more complex, which places this region outside the scope of the estimator. Second, the leader wake is only one of several possible aerodynamic effects. Ground effect near the floor and recirculation of the follower's own wake in confined spaces add further attitude perturbations. This work isolates the leader-follower interaction and operates in conditions chosen so that the leader downwash is the dominant external disturbance, flying away from the ground and in an open area to limit these other effects. An ambient baseline test performed before the flights confirms a low-turbulence environment and sets the noise floor against which wake events are detected.

1.3 From Disturbance to Measurement: Interpreting Wake-Induced Signals

Turning a wake interaction into a usable measurement requires a clear separation between two physical quantities, since conflating them would misplace the region of highest sensitivity and compromise the state estimator: the first quantity is the velocity of the air in the wake, while the second is the force and moment that this air exerts on the follower. These two quantities do not reach their maximum at the same place. As highlighted by Bauersfeld et al. [1], the mean axial velocity $\langle U \rangle$ is a function of the vertical distance z and the radial offset r: the velocity is largest on the axis directly below the leader, at $r = 0$, and decreases with radial offset. A follower carrying an airspeed probe would therefore read its strongest signal on the axis. However, the follower studied here carries no airspeed probe, since it exceeds the scope of this work, and relies solely on its Inertial Measurement Unit, which measures specific force and angular rate, and the onboard Extended Kalman Filter (EKF) turns these into an attitude estimate. The quantity that reaches the follower is therefore the aerodynamic moment produced by the wake, not the wake velocity. The moment follows from the radial gradient of the jet introduced in Section 1.2. When the follower sits on the axis, directly below the leader, the inflow across its rotor plane is symmetric, the thrust loss is the same on both sides, and the net moment is zero. As the follower moves off the axis, this symmetry breaks. The rotors on the side nearer the leader meet a faster inflow than those on the far side, the thrust reduction becomes unequal across the frame, and this asymmetry,

Table 1.2: The two fluid-dynamic views of the wake. Because the follower relies on inertial sensing, the estimator is built on the moment model, which places the region of highest sensitivity off the axis and a null on the axis.

	Velocity field	Moment field
Primary variable	Mean axial velocity $\langle U \rangle(r, z)$	Induced force and aerodynamic moment
Source	Bauersfeld et al. [1]	Kiran et al. [15]
Sensing	Airspeed probe	Inertial Measurement Unit
Spatial maximum	On-axis ($r = 0$)	Annular region ($r \approx r_{\text{peak}}$)
On-axis response	Maximum	Approximately zero

acting through the follower half-wheelbase, produces a pitching or rolling moment. This behaviour is confirmed by the force-torque measurements of Kiran et al. [15] on a quadrotor pair: starting from the stacked configuration and increasing the horizontal offset, the induced pitch moment grows, peaks at the offset of maximum moment arm, and then decays to zero as the follower leaves the wake; hence, the moment is null on the axis and reaches its maximum in a ring around it. The offset at which this maximum occurs depends on the wake geometry and the follower size: for the Parrot Minidrone Mambo it is estimated directly from the flight residuals in Chapter 4 and is denoted r_{peak} in what follows. This is the distinction that shapes the estimator. The velocity field and the moment field have inverted spatial features. Where the velocity is largest, on the axis, the moment is zero, and the moment is largest in the annular region around the axis. An estimator that relies on the IMU must therefore be built on the moment model, and must expect its strongest signal off the axis, at r_{peak}, rather than on it. A comparison of the two views is given in Table 1.2.

The operational implications of this distinction are critical. A control law synthesized under the assumption that the inertial response correlates linearly with the local fluid velocity would actively drive the follower toward the central axis ($r = 0$). Upon converging to this axis, the induced aerodynamic moment vanishes, leading to a complete loss of the sensory signal and the subsequent destabilization of the closed-loop system. By explicitly formulating the state estimator based on the aerodynamic moment topology, the control strategy ensures that the follower is regulated along an annular region of maximum sensitivity, deliberately avoiding the central null zone. Following the identification of the appropriate physical variable, the signal processing architecture can be established. The induced aerodynamic moment dynamically perturbs the follower's attitude, manifesting as an exogenous deviation within the Extended Kalman Filter (EKF) state estimate. However, the raw attitude telemetry conflates this aerodynamic disturbance with the deliberate corrective efforts generated by the flight controller.

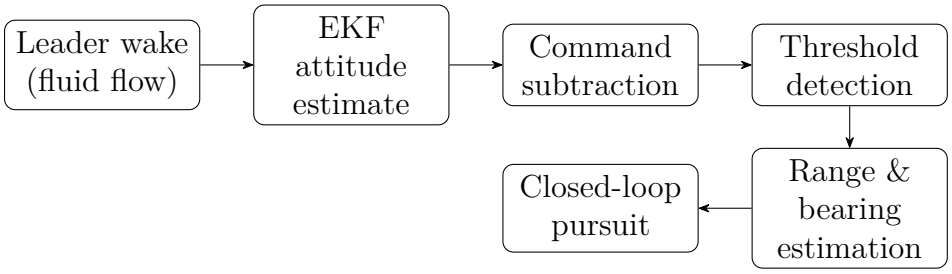

Figure 1.1: The implicit sensing pipeline. The aerodynamic interaction appears as uncommanded attitude perturbations. Subtracting the controller setpoints isolates the disturbance, whose peaks are processed into relative range and bearing and passed to the pursuit loop.

To isolate the purely aerodynamic signature, the commanded reference attitude setpoints are subtracted from the estimated states, yielding the attitude tracking residuals:

$$\Delta\theta = \hat{\theta} - \theta_{\text{cmd}}, \tag{1.2}$$

$$\Delta\phi = \hat{\phi} - \phi_{\text{cmd}}, \tag{1.3}$$

where $\hat{\theta}$ and $\hat{\phi}$ are the EKF pitch and roll estimates and θ_{cmd}, ϕ_{cmd} are the setpoints issued by the trajectory controller. The direction of the resulting moment is tied to the bearing of the leader. A leader directly ahead induces a pitching moment, a leader to the side induces a rolling moment, and intermediate angles give a coupled response. The relative bearing can therefore be recovered from the two-argument arctangent of the residuals, which supplies the angle needed to close the polar pursuit loop. The full signal-extraction chain is shown in Figure 1.1.

1.4 The Baseline One-Way Leader-to-Follower System

Wake-based leader sensing builds on earlier work on the same hardware. Cerruti [5] implemented a first proof-of-concept system, showing that a resource-constrained UAV can track a leading vehicle using only the pitch-axis perturbation induced by the leader downwash.

The baseline uses a discrete, state-based pursuit strategy. The follower moves forward at a constant reference velocity of about $0.12\,\mathrm{m\,s^{-1}}$. When it enters the leader wake, the induced pitching moment drives the pitch residual (Equation 1.2) out of a fixed interval, and this crossing is registered as a detection event. Cerruti [5] first set this interval to $(-0.06, 0.1)$ radians and later moved the lower bound

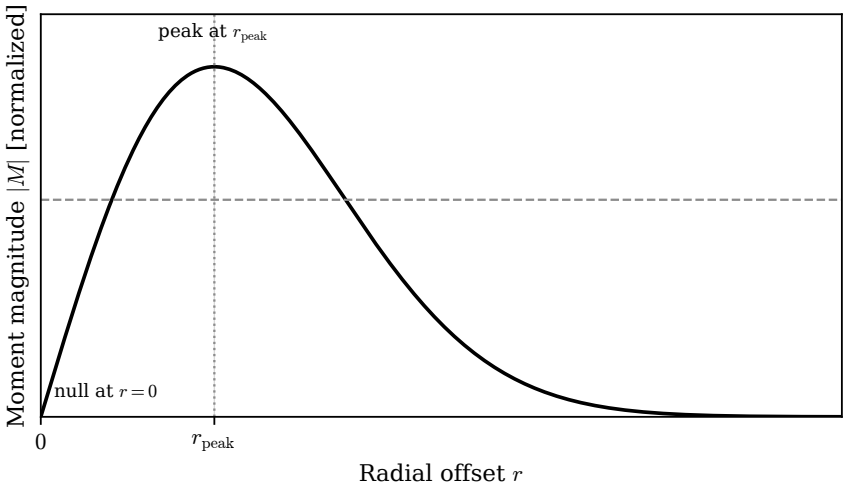

Figure 1.2: Non-monotonic relation between radial offset and induced moment magnitude. A given amplitude corresponds to two offsets, one on each side of the peak, so the inversion is not unique. To obtain a single-valued mapping, the numerical inversion is restricted to the monotone outer branch beyond r_{peak}. The peak location r_{peak} is determined empirically from the flight residuals.

to -0.08 radians, because a small oscillation follows each detection and its peaks could fall below -0.06 and trigger false events. After a detection the follower records its position and time, brakes, and starts a new approach to produce the next detection. The leader scalar velocity is then estimated by finite differences between two successive detections:

$$S = \frac{X_{\text{present}} - X_{\text{past}}}{t_{\text{present}} - t_{\text{past}}}, \tag{1.4}$$

where X and t are the position and time at each detection. The follower then adopts S as its own reference velocity to match the leader pace. A finite-state machine runs this loop, and a timeout forces a new leader search when detections stop, which protects the tracker against continuous leader acceleration.

The same system also carries a simple one-way command channel. Since the rate of detection events depends on the leader motion, the leader can modulate its own velocity to send a small set of discrete commands, such as land, spin and vertical oscillation. This gives an implicit command interface that still works when explicit radio communication is unavailable.

Testing of the baseline gave clear data on where detections occur in space. With a safety stop triggered at the first detection, the point of intersection was found to vary widely, from directly below the leader at $r = 0$ up to about 60 cm ahead,

depending on the relative approach velocity. In steady tracking the inter-vehicle separation settled in a wide band, between 20 cm and 80 cm behind the leader, and the motion split into a repeating four-phase cycle of approach, detection, braking and re-approach.

These results show two things about the baseline. First, they confirm that an IMU-only follower can perceive a preceding aircraft through aerodynamic coupling. Second, the architecture works as a discrete, iterative tracker, not as a continuous regulator of relative position. The channel is also one-way. Moving from iterative tracking to continuous regulation, and adding a follower-to-leader feedback path, is the objective of this thesis and of the paired work on the return channel.

1.5 Open Issues in the Baseline System: Blind Zones, Range Ambiguity and Robustness

The baseline tracks a leader, but it does so through repeated approaches rather than by holding a fixed relative position. The reasons it cannot settle at a defined setpoint follow from the structure of the baseline estimator.

The first limit is an aerodynamic blind zone. Detection depends on the follower moving through the spatial gradient of the wake, so it needs a non-zero relative velocity. If the follower slows down to match the leader, the differential flow fades, the detection signal sinks into the noise floor, and the state estimate is lost. When Cerruti [5] tried a continuous Proportional-Derivative controller on the raw pitch residual, the loop showed diverging oscillations and held position only when the leader stayed still, with the gains reduced so far that the follower barely moved. The moment model of Section 1.3 explains this. The induced moment tends to zero as the offset r tends to zero, so a controller that drives the tracking error to zero drives its own sensing signal into the on-axis null and loses the leader.

The second limit is range ambiguity. The residual magnitude depends on both relative distance and relative velocity. As shown in Figure 1.2, the moment magnitude is also not monotonic in radial offset. It rises from zero on the axis to a maximum at r_{peak} and then decays. A given magnitude therefore matches two offsets, one inside and one outside r_{peak}, so a direct inversion from amplitude to range is ill-posed unless it is restricted to the monotone outer branch. A controller that reads only a falling amplitude cannot tell whether the follower is receding safely or closing toward the on-axis null.

The third limit is robustness of the detection logic. The fixed threshold is sensitive to parameter drift, in particular to battery voltage, which changes the response of the flight controller and shifts the tuned threshold. Uncommanded tracking errors during fast lateral manoeuvres also produce residuals with a spectral signature close to that of a real wake encounter, which causes false positives. A static amplitude

threshold cannot separate these two cases on magnitude alone. The baseline state machine also chatters when the residual sits near a threshold and crosses it several times within one execution cycle.
Finally, the baseline returns only a scalar velocity estimate. It does not reconstruct the full relative state, that is both range and bearing, so it cannot support a two-dimensional pursuit law. Removing these limits calls for an estimator that maps the biaxial IMU residuals onto the physical wake model and returns a continuous estimate of the relative range and bearing.

Chapter 2

State of the Art on Wake-Based Sensing and Relative Estimation

2.1 Aerodynamic interaction models for multirotor UAVs

During the flight, the multirotor Unmanned Aerial Vehicle behaves not as an isolated rigid body but as fluid-dynamically-coupled system because the propulsion system accelerates a column of air downwards generating a flow field that interacts with other entities including the surrounding vehicles and even fixed boundaries. If they are in close formations, the aerodynamic downwash that is generated by the upper (Leader) vehicle reaches the lower (Follower) vehicle, perturbing its flight dynamics, with maximum effect when the vehicles are vertically aligned with the perturbations being of sufficient magnitude to destabilize the trailing vehicle and hinder its tracking performance [16]. A rigorous mathematical modeling of the interaction is a fundamental requirement in dense formations as a reactive controller cannot reject such disturbances without first characterizing it [16]. This section reviews the existing models of the aerodynamic interaction classifying them in three main categories: physical force-based models, learning-based paradigms and direct characterizations of the flow field. It also examines similar models of boundary-induced aerodynamic effects as they share the same fluid dynamic principles.

The first category consists of physical, force-based models. Gielis et al. [16] devised an experimental apparatus to directly characterize the six-degree-of-freedom exogenous forces, which act upon a multirotor's airframe. In doing so, they

characterize these forces as a function of the relative spatial state between vehicles executing different flight trajectories, thereby extending single-vehicle interaction paradigms to the aggregate aerodynamic forces felt by an ego-vehicle within a multi-agent system. Similarly, Jain et al. [17] formulate an aerodynamic disturbance model for multirotors in proximity flight, which they calibrate with onboard empirical measurements. In previous foundational work, Yeo et al. [18] derived an empirical model of the rotorcraft's downwash, tailored for disturbance localization and avoidance, an objective that is very similar to the one of the wake, by using a methodology advanced in this thesis. However, the limitation of many of these works within this category, as pointed out by Gielis et al. [16], is to report results in dimensional units specific for a particular hardware configuration, which severely limits their applicability across different UAVs.

The second class consists of learning-based models, mostly developed to address the shortcomings of the first category. Mathematical models proved insufficient in capturing the three-dimensional structure with the required accuracy, especially for a dense formation, and alternative approaches had to be explored. In particular, the authors of Smith et al. [2] identify the downwash topologies directly from flight data. They predict the downwash forces by means of a machine learning algorithm, which was trained on a limited amount of data, while accounting for rotational symmetry of the system. Their results demonstrate that a geometry-aware neural network that incorporates limited amounts of empirical information outperforms the standard model trained on a significantly broader dataset. Similar results are reported by Shankar et al. [19], who utilize a learned downwash model to design a feedback loop for the vertical docking of two multirotors. Particularly, in the leader-follower formation, the follower vehicle experiences stronger disturbances induced by the leading aircraft during the final stage of approach, and the utilization of a learned feedforward loop reduces the necessary distance between the vehicles. Learning-based methods also underlie the heterogeneous swarm interaction models proposed by Shi et al. [3], which predict inter-vehicle aerodynamic forces between dissimilar UAV platforms, to augment the overall control authority.

The third category focuses on characterizing the fluid flow field instead of estimating the forces directly. Thus, the induced flow of air underneath a quadrotor has been characterized as a fully developed turbulent jet [1]. These formulations provide a compact, parametric description of the mean fluid velocity as a function of space, from which downwash models for both high-fidelity simulations and closed-loop control synthesis can be synthesized. This view provides the most prevalent paradigm used throughout this thesis, and a detailed exposition of the model, including the scaling laws to generalize it to different dimensioned vehicles, is presented in Section 2.2.

A related corpus of the literature addresses the aerodynamic interaction of a single vehicle with ground or other surfaces. This is an important consideration, as

ground and wall effects are significant aerodynamic disturbances, which need to be separated from the influence of the leader's wake. The classical models of near-boundary flight have been developed by applying the method of the mirror image to the rotorcraft aerodynamics, resulting in a reduction of the induced velocity at the plane of rotation and, consequently, an increase in thrust near the ground [20]. Carter et al. [21] tested the applicability of such large-scale models to the micro-quadrotors by measuring, with particle image velocimetry, the flows around the quadrotor near the ground, ceiling, and sidewalls. The experiments confirmed that the lift increase dominates at the ground and ceiling while a minor lift decrease is present near the sidewalls, thus validating the assumptions of the classical models. The flow singularities observed in the case of the ceiling were explained by the momentum theory developed for rotorcraft [22] while the similar effects were found for the micro-quadrotor near the ground, ceiling, and sidewalls [21]. Thus, the aerodynamic interaction of a vehicle with nearby surfaces can be considered as a manifestation of the same effect, modifying the inflow of the rotor according to the motion of the neighboring boundary. This phenomenon has been shown to affect both the leader-follower formations and the vehicles operating close to the ground.

In all three cases, the aerodynamic interaction was studied to characterize the coupling of a rotor wake with a nearby vehicle or a surface. In all cases, this interaction was treated as a disturbance, which needs to be estimated and compensated for by the flight controller. However, the present thesis pursues a different goal, using this interaction as an information channel. Specifically, the disturbances induced into the follower by the leader's wake have been used to explicitly estimate the relative position of the follower with respect to the leader.

2.2 Turbulent jet models for rotor downwash characterization

While the majority of the previous section concerns methods quantifying the aerodynamic forces exchanged between vehicles, using the wake of a leader for relative state estimation purposes requires a detailed description of the fluid flow field itself. Below a quadrotor, after the combination of the slip streams from the individual rotors, the resulting downwash is well approximated by classical turbulent round jet theory, allowing to directly apply the fluid dynamic properties of this canonical structure to the rotorcraft case. In this section, the theoretical foundations of the turbulent round jet are reviewed, along with its application to multirotor downwash topologies.

The free turbulent round jet is the flow which is self-sustained, i.e., there are no walls restricting the jet motion, which makes its statistical properties defined solely

by the local flow scales and not by the boundary conditions [23]. The free jet possesses two key characteristics. First, the mean axial velocity at the jet centerline decays as the inverse distance from an origin several nozzle diameters upstream, and the jet boundary is expanding with a nominal constant half-angle [23]. Second, the radial profile of the axial velocity component is self-similar at any downstream location. Specifically, the velocity profiles collapse when the local axial velocity is non-dimensionalized by the centreline value, and the radial coordinate is scaled by the half-width of the jet. In this way, the transverse velocity profiles can be well approximated at different depths by a single non-dimensional curve, the Gaussian distribution [24]. That property is essential for the mathematical description of the jet because it allows to parameterize the entire profile of the flow by a single self-similar curve. In particular, the use of the Gaussian distribution is justified by assuming fluid dynamic similarity of the axial profile, flow axisymmetry, and domination of the radial velocity gradients over the axial ones [25].

Applying this theoretical framework to a quadrotor platform requires consideration of additional geometric effects due to the system being composed of four independent rotors rather than a single nozzle. Within the near-field region, each of the rotors produces a unique slipstream, making the macroscopic turbulent-jet approximation unsuitable. However, the flows do merge to produce a single round jet, once the individual slip-streams have fully coalesced. Empirical results using Particle Image Velocimetry (PIV) indicate that the wake of the quadrotor only becomes self-similar at the merged far-field region, specifically at a distance greater than $z/l > 6.5$ vertically below the platform, where l is the diagonal motor-to-motor baseline [26]. As such, the single jet approximation is not suitable to model the aerodynamic near-field interactions occurring below this threshold.

A highly comprehensive formulation of this multirotor flow field is provided by Bauersfeld et al. [1]. By extensive empirical measurement of the induced downwash of various free-flying quadrotors, they show that when normalised by a characteristic exit velocity and inter-rotor spacing l, the mean velocity profiles of dimensionally inconsistent systems collapse onto a self-similar curve, when measured within the far-field regime. This dimensional collapse is highly beneficial, as it allows aerodynamic models to be extrapolated across different hardware platforms, utilising only the mass, characteristic length and propeller diameter of the vehicle. The fit parameters of this empirical formulation have subsequently been validated by independent studies. Specifically, Kiran et al. [26] report far-field kinematics that are consistent to those found by Bauersfeld et al. [1], recording a centreline decay constant of approximately 9.5 (vs 10.1) and a spatial spreading rate of 0.067 (vs 0.077). This agreement supports the generality of the turbulent-jet approximation across different aerial vehicles and experimental setups. The primary value of this formulation is that it allows us to capture highly complex 3D fluid dynamics within a simplified parametric model. The combination of the centreline decay law, which

determines the axial velocity as a function of downwash, and the self-similarity of the radial distribution allows us to compute the velocity of the fluid at any given offset. Modern simulation software explicitly takes advantage of this efficiency. For example, a particle-based downwash simulator built for the purpose of high-speed, close-proximity UAV racing places particles beneath each rotor using momentum theory; these particles then move and decay strictly according to turbulent-jet kinematics, allowing us to extract the local airspeed by way of spatial interpolation [27]. This highlights the fact that turbulent-jet kinematics are light enough to be used inside real-time, multiple-agent control loops.
However, despite the model's accuracy, there are two significant limitations to using this model in the scope of this thesis. First, the model captures only the time-averaged mean flow field. That is, the steady-state downwash perceived by a following vehicle, while not capturing the high-frequency turbulent fluctuations over the mean. Second, the model is strictly valid only in the far-field, fully merged region. As the follower approaches the leader in an ultra-close formation, it may actually enter the near-field region where the independent rotor slipstreams have not yet merged in a highly non-linear manner. These two limitations together define the spatial boundary within which the aerodynamic sensing approach proposed in this thesis is applicable. Finally, these boundaries are discussed in more detail in Chapter 3 where the channel model for the geometry evaluated in this thesis is derived.

2.3 Downwash-aware control and compensation strategies

The fluid-dynamic and force-based models described in the previous subsections are primarily driven by the need for closed-loop control. By defining the interaction between vehicles in-flight with a mathematical model, control strategies can be engineered which compensate for these effects, rather than ignoring them as unknown disturbances. The purpose of the reviewed methods is to reduce the distance required for stable flight, which would otherwise be conservatively large for disturbance-blind controllers, by estimating, predicting and compensating for the influence of the wake. The compensation strategies reviewed generally differ through the manner of disturbance estimation and control input injection.
The most basic mitigation approach is large standoff distances, such that controllers without an explicit aerodynamic model must maintain large separation distances to keep the perturbations within the bounds of the baseline controller's stability margins, restricting tightly spaced formations for any but the most robust controllers [3]. All following approaches attempt to mitigate this by reducing the required standoff distance.

A primary class of compensation strategies is based on estimating and feedforward injection of the disturbances directly estimated from the inertial telemetry: in such approaches indeed, aerodynamic disturbances are analytically reconstructed from the IMU state estimates. Hence the estimated disturbance profile resulting from an initial trajectory can be feedforward injected during the next iterations to compensate for the disturbance experienced, especially beneficial when dealing with repetitive trajectory tracking [17]. Such strategies are incredibly advantageous in that they do not require additional exteroceptive sensors nor inter-vehicle communication, only relying on the inertial signature of the disturbance in the wake. This philosophy is very much reminiscent of the sensing paradigm proposed in this thesis, with one fundamental difference: while the former approach uses the estimate to actively cancel the disturbance, we use it to localize the disturbance in space.

A second growing class of techniques focuses on learning downwash models within the control loop, allowing a closed-loop system to incorporate learned aerodynamic effects and track desired trajectories while maintaining stability. Neural-Swarm and derivatives thereof utilize permutation-invariant neural networks to approximate the multi-vehicle aerodynamic interaction matrices, which are then combined with stable trajectory tracking controllers. The resulting algorithm demonstrated a four-times lower worst-case altitude tracking error compared to a disturbance-agnostic baseline [3]. Similarly, Smith et al. [2] learned the downwash interaction forces, enforcing rotational symmetry during the learning process. The results showed that a physics-informed model, trained on five minutes of flight data, surpassed the performance of a generic model trained on a dataset three times larger by their own metrics. Gielis et al. [16] incorporated a learned aggregation of multirotor downwash into the Nonlinear Model Predictive Controller (NMPC) to demonstrate improved trajectory tracking performance. The results showed a 40 percent decrease in global tracking error and a 57 percent decrease in maximum altitude error with a training set containing only 46 seconds of data. Shankar et al. [19] used a similar approach, applying the learned aerodynamic model within an optimal feedback controller to perform vertical docking of two quadrotors. The learning process decreased the follower drone's approach tracking error to below 0.06 metres, a three-to-four-times improvement, enabling the docking manouever to take place in the first place.

A third category avoids explicit modeling of parametric downwash and instead employs online adaptive control and reinforcement learning techniques. For example, a residual reinforcement-learning module can be added to a conventional cascaded flight controller; hence this architecture autonomously generates compensatory control inputs to counteract thrust degradation and aerodynamic moments from nearby vehicles, based purely on the ego-vehicle's state and baseline control commands without requiring inter-vehicle communication [28]. Classical adaptive control theory has also been applied to this domain, employing online adaptation

mechanisms to compensate for unmodeled aerodynamic disturbances and successfully maintaining strict vertical alignment of a three-vehicle stack during dynamic flight manoeuvres [29]. These approaches sacrifice analytical insight from an explicit fluid-dynamic model in favor of operational flexibility to counteract stochastic and unmodeled perturbations in real-time.
Two critical conceptual intersections relate this existing literature to the core thesis of this work: first, virtually all contemporary methodologies, be they feedforward, learning-based, or adaptive, consider the aerodynamic wake as an adverse error signal to be actively attenuated, allowing the follower to safely track an independent reference trajectory while ignoring the physical presence of the leader. The thesis proposes that this same aerodynamic signal can instead be decoded as a direct, continuous measurement of the leader's relative spatial coordinates. Secondly, the compensation strategies which most closely parallel the proposed architecture are those which rely exclusively on endogenous inertial telemetry, without the use of additional sensors or radio-frequency exchanges [17]. These prior works have empirically validated the core hypothesis that wake-induced aerodynamic disturbances are observably present through proprioception. The theoretical leap taken by this research is to transform from considering this observable as an exogenous rejection input to its being the primary physical carrier of relative positional information.

2.4 Relative localization methods in GPS-denied and communication-limited settings

The effectiveness of any formation control structure fundamentally relies on an accurate estimation of neighboring agents' relative state. Although Global Navigation Satellite Systems (GNSS) are widely used for positioning in open-sky environments, close proximity UAV operations are often carried out indoors or require inspecting complex infrastructure such as tall buildings where satellite positioning is not available due to signal attenuation or is otherwise inaccurate. Thus, the problem setting changes to require the relative positioning of formations and their members using only onboard resources. This section reviews the relevant non GNSS relative positioning techniques, focusing on those that employ different sensors, and provides motivation by comparing the proposed wake sensing method to them.
The most common types of exteroceptive localization methods are based on some kind of exteroceptive infrastructure, such as optical motion-capture systems, which allow millimetre-level precision in positioning but are severely limited by the working volume of the system. Distributed Ultra Wide Band beacons can be placed strategically in the environment to provide positioning information with decimetre-level precision over a much larger area than motion capture systems [30]. However,

both systems are characterized by a dependence on stationary infrastructure and a specific coordinate frame, which makes these approaches unsuitable both for the localization of unexplored environments and for autonomous decentralized multi-robot systems [31].

To bypass infrastructural dependencies, a primary class of onboard methodologies utilizes inter-agent radio-frequency ranging: by providing each UAV with a UWB transceiver, multi-agent systems are capable of estimating relative spatial configurations based on time-of-flight distance measurements; thereby constituting a lightweight and low power architecture particularly conducive to swarm scalability (Liu et al., 2024). Scalar range measurements offer omnidirectional spatial awareness with minimal payload or computational overhead, thus reserving higher bandwidth optical sensors for mission specific purposes as opposed to inter-agent tracking (Nguyen et al., 2021). Nonetheless, the fundamental shortcoming of scalar ranging lies in weak state observability, wherein full relative pose reconstruction requires sufficient dynamic excitation of the formation geometry in order to allow mathematical triangulation of positions. Furthermore, the asymptotic accuracy of such estimators is bound by the stochastic noise of the RF hardware, which in practice results in planar localization errors of tens of centimetres for convoy implementations (Chen et al., 2022).

A second major class of onboard localization makes use of computer vision: electro-optical sensors that track the position of nearby vehicles by virtue of fiducial markers, deep learning-based object detectors, or collaborative visual-inertial odometry in a common frame of reference [31]. Vision-based approaches provide rich and expressive state information in the form of high-dimensional per-pixel feature descriptors. They do not require any RF emissions, but rather suffer from the limitations of requiring a direct line of sight and sufficient illumination, as well as demanding image processing pipelines. Particularly concerning for low-cost and small-sized drones is the heavy processing load these approaches impose on the onboard microcontroller, which can introduce unacceptable delay and sluggishness in the state update rate [9]. The costs associated with weight, power and computational budget for such approaches are prohibitively large, especially for nano drones, as discussed in Section 1.1

A third, more restrictive category includes proprioceptive and implicit sensing techniques that eliminate the need for dedicated exteroceptive sensors for inter-agent tracking. For example, Helm et al. [7] showed the localization of a leading UAV by a follower in an indoor leader-follower setup using only scalar inter-vehicle ranging combined with the follower's own estimate of its inertial state. They show that stable formation flight can be achieved mathematically in the absence of GNSS and without external tracking infrastructure. This category is the closest philosophical predecessor to this work, as it actively reduces the reliance on exteroceptive hardware and completely decentralizes the relative estimation

problem.
The aerodynamic wake-based sensing approach introduced in this thesis represents the culmination of this technology chain. Subsequent localization approaches progressively strip away any reliance on a surrounding infrastructure, first external infrastructure, followed by dedicated RF ranging hardware, and lastly even electro-optical hardware. The forward channel implicitly analyzed in this work completes this process by only utilizing the unavoidable aerodynamic downwash exerted onto the follower's fuselage to determine the relative position of the leader. As such it constitutes the most infrastructure- and hardware-lightweight approach to date. This comes at the price of a limited maximum range and a complicated, non-linear sensing topology dictated by fluid mechanics which will be detailed and mathematically formalized in the following chapters.

2.5 Signal detection methods for weak aerodynamic signatures

The identification of a leading vehicle based on the aerodynamic wake is a classic signal detection problem; the trailing UAV must autonomously decide whether there is a leader wake present in the field of view given the noisy onboard telemetry, and then extract the relative state information. The signal of interest is the wake, which is significantly attenuated by the sensor noise, active controls, and the dynamics of the ego vehicle. This section reviews signal detection methods that are relevant to low SNR (Signal-to-Noise Ratio) and covers a burgeoning literature that adapts the signal detection literature to aerodynamic signatures.
The theory of detecting weak signals in the background of noise is quite developed. The basis for this was laid down in the fields of radar, sonar, and telecommunications. There are two main approaches to signal detection, the matched filter and the energy detector. The former is used to maximize the signal-to-noise ratio (SNR) at the receiver output by comparing the input signal with a known stored version of the signal. It is optimal when the signal structure is known and, therefore, constitutes the best linear filter [32]. The energy detector, in turn, accumulates the energy of the input signal over a certain time window and compares it with a threshold. This method does not require knowledge of the signal structure but is more sensitive to wideband noise [33]. However, in both cases, the essence of the detection process is reduced to determining a test statistic that exceeds the specified threshold value.
A static detection threshold is easy to compute, but can be very fragile in practice, since it implicitly assumes the noise variance to be both stationary and known with sufficient accuracy. In a varying noise environment, this type of threshold either produces too many false alarms or, more importantly, too many missed

detections [33]. The standard solution to this problem is the Constant False Alarm Rate (CFAR) detector: it computes the threshold on a per-sample basis, based on an estimate of the noise floor in a local neighborhood around the sample, and thus maintains a statistically stable false alarm rate regardless of noise variations: specifically, in the classical cell-averaging case, it uses the mean of neighboring noise-only samples as a reference [34]. This approach is widely used in many different contexts; for instance, it is applied to time-of-arrival estimation in Ultra-Wideband (UWB) receivers using an energy detection front-end [34], and in spectrum sensing for cognitive radio, where the threshold is dynamically computed from neighboring frequency bins [35]. One common variation on this theme is to iteratively estimate the noise distribution by removing signal-containing outliers, until the false-alarm rate is within acceptable limits, even in the presence of multiple signal samples in the noise estimate [36].

A fundamental principle from the cited literature has direct application to our proposed architecture: the threshold of any threshold-based detector is fundamentally limited by the SNR and the noise floor estimation uncertainty [33]. To achieve maximal robustness in such systems it is therefore crucial to dynamically bind the threshold to the measured value of the noise floor instead of using a hard-coded design-time value. The noise characteristics of the multirotor UAV used in the present study are undeniably non-stationary: large variations in the controller output and sensor noise are observed as the battery voltage decreases; a dynamic threshold tracking the noise floor is therefore the most viable option, which was formally developed in Chapter 4.

The application of these methods to aerodynamic signatures of interest remains largely unexplored. As discussed in Section 2.3, most of the literature on multirotor downwash focuses on continuously estimating the disturbance for the purpose of rejecting it actively, not detecting distinct events. The closest work in this domain is the reconstruction of the downwash effects from the proprioceptive inertial data of the observer [17]. This paper, however, tackles a slightly different problem of continuous state estimation rather than a statistical hypothesis testing one. One of the few papers explicitly mentioning the information contained in wakes and needing to process it is the initial version of the baseline system by Cerruti [5]. The original design used a simple threshold on the follower's pitch residual to detect occurrences of wake following — with this threshold being a hard-coded geometric bound. The initial design suffered from the same issues of fragility caused by the use of such a static bound. This thesis provides an improvement over the baseline by using a more expressive thresholding function informed by an online noise-floor estimation, inspired by the CFAR detection principles reviewed here in this section.

2.6 Positioning of this thesis with respect to existing work

The previous four sections present research that stems from a common paradigm: the multirotor downwash is treated as an undesired disturbance. Accordingly, the central research aim in this space was to understand and mitigate the effect of rotor-induced flows on the target vehicle in order to allow for closer proximity operations. The interaction models in Section 2.1, the jet fluid mechanics in Section 2.2, and the control methods from Section 2.3 all contribute towards disturbance rejection. On the other hand, the localization approaches in Section 2.4 tackle the problem from a different direction — the need for relative position information is met with solutions that rely on exteroceptive sensors such as satellite positioning receivers, RF range finders, or electro-optic cameras. This thesis explores concepts in both of these areas. The current work is situated at the intersection of vehicle dynamics and localization; this fact becomes evident in the following section.

The main difference between the proposed architecture and the previous works is the exploitation of the downwash phenomenon. In the control literature, the wake is considered to be an error signal that must be attenuated, and most of the methods dedicated to the estimation of these perturbations rely on using onboard inertial telemetry to reconstruct them [17]. Similarly to [5], the inertial signature is analyzed, yet, instead of synthesizing a compensation signal, the perturbation is decoded to localize the leading vehicle. In other words, the physical disturbance is directly used as a relative state measurement. In the author's opinion, the concept of transforming aerodynamic disturbances into a physical equivalent of a position/orientation relative state measurement (i.e., relative position/orientation) has not been previously explored in the literature. It forms the core idea of the first baseline architecture [5], which is furthered by this thesis.

The secondary defining feature of this research is that it adheres to the principle of sensor parsimony. Communication-limited estimation and control have been a hallmark of infrastructure-free swarm robotics. One representative example of such an approach is the compact-swarm formation by Petráček et al. [37], which employs a decentralized control policy that does not require communication or external localization cues. Instead, the individuals communicate their state and cooperate via vision-based perception of a single UV active marker and a camera. While similar in spirit to the bio-inspired approach of the current paper, the method of Petracek et al. requires additional positioning hardware, namely, the UV marker and camera, for relative localization. By contrast, the implicit aerodynamic forward channel utilized in this work does not require any computer vision or communication hardware. The leader's position is estimated by the follower's flight control IMU. From the perspective of infrastructure parsimony, the current approach is even

more infrastructure-free than the communication-implicit formations of Petráček et al. [37] because it only utilizes the standard autopilot IMU of the follower drone for implicit aerodynamic sensing.
The trade-off to pay for such an extreme hardware reduction is a reduced range of applicability; while RF ranging and computer vision yield a 3D-position estimate that is valid regardless of the distance to the target, aerodynamic wake interactions are inherently restricted to short ranges due to fluid mechanics. As a result, the wake-based position sensing channel comes with an inevitably large central blind spot and a limited range of operation. Hence, this work does not suggest to abandon conventional position sensing approaches in favor of wake vortex position sensing. Instead, it is proposed as a supplementary method that has its main applicability in the close range where standard position sensing approaches fail due to both aerodynamic coupling and reduced observability. The processing of the signal data described in Sect. 2.5 plays an important role in this context. Its careful adaption to the characteristics of aerodynamic wake vortexes is the major technical contribution of this work and the following chapters.
This work therefore resides in a particularly specialized and until now unexploited portion of the UAV formation design space, as it recasts aerodynamic downwash as an information bearing signal rather than a parasitic disturbance, wholly eliminates the need for dedicated inter-agent sensing hardware and formally accounts for the resulting geometric constraints. The following chapters mathematically formalize the model of this aerodynamic channel, derive the detection and state-estimation pipelines required to decode it and synthesize the closed-loop pursuit logic necessary to exploit it.

Chapter 3

Modeling the Leader-to-Follower Aerodynamic Channel

3.1 Physical assumptions and operating conditions

This chapter provides a set of mathematical models describing the aerodynamic channel from the leading to the following unmanned aerial vehicles (UAVs). The model links the relative position of both aircraft to the set of disturbance signals observed by the follower inertial measurement unit (IMU) and provides the theoretical basis for the detection logic and state-estimation architectures described in Chapter 4. The chapter begins with a description of the physical assumptions and operational envelopes of the presented model.

The primary consideration pertains to a geometrical setup of the multi-agent formation. The follower aircraft is positioned within the lower half-space of the leader, with its location described with the use of cylindrical coordinates: the distance to the central wake axis r and the depth below the plane of the leader's rotor z. The aerodynamic channel is not defined with a particular input-output relation but is given as a continuous vector field over the (r, z) domain. The choice of this particular domain follows from the identification of the aerodynamic field of the follower aircraft subject to the downwash field of the leader aircraft. More precisely, a set of longitudinal cross-sections of the wake vortex (along the centerline) and its transverse cross-sections at various depths are used in the model. Thus, the location of the follower aircraft in the wake field is uniquely identified by the coordinate (r, z), and for each such position, the aerodynamic disturbance can

be determined.

The secondary assumption is related to the properties of the wake: In the assumed volume of activity, the downwash is approximated to a stationary flow field with its symmetry axis aligned. The assumption implies that the model accounts only for a time-averaged mean velocity field that is imposed on the follower. High-frequency turbulent fluctuations are filtered out. This simplification is consistent with the temporal resolution of the anemometric instrumentation used for field characterization, as well as the established mean-flow paradigms formulated by Bauersfeld et al. [1]. The assumption of axisymmetry assumes that the fluid properties are strictly functions of the radial offset r and axial depth z, and are completely invariant with respect to the azimuthal angle. This structural property was empirically validated during testing: anemometric readings acquired at uniform radii across azimuths coincided within the sensor's measurement uncertainty threshold. As such, this confirms that the macro-scale wake behaves as an isotropic round jet, as theoretically expected at spatial scales significantly downstream from the rotor source.

The tertiary assumption relates specifically to the hardware platform since the experimental vehicle corresponds to the Parrot Mambo nano-scale quadrotor, which is configured on an X-frame geometry, possesses a total mass of approximately 63 g, and a half-wheelbase (rotor arm) of nominally 0.06 m. The latter attribute is significantly lower than the dimensional bounds (between 230 g and 6.3 kg) of the macro-scale vehicles originally used to validate the turbulent-jet model by Bauersfeld et al. [1]. As a consequence of this, the fluid wake of the Mambo possesses a higher degree of collimation compared to the baseline macro-model, which is analyzed in detail in Section 3.2 and is treated strictly as a fluid-dynamic characteristic of the nano-scale platform.

The final assumption provides specific environmental boundary conditions and an empirical measurement approach, since the spatial characterization of the aerodynamic field was realized in a controlled indoor environment where the precisely measured floor grid provided an absolute position reference. During the profiling, the leader UAV was constrained to a relative position to a measurement apparatus, with rotors at equilibrium hover RPM producing a continuous steady-state downwash from a deterministically known coordinate. The resultant mathematical model explicitly isolates this pair's leader-follower aerodynamic coupling, which implicitly assumes the leader wake to be the dominant exogenous disturbance, classifying secondary aerodynamic effects such as ground-effect or follower's slipstream recirculation as higher-order terms. Such an approach is dictated by ambient baseline measurement taken before the test campaign: without active rotors, the ambient room airflow registered at roughly $0.2\,\mathrm{m\,s^{-1}}$, coinciding with the absolute noise floor of the instrumentation. As such, ambient environmental disturbances are mathematically insignificant compared to the active downwash stream, whose

axial velocities are an order of magnitude greater than background noise.
Given the four assumptions, the implicit channel may be expressed as a deterministic spatial transformation from the relative coordinates (r, z) to the moment generated on the follower airframe. This deterministic transformation is contaminated by endogenous noise, characterized in detail in later sections. The following sections derive this channel model in three steps. First, characterizing the fluid propagation and spatial attenuation of this downwash field (Section 3.2), then, defining the transduction mechanism by which this fluid field manifests into an inertial moment on the follower (Section 3.3), and finally, quantifying the system noise and Signal-to-Noise Ratio (SNR) that theoretically bound this channel's operating capacity (Sections 3.4 and 3.5, respectively).

3.2 Downwash propagation model and spatial attenuation

The development of the aerodynamic channel model required an in-depth analysis of the flow structure below the leading vehicle. This field of flow was measured using a current meter. The probe was moved from one point to another following the centerline of the wake and at various depths, as sketched in Figure 3.1. A description of the flow structure was then separated into a longitudinal decay along the center of the wake and the distribution across the wake.

Axial decay

Along the central wake axis, where $r = 0$, the mean fluid velocity attenuates as a function of the axial distance z. Adopting the turbulent-jet framework established by Bauersfeld et al. [1], the centreline velocity is parameterized using a hyperbolic decay law incorporating a virtual origin,

$$U_C(z) = \frac{C}{z - z_0}. \tag{3.1}$$

Applying a least-squares fit of this functional form to the empirical axial profile yields $C = 444$ and $z_0 = -34\,\text{cm}$, resulting in the specific relationship:

$$U_C(z) = \frac{444}{z + 34}, \tag{3.2}$$

where z is expressed in centimetres and U_C in metres per second. The negative virtual origin mathematically positions the apparent flow source above the physical rotor plane, a standard characteristic of developing turbulent jets. Figure 3.2 juxtaposes this analytical fit against the empirical measurements: the hyperbolic

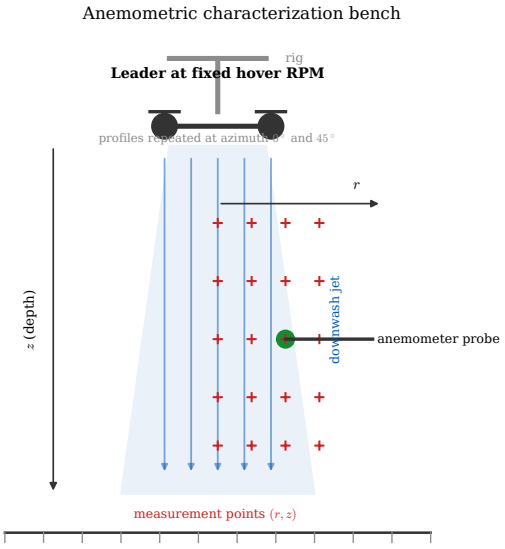

Figure 3.1: Anemometric characterization bench. The leader is held at fixed hover RPM on a rig; a probe traverses the downwash jet, sampling the mean velocity at points (r, z) along the centreline and at increasing radial offsets. Radial profiles are repeated at azimuths 0° and 45° to check axisymmetry.

law demonstrates high fidelity across the entire evaluated domain, from 10 cm to 50 cm, wherein the peak velocity attenuates from approximately $10\,\mathrm{m\,s^{-1}}$ down to $5\,\mathrm{m\,s^{-1}}$.

Radial profile

At any constant axial depth, the local fluid velocity attenuates as a function of the radial offset. The standard turbulent-jet paradigm predicts a bell-shaped, self-similar radial profile, formulated as:

$$U(r,z) = U_C(z)\left[1 + (\sqrt{2}-1)\left(\frac{r}{r_{1/2}(z)}\right)^2\right]^{-2}, \tag{3.3}$$

where $r_{1/2}(z)$ denotes the spatial half-width of the jet at depth z. In this regime, the empirical wake topology exhibits a pronounced physical divergence from the generalized model. The Mambo platform, with a mass of approximately 63 g, operates significantly below the inertial range of the vehicles (230 g to 6.3 kg) utilized by Bauersfeld et al. [1] to derive the baseline model. Consequently, its aerodynamic wake is significantly more collimated than the classical turbulent

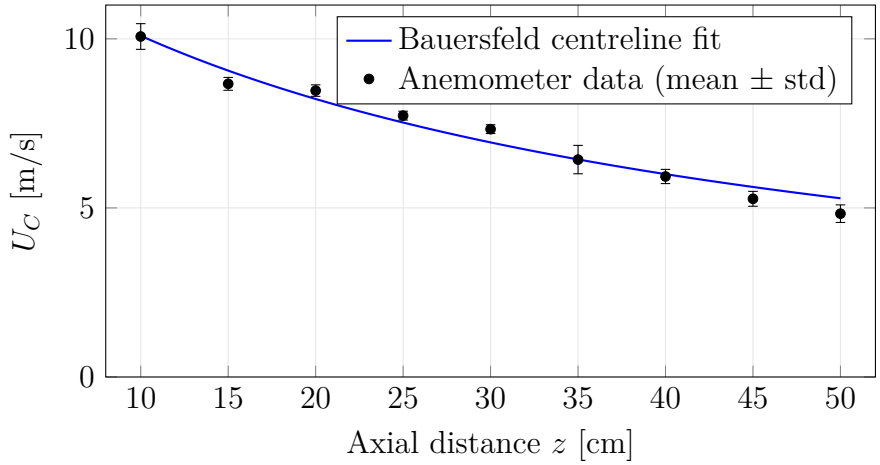

Figure 3.2: Longitudinal axial decay of the leader's downwash measured on the central wake axis ($r = 0$). Markers indicate the empirical mean velocity, with error bars representing the standard deviation across repeated measurements. The solid curve denotes the fitted hyperbolic centreline law $U_C(z) = 444/(z + 34)$

profile predicts. This structural collimation is evident in Figure 3.3. At all profiled depths, the flow sustains near-maximum centreline velocity up to a radial offset of approximately 10 cm, followed by a sudden spatial collapse between 10 cm and 15 cm. Beyond this collapse boundary ($r \geq 15$ cm), the wake velocity attenuates to the absolute detection floor of the anemometer, approximately $0.2\,\mathrm{m\,s^{-1}}$. In instances where direct readings at these extreme radii were omitted, the data points are formally bounded at the instrumentation floor, as measurements from adjacent depths verify that the collapsed wake possesses no kinetic energy resolvable above the sensor's sensitivity threshold. The classical turbulent-jet morphology, which dictates a half-width expanding linearly with depth, yields a far more gradual decay and is mathematically incapable of reproducing this severe spatial truncation. This narrow, highly collimated wake structure is an intrinsic fluid-dynamic property of the nano-scale quadrotor; it is documented here as a fundamental empirical finding rather than a measurement artifact, directly attributable to the application of the model at a previously uncharacterized micro-scale.

To ensure that the spatial model accurately replicates this observed physical behaviour, the half-width parameter $r_{1/2}(z)$ is determined from an independent empirical fit at each discrete depth, rather than being constrained to adopt a linear growth profile in z. This unconstrained free-width parameterization then successfully captures the severe radial truncation of the Mambo wake at all depths and is rigidly constrained by the $r = 0$ boundary conditions imposed on the axial dataset. The resulting continuous bivariate map $U(r, z)$, constructed from the axial

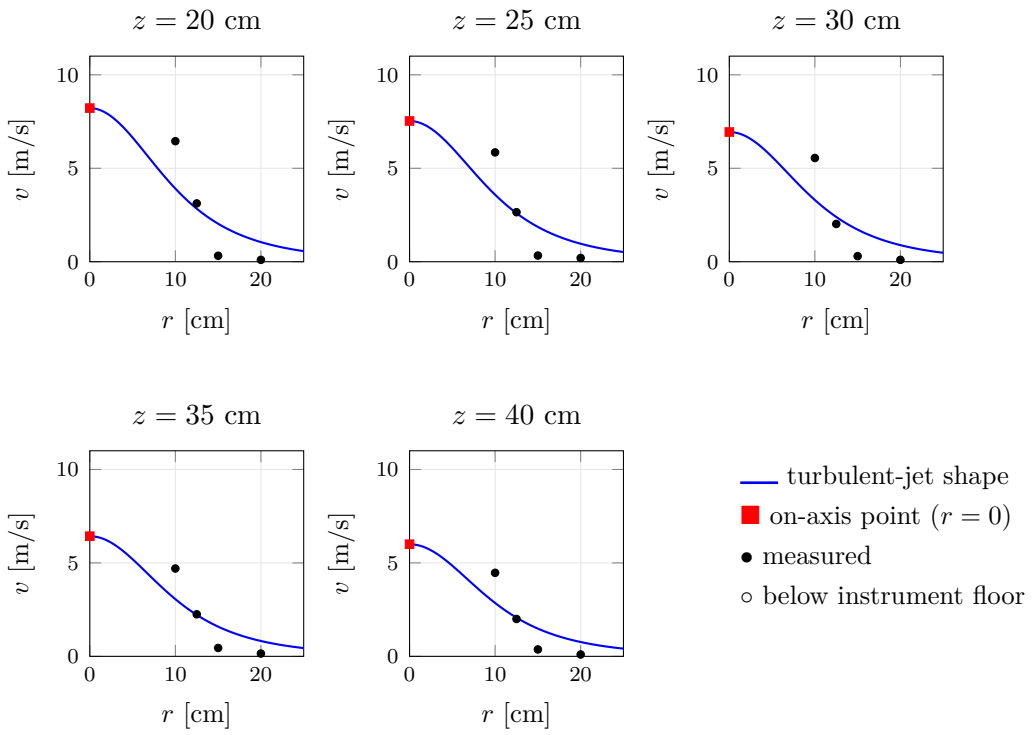

Figure 3.3: Transverse radial profiles of the leader's downwash, discretized by axial depth. Solid markers indicate valid empirical measurements; open markers signify spatial coordinates where the local velocity falls below the sensor's absolute detection floor ($\approx 0.2\,\mathrm{m\,s^{-1}}$), denoting a complete structural collapse of the wake. The continuous blue curve illustrates the theoretical turbulent-jet distribution anchored to the respective on-axis peak velocity. Notably, at every depth, the flow sustains peak magnitude until roughly 10 cm, precipitating a severe spatial collapse between 10 cm and 15 cm.

formulation of Equation 3.2 combined with the depth-dependent radial profiles of Equation 3.3, provides an accurate representation of the mean velocity field across the empirically validated operational volume.

Axisymmetry

A foundational premise of the proposed model is azimuthal invariance: the fluid properties are postulated as strict functions of r and z only, being completely independent of the azimuthal angle. This geometrical assumption has been verified by replicating the transverse radial profiling at discrete azimuthal orientations of 0° and 45°. The independent datasets are found to be equal within the instrument's

resolution limit, with the maximum variance between the azimuthal replicates being $0.31\,\mathrm{m\,s^{-1}}$, which is below the anemometer's noise floor. Thus, the wake is demonstrated to be axisymmetric within the available limits of precision, in accordance with round-jet theory. The divergent azimuthal sweeps are therefore considered as isomorphic replicates of the same spatial coordinates in the analytical model.

3.3 Channel observability from onboard follower measurements

In the last section, we obtained the linearized expression for the leader wake velocity field, $U(r, z)$, which is not directly accessible to the follower, since the follower does not possess air-data sensors: it has only an inertial sensor for flight control, and the question arises as to whether the relative leader's position can be reconstructed from the inertial measurements. This section provides the connection between the wake velocity field and the measurement available to the follower.

The follower is subject to a moment the wake applies to its airframe, due to the asymmetry in the jet load; the moment is a result of the uneven distribution of jet load, and not the speed of the jet at a point. The follower must generate a counter-moment to the moment produced by the jet to maintain level flight. The inertial unit detects this counter-moment Kiran et al. [15] document the same mechanism in a pair of quadrotors: a horizontal offset of the follower leads to a shift in the strongest part of the wake away from the centre of the lower vehicle, inducing a downwash-induced pitch moment which the vehicle must counter to stay stable.

The size of the moment is determined by the geometry of the asymmetry. Directly below the leader, the follower is at the center of the wake, the load is distributed evenly (or canceled) over the airframe, and the moment is small (possibly zero). However, as soon as the follower deviates from the centerline, the situation changes: one side is under the jet, while the other is exposed to the clear air, so the imbalance grows, resulting in a larger moment. Finally, a follower that is too far away from the centerline will stop being exposed to the jet. The moment decreases back to zero. Hence, the moment is smallest at the centerline, reaches maximum imbalance at a certain offset, and diminishes to zero as the follower moves far away from the leader. This pattern can be used to indicate the direction of the follower, but it has a peculiarity: it peaks not on the fastest flow (the axis) but on the strongest imbalance of loading between the left and right sides of the aircraft.

The moment appears on the follower as an uncommanded change in attitude. The onboard Extended Kalman Filter (EKF) estimates the follower pitch and roll, and the flight controller issues commanded setpoints to track its own reference. The

wake-induced part is isolated by subtracting the command from the estimate, giving the attitude residuals

$$\Delta\theta = \hat{\theta} - \theta_{\text{cmd}}, \qquad \Delta\phi = \hat{\phi} - \phi_{\text{cmd}}, \tag{3.4}$$

where $\hat{\theta}$ and $\hat{\phi}$ are the EKF estimates and θ_{cmd}, ϕ_{cmd} the setpoints. These residuals are the observable output of the channel: the disturbance the wake writes onto the follower attitude, read through the sensing already present on the vehicle.
Two properties of the channel follow. First, the bearing to the leader is observable; a leader ahead produces a front to back load on the airframe which causes a pitch residual, while a leader to the side produces a roll residual. An in-between bearing would result in a moment split between the axes. Thus, the combined action gives a bearing which can be found from the two components. This is taken up in Chapter 4.
Second, the range is only observable in part. The magnitude of the residual depends on the moment, which increases from a minimum near the axis to a maximum at some intermediate offset, and then decreases again. Thus, a given residual magnitude corresponds to two offsets, so the mapping from magnitude to range is not globally invertible. The range is only observable on the monotonic portion of the curve, the exterior of the peak, which defines the useful sensing region of the channel. The inversion and the valid branch are discussed in Section 3.6. Taken together, the channel makes the bearing fully observable and the range observable on a restricted branch, from the inertial signal alone without any dedicated sensor. This is the observability on which the estimator of Chapter 4 is constructed.

3.4 Noise sources and uncertainty in close-proximity flight

The aerodynamic channel encodes the induced aerodynamic moment about the center of mass via attitude-tracking residuals defined in Section 3.3 in its spatial coordinates. However, the signal is corrupted by stochastic disturbances inherent to the dynamics of the vehicle and its instrumentation. These disturbances define a noise floor, which the wake-induced signal has to overcome to be detected. The goal of this section is to characterize the statistics of this noise, as its variance ultimately defines the minimum aerodynamic disturbance which can be detected, i.e., the operational envelope of the implicit channel.
The baseline attitude tracking residual of the follower contains a non-zero variance even when there is no leading vehicle. The main reason is an inherent stochasticity of the onboard inertial sensors. In particular, commercial-grade gyroscopes have several noise sources which are inherent to specific implementations, e.g., quantization noise and white noise processes which are dominant in the high-frequency band,

according to the Allan-variance specification, with a particular manifestation as an angle random walk [38]. Another source is the high-frequency mechanical vibrations from the propulsion system, which are not always effectively attenuated by the digital low-pass filters in the attitude control loop, thus manifesting in the residual as another oscillatory disturbance. Moreover, the high-frequency corrections from the flight controller to reject low-frequency disturbances during steady-state hovering are not entirely canceled by subtracting the commanded attitude rates. Overall, these effects combine in a manner that resembles a Gaussian distribution on each axis, which will be confirmed with the following analysis of the distribution.

In order to assess the noise, the data were gathered from 11 different flights in the same steady-state hover in a clean environment, such that the attitude residuals represent only the inherent noise of the system. The noise distributions are then characterized by the median absolute deviation, rather than the standard deviation, as a means of noise dispersion estimation. This is done because, during the flights, occasional large outliers were observed in the data. The standard deviation, as an estimator of dispersion, would penalize those outliers by artificially increasing the variance, while the MAD estimator would be less sensitive to them.

Considering the axes separately, the robust dispersion of the residuals is about 9.5 mrad in pitch and 14.5 mrad in roll. This significant asymmetry, which dominates the variance, is even more prominent in the robust estimation and thus appears to be a physical characteristic of the airframe dynamics, not an artifact of the statistical model.

Since the proposed detection architecture is considering the Euclidean norm of the biaxial residuals, and not the independent axial components, under undisturbed hover conditions, this norm typically has a magnitude of around 14 mrad, with a healthy dispersion of about 9 mrad. This scalar value therefore represents the noise floor threshold which the aerodynamic wake signature has to overcome in order to be detected. The detection logic, therefore, as introduced in Chapter 4, is based exclusively on the statistics of this aggregated norm statistic, and not the axial components separately.

Furthermore, the stochastic noise profile displays strong temporal autocorrelation. At the nominal 200 Hz control loop frequency, consecutive residual samples are by far not independent. As the closed-loop attitude dynamics act as a low-pass filter on the residual signal, a significant spectral shaping can be observed. Empirical decorrelation time-constants amount to 400 ms for the pitch and 460 ms for the roll channel. Thus, the effective degrees-of-freedom, i.e., the number of statistically independent samples, is much lower than the nominal sampling frequency suggests. This strong temporal correlation needs to be taken into account when the noise floor is dynamically estimated over a finite time window.

Finally, the absolute theoretical limit to system detectability is set by the empirical noise floor. The aerodynamic wake is therefore only detectable (in a region of space)

Table 3.1: Empirical noise floor of the follower in quiet hover, aggregated over eleven leader-free flights. Dispersions are robustly estimated utilizing the median absolute deviation.

Quantity	Value	Unit
Pitch residual dispersion	9.5	mrad
Roll residual dispersion	14.5	mrad
Norm typical level	14	mrad
Norm dispersion	9	mrad
Pitch decorrelation time	400	ms
Roll decorrelation time	460	ms

if the induced rotation moment exceeds the established threshold of approximately 14 mrad in magnitude. By combination with the fluid-dynamic spatial attenuation models as described in Section 3.2, this value forms the limit of the channel's operational sensing volume by the stochastic noise boundary. Beyond a specific radial offset (in terms of moment), the wake-induced rotation moment is no longer distinguishable from the stochastic noise. Thus, in terms of forward traffic flow, the leading vehicle becomes fully undetectable to the follower beyond a certain distance from the host vehicle. The following section will formalize this operational limit by a quantitative evaluation of the signal-to-noise ratio as a function of the relative spatial vehicle position.

3.5 Signal-to-noise ratio as a function of relative position

The operational envelope of the aerodynamic channel is governed by the signal-to-noise ratio (SNR), i.e., the ratio between the wake-induced signal perceived by the follower and the noise floor described in Section 3.4. Since in both cases the quantities are measured in the same unit, namely, the Euclidean norm of attitude tracking residuals, the ratio is simply calculated without any conversion: the signal is the residual norm induced by aerodynamic wake while the noise floor is estimated during hover at a fixed position.

As established in the model observability section, the trailing UAV is designed to measure the aerodynamic moment induced by the wake, and not the actual local fluid velocity. This moment is due to the asymmetric inflow distribution across the rotor plane [15]; hence, whenever the UAV flies in a region of locally uniform flow field, the follower is subjected to symmetric inflow, which induces a null aerodynamic moment. On the contrary, if the fluid velocity shows a significant radial

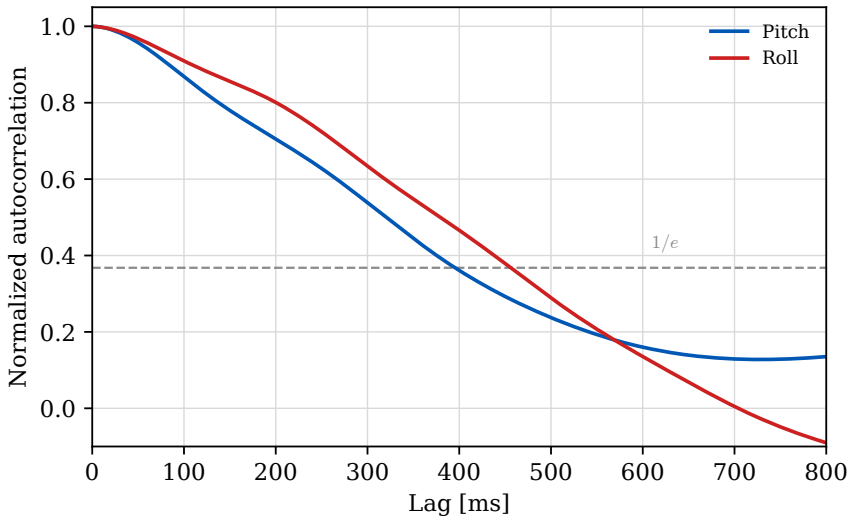

Figure 3.4: Normalized autocorrelation of the pitch and roll residuals in quiet hover, averaged over the eleven leader-free flights. Both axes stay correlated for several hundred milliseconds, crossing the $1/e$ level at about 400 ms in pitch and 460 ms in roll. The residual is therefore strongly coloured by the attitude loop, and the number of independent noise samples per second is much smaller than the 200 Hz sample count.

gradient, the inflow experienced is also non-uniform, and a rotational perturbation is induced. The magnitude of this signal is proportional to the radial gradient of the velocity field mentioned in section 3.2, reaching a minimum at the centerline and a maximum at an intermediate radial offset; thus, a characteristic geometric shape as shown in Figure 3.5 is obtained.

Within the main wake structure, the signal magnitude dominates the noise level by a considerable margin. The results of the flight test with deliberate inter-vehicle offsets demonstrate that the upper limit of the residual norm (90th percentile) varies from approximately 60 milliradians near the middle and at the edge of the sensing band to about 90 milliradians at the intermediate peak offset. Compared to a baseline noise level of approximately 14 milliradians, the operational signal-to-noise ratio (SNR) is on the order of 13 to 16 decibels within the wake. This provides an ample detection margin: the wake signal magnitude is several times higher than the noise level throughout the volume of interest, ensuring reliable leader detection when the follower is within sensing range. It is notable that the location of the signal maximum is not found on the central wake axis. Whereas the induced perturbation is relatively weak directly beneath the leader, taking values of the order of 60 mrad, it reaches its peak value of nearly 90 mrad at an

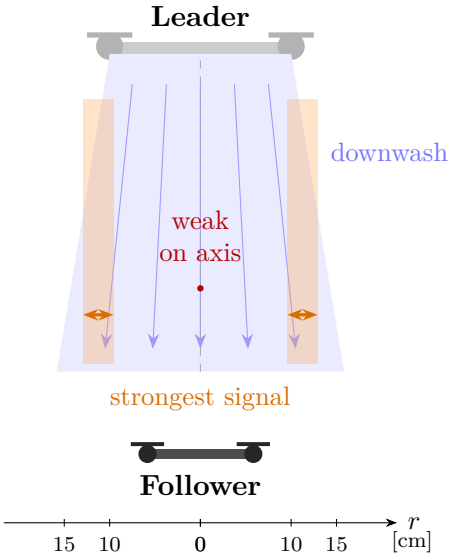

Figure 3.5: Geometry of the wake signal. The leader downwash descends as a jet that reaches the follower below. The induced moment is weak on the axis, where the inflow is symmetric, and strongest at an intermediate offset inside the wake, where the radial gradient of the velocity field is largest. Beyond the edge of the jet the signal falls to the noise floor. The useful sensing region is the shaded band within the downwash, not the axis and not the far field.

intermediate offset, before decaying outwards. This evolution is entirely consistent with a moment due to a radial velocity gradient, which would take its maximum on either side of the central axis, be zero directly on the axis due to symmetry in the inflow, and be maximal at the intermediate offset in the gradient across the vehicle's footprint.

Consequently, the effective sensing region takes the form of an annulus, rather than a continuous disc extending to the edges of the measurement volume. This annulus is bounded in the interior by an on-axis region in which the signal has been attenuated to the noise floor due to fluid-dynamic symmetry, and an exterior boundary beyond which the jet structure has sheared and decayed to the noise floor. The anemometric characterization of the wake in Section 3.2 identifies this outer boundary as occurring at approximately 15 cm from the jet axis, beyond which the wake structure effectively collapses. This annular geometry is the distinctive feature of the aerodynamic channel. The main restriction it imposes is not a restriction on the absolute SNR (which is acceptable within the wake boundary), but a geometric restriction: the phenomenon is confined to a narrow annular region, restricting acquisition to times when the follower is traversing within this annular corridor. Furthermore, the non-monotonic nature of the spatial distribution is the defining

challenge for range estimation described in the following section: given that the signal magnitude rises then falls with radial distance, any single measurement defines two possible radial distances, and the function relating the two is single valued only on the descending branch.

3.6 Inversion of the wake model for range estimation

The relative bearing estimation procedure described in Section 3.3 is well-posed: it is fully determined by the direction of the attitude tracking residual vector. The radial range estimation, however, is not that straightforward, as it involves the inverse transformation of the scalar magnitude of the measurement residual into an estimate of position offset, which is only possible locally, due to the function's range not being invertible globally. In this section, we formalize the estimation procedure, describe the branch of the function used, derive the analytical expression of the estimator, and verify its correctness using the flight telemetry.
The magnitude of the signal is proportional to the gradient of the aerodynamic velocity field in the radial direction (see Section 3.5), which has a non-monotonic dependence on the radial offset r: the gradient is minimal on the centerline, rises to a local maximum at some offset, and decreases monotonically beyond that. A function that has a local maximum cannot be globally inverted: a given signal magnitude corresponds to two different radial offsets, one on the increasing inner side of the peak and one on the decreasing outer side. Therefore, inverting the magnitude directly yields significant spatial ambiguity at the peak.
In order to obtain mathematical consistency, only one branch of the curve, either left or right, is considered for the inversion with respect to the peak, so that the magnitude defines the offset uniquely. This may be achieved, for instance, by setting the lower bound on the reconstructed range to be close to the theoretically possible maximum offset, thus eliminating ambiguity caused by the second branch. In this way, the follower UAV can approach the leader up to the distance of maximum offset safely, maintaining the one-to-sixty correspondence between the states.

The aerodynamic signal inside the monotonic outer branch decreases at an increasing rate with the distance from the centerline. The estimator fits an exponential decay function to the residual norm versus range

$$\mathrm{norm}(r) = A\, e^{-Br}, \tag{3.5}$$

and subsequently extracts the estimated range $\hat{r}$ via analytical inversion,

$$\hat{r} = \frac{1}{B} \ln \frac{A}{\mathrm{norm}}. \tag{3.6}$$

To suppress the effect of high-frequency transients, the reconstructed range is passed through a three-point median filter and clamped to the operating range defined by the peak offset and a maximum range. If the signal magnitude falls below the specified detection threshold, the current estimate is held for a given period of time and then relaxed towards a default rest value. This procedure prevents an indefinitely held outdated state estimate due to signal loss. It should be noted that such a mathematical relaxation introduces a fluctuation of the estimated range between individual detection events, independent of the actual vehicle dynamics. The onboard state estimator telemetry shows that our assumption of monotonic inversion is valid on the outer branch. This was verified through dynamic pursuit experiments on several flight logs. Plotting the filter's range output against the instantaneous residual norm simultaneously confirmed the conjecture. Following a successful leader detection, the correlation between the sensed signal and reconstructed range is largely monotonic, as sensor signal strength is positively correlated with decreases in radial separation. The Pearson coefficient of this correlation is approximately -0.95, indicating that there is significant repeatability of this relationship throughout the various experiments. The signal's range correlation converges to the same functional relationship within a span of a few milliradians. Also, the estimated ranges are clamped to the lower bound of the peak offset to prevent the closed-loop dynamics from encountering the non-monotonic branch, even at the minimum inter-vehicle separation.

3.7 Limits of the model and expected failure regions

The model of this chapter maps the relative leader's position to the signal the follower reads, and it holds over a bounded region under stated assumptions.

The first limitation is a geometrical one. The signal that can be carried by the channel is only useful within the wake, in the annular band of Section 3.5. On the axis, the induced moment is weak due to symmetric inflow across the follower, and the model is unable to distinguish position. Beyond the collapse of the jet, close to 15 cm of radially offset the anemometer field, the flow drops to the instrument floor and no signal is incident on the follower. Outside this band, on the axis or the far field, the model predicts no observable signal and the follower is blind.

The second limitation is non-monotonicity in the range map. The follower's moment is not created by the intensity of the downwash at the follower's location but by an asymmetric loading of the wings. When the follower is displaced, the wake affects only one side of the vehicle, causing an asymmetric load on the aircraft. This induces a moment, which the follower must counteract with a restoring moment, the magnitude of which is detected by the inertial unit. Kiran et al. [15] describe

the same mechanism: a horizontal offset produces a horizontal shift in the highest velocity region, causing the downwash-induced pitch moment to require a restoring moment for stability.The size of this moment is defined by the geometry of the asymmetry, not the local flow speed alone, and is therefore largest at the offset where the jet loads one side of the vehicle while leaving the other unloaded, falling off on either side of that offset. The signal magnitude is therefore non-monotone in radial offset, taking a given magnitude twice, once on each side of the offset of maximum, and range recovery is single valued only beyond the maximum on the outer side. The estimator remains on this side by clamping the range at the offset of the maximum, and closer in it uses bearing only.

The third limitation is the scale of the platform. The velocity field of Section 3.2 was characterised on a nano-quadrotor of about 63 g, below the range on which the turbulent jet model of Bauersfeld et al. [1] was established. The measured wake is more collimated than predicted by the model, so the radial profile of the model is a fit to the measurements at this scale, not a transfer of the published model. The model should not be extrapolated to vehicles of a different size without re-characterising the field.

The fourth limitation is that there is no absolute calibration: The signal-to-noise ratio of Section 3.5 is stated on the measured attitude residual, and the range inversion of Section 3.6 is a consistency check against the onboard estimate. Neither is tied to an external measurement of the true relative position, as there is no external tracking available on this platform. The model is therefore internally consistent and controllable, but has not been established as absolutely accurate, and should be read as a description of the channel rather than a calibrated position.

These limitations define the envelope within which the channel model is used in the rest of the thesis. Within the wake, on the outer branch, at this platform scale, and for a slowly moving leader, the model relates position to signal well enough to support detection, bearing estimation, and range recovery on the outer branch. The chapters which follow develop the estimator and the pursuit loop within the channel within this envelope, treating the regions outside as regions of expected failure rather than as cases the model is meant to represent. CAPITOLO 5

Chapter 4

Detection and Estimation of the Leader Position

Chapter 3 provided a theoretical depiction of the underlying aerodynamic channel. It included the theoretical link between the relative spatial state of the leader and the induced aerodynamic moment affecting the follower attitude dynamics, and delineated the geometric region over which this deterministic relation is valid. This chapter takes this relation into an operational real-time estimator, and it validates the design using dynamic simulations. The pipeline takes as input the raw endogenous attitude residuals, and it autonomously estimates onboard a binary presence detection flag, a relative bearing vector, and a radial range estimate.Combined these variables describe the leader relative position in a polar coordinate frame that provide the state feedback of the closed-loop pursuit controller described in Chapter 5. Notably, each step in the pipeline is executed using only the baseline Inertial Measurement Unit (IMU) avoiding the use of additional exteroceptive sensors or active radio-frequency communications links.

The proposed architecture was evaluated using the Simulink model of the follower UAV, where the model of the leader's aerodynamic wake was injected as an external disturbance acting on the follower's airframe. At the time of the current study, the experimental characterization of the follower's airframe from Chapter 3 had not yet been completed, hence it was replaced by the turbulent jet law from Bauersfeld et al. [1] as the model for the rotor-induced flow field. Subsequent measurements confirmed that the injection of the flow field characterized from the anemometry data leads to an induced moment of the same topology as that from the current model. With respect to the setup of the Simulink simulation, it employs a virtual leader following a pre-defined trajectory, which allows for determining the true relative position at each point in time. This is in contrast with the experimental setup described in Chapter 7, where such a position is not available, only the

estimated position from the vision algorithm can be compared to the ground truth.

This chapter is organized to follow the processing chain used for the estimation. First, a rotationally invariant detector that can be used to detect wake interactions at any approach angle is formulated in Section 4.1. Then, a dynamic threshold is added to the detector using an online noise-floor estimator, which is described in Section 4.2. Temporal signal conditioning is applied to the detector output in Section 4.3 to provide reliable event classification. The computation of the relative bearing is then detailed in Section 4.4, with the radial range estimation found in Section 4.5. A relative position state is formed from the two orthogonal measurements in Section 4.6.Section 4.7 contains a theoretical comparison of different filtering approaches, which forms the basis for the choice of the estimator. In Section 4.8, performance trade-offs specific to the considered system are discussed.

4.1 Vector-norm detector for rotationally invariant wake sensing

The main aim of the estimation pipeline is to determine the physical presence of the leader's aerodynamic wake at each discrete computational instant. The legacy baseline architecture proposed by Cerruti [5] performed this binary classification either using the pitch residual exclusively, or by combining the pitch and roll residuals using a disjunctive logical OR. A detection was thus declared if either independent residual independently exceeded a static geometric interval. This topology introduces a severe geometric blind spot, since approaching a leader vehicle along an intermediate diagonal bearing (e.g., nearing forty-five degrees) causes the aerodynamic moment to distribute itself symmetrically with respect to the axes, such that each orthogonal component individually registers only a fractional amplitude of approximately 0.707 relative to the total disturbance vector. In this scenario, both isolated components can fail to individually trigger the static single-axis thresholds even when the aggregate disturbance is well beyond the detection boundary. This is true precisely at the lateral approach vectors that a robust multidirectional pursuit controller must actively navigate.

In order to solve the aforementioned problem, the proposed estimator evaluates both attitude tracking residuals simultaneously, namely: the pitch residual $\Delta\theta$, and the roll residual $\Delta\phi$. The resulting aerodynamic disturbance forms a vector within the continuous $(\Delta\theta, \Delta\phi)$ space. Its scalar value encodes the absolute magnitude of the wake interaction, independent of its azimuth. Replacing the disjunctive logic of the previous estimator by the Euclidean-induced norm of the two residuals,

$$\rho = \sqrt{\Delta\theta^2 + \Delta\phi^2}, \tag{4.1}$$

undamentally eliminates the geometric blind spot. The Euclidean norm ρ preserves

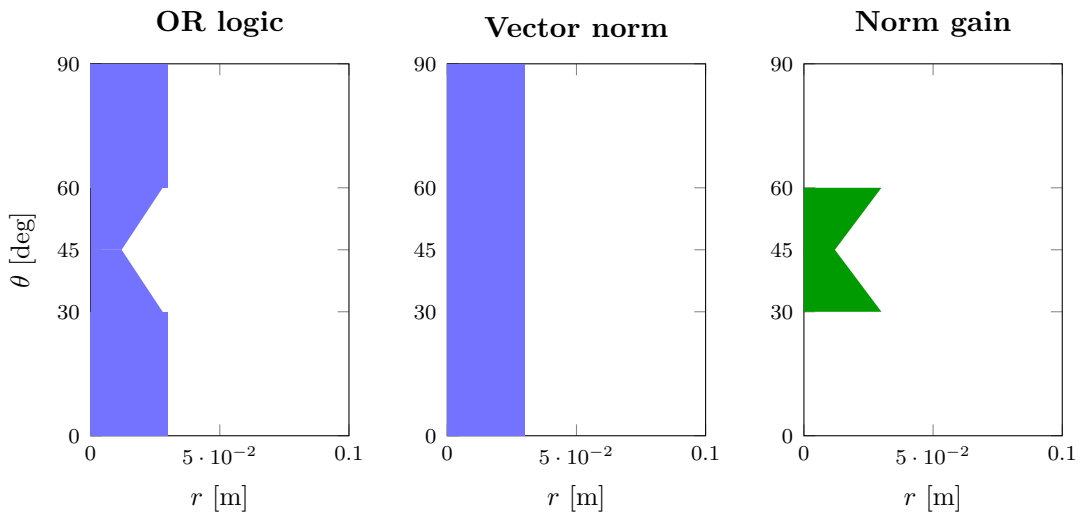

Figure 4.1: Detection zone in the relative plane (r, θ) from simulation, for the OR logic of the baseline and for the vector norm. The OR logic detects (blue) across most angles but leaves a gap around 45°, where the wake splits between pitch and roll and each component falls below the fixed interval. The vector norm detects across the whole disc with no gap. The right panel shows the region the norm recovers (green), a band centred on 45°.

total disturbance amplitude for any bearing, so a leader at 45 degrees has the same norm as a leader directly ahead at the same distance. Therefore, the range limit takes on the familiar shape of a circular sensing disc around the follower rather than the cross-shaped pattern with its distinctive dropoff along the diagonals.

Simulations confirm this claim. We performed a two-dimensional sweep of the leader's position on the relative Cartesian plane to test the tracking ability of the norm-based detector compared to the classical OR logic. The results show that the vector norm detector can correctly identify the leader in all directions of the sensing disk, while the OR logic fails to detect the leader in diagonal directions (around 45 degrees). On the contrary, the case when the OR logic detects the leader, but the norm-based detector does not is negligibly small. Thus, the vector norm detector provides complete sensing coverage with the same proprioceptive data as the baseline IMU without additional computational overhead.

4.2 Adaptive thresholding based on online noise-floor estimation

The effectiveness of a detector based on the residual norm ρ depends on its threshold. The baseline used a threshold of 0.06 rad, chosen under full battery conditions. However, this threshold proves too sensitive as the voltage supplied to the UAV decreases throughout a flight. As the UAV's battery capacity depletes, the sensor exhibits increased variance and requires more control effort, thus increasing the baseline value of the quiet-hover norm. Thus, a threshold based on the voltage at peak capacity will be exceeded prematurely during a flight, resulting in many false positives. Consequently, an effective threshold must track the time-varying noise floor of the system.

The first version of the adaptive threshold algorithm estimated online the variance of the ambient noise during periods of undisturbed hover. It accumulated norm samples below a given scalar gate that were considered to be due to isolated noise and calculated their standard deviation σ. Then it set the threshold equal to four times the standard deviation:

$$T = 4\,\sigma. \tag{4.2}$$

This eliminates the need for the user to adjust the threshold relative to the noise level because the threshold is automatically adjusted according to the background noise. In the idealized case, this tremendously decreased the false-alarm rate while keeping the true-positive detection rate the same; for the standard 3-scenario test suite, the static threshold resulted in about 9 false alarms, whereas the dynamic threshold only had 1 false alarm (a decrease of 92 % in false alarms).

While this four-sigma threshold architecture worked well in simulation, in practice, on a real flight log, it has a critical structural flaw that had to be addressed in the final algorithm. To avoid corrupting the baseline noise estimate by large wake signatures, the algorithm only allows samples to enter its temporal buffer if they fall below a restrictive geometric gate. In reality, however, the true quiet-hover norm fluctuates around this theoretical gate. As a result, when the leader is present, only a small percentage of samples (in the order of a few percent) successfully pass the gate. The temporal buffer is quickly drained and halts updating; the dynamic threshold freezes permanently to the absolute bottom of its allowable range for the rest of the flight. This design flaw is circular; the scalar gate required to safeguard the noise estimate logically needed the very adaptive threshold it is needed to continually generate.

The finalized, flight-proven architecture mathematically resolves this circularity via full excision of the gating mechanism, using an order statistic of the recent residual norm as a threshold, implementing an order-statistic constant false alarm

rate (OS-CFAR) detection paradigm [39]. A receding time window containing the most recent W norm samples is continually sorted by magnitude; the localized noise level is taken to be the q-th quantile of this array, and a dynamic threshold is formulated as a deterministic scalar multiple thereof:

$$T = \alpha\,\hat{\rho}_q, \qquad \hat{\rho}_q = \text{quantile}_q\Big(\rho_{n-W+1}, \ldots, \rho_n\Big). \tag{4.3}$$

The employment of an order statistic dispenses with the requirement of a scalar gate. Given that the sliding window houses the high-level interactions of the leader, said deterministic perturbations are enclosed within the upper extreme tail of the ordered statistics, strictly above the q-th quantile boundary. As such, they are excluded from the baseline noise calculation. The OS-CFAR estimate conforms to the true stochastic noise floor, wholly independent of the presence or absence of transient aerodynamic signals. For the current nano-quadrotor platform, this architecture is parameterized with a window of $W = 400$ samples (two seconds of data at 200 Hz), a statistical quantile of $q = 0.75$ and a scaling multiplier of $\alpha = 3$. When applied to empirical flight logs, this formulation maintains a false-alarm rate of approximately one percent while detecting the leader in every trial, unlike the four-sigma threshold which was held fixed in its lower clamp. This progression is a transparent reminder that, although the four-sigma mechanism worked in simulation, the buffering failure was unveiled in flight, necessitating the theoretically rigorous resolution of OS-CFAR.

4.3 Debounce logic for robust event detection

The OS-CFAR threshold produces a raw binary detection flag at each discrete sampling instant, but rarely is there a continuous Boolean signal corresponding to a single physical aerodynamic encounter. Due to inherent turbulent fluctuations in the downwash, the residual norm often crosses the threshold boundary, resulting in a single wake interaction producing a high frequency burst of transient detections. If sent downstream to a state estimator, this burst would be erroneously interpreted as multiple targets quickly appearing and disappearing, which could destabilize the pursuit control logic.

In order to address this issue, the architecture uses a confirm-and-hold temporal debounce filter: a candidate detection is only confirmed by the filter once the raw binary flag has been continuously asserted for a certain number of consecutive samples, in order to avoid isolated single-sample noise spikes erroneously crossing the threshold. The confirmed detection is then held active for a cooldown interval (until the cooldown counter has expired) during which any additional raw threshold crossing will be merged with the current event; the merged event is only ended when the cooldown counter reaches zero again after such a crossing. For this

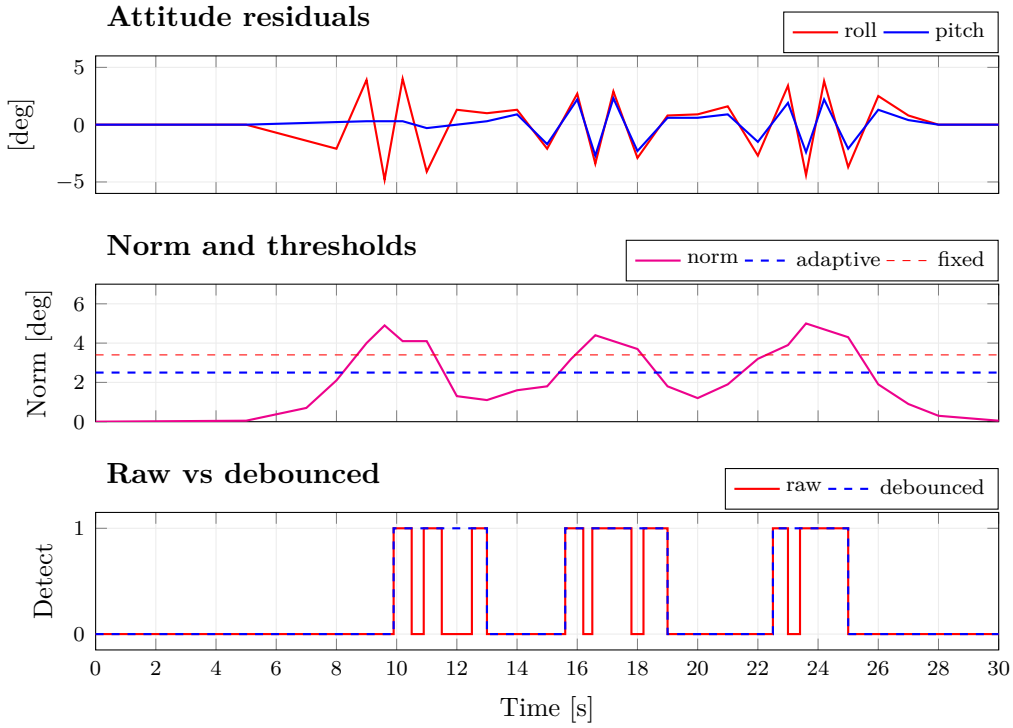

Figure 4.2: Three simulated approaches at $\theta = 0°$ (t=10 s), $\theta = 45°$ (t=17 s) and $\theta = 90°$ (t=24 s), with the wake injected through the Bauersfeld model. Top: the pitch and roll residuals. Middle: the residual norm against the adaptive and fixed thresholds. Bottom: the raw detection and the debounced output, which reduces the fragmented raw crossings to three clean events, one per approach.

platform, the confirmation threshold was set to $N_{\text{confirm}} = 5$ samples (25 ms) and the cooldown window to $N_{\text{cooldown}} = 1000$ samples (five seconds), so as to bridge sub-second gaps between encounters without merging different actual encounters. On the standard three-scenario simulation, this results in the nine raw detections being merged into three solid events, corresponding to the three actual approaches.

In practice, this temporal debounce is combined with a distance-based latch for the physical deployment. The latter arms immediately upon detection confirmation and holds the activation until the analytically estimated range has passed a certain threshold. This geometric cross-check binds the temporal duration of the event to the spatial dynamics of the encounter, as opposed to a fixed timer. The effectiveness of the spatial latch is then exercised and validated on the empirical flight logs described in Chapter 7, which feature continuous range estimation telemetry.

4.4 Approach-angle estimation from attitude residuals

Having established the presence of the leader, its relative bearing is obtained immediately from the vector direction of the aerodynamic disturbance. Since the moment due to the wake is collinear with the spatial vector between the two vehicles, the orthogonal attitude residuals contain the angular information: A longitudinal spatial offset predominantly disturbs the pitch axis, a lateral offset isolates the roll, and a linear combination of both offsets yields a corresponding mixture of disturbances. Thus, relative bearing is defined by the phase angle of the residual vector:

$$\hat{\theta} = \operatorname{atan2}\Big(\Delta\phi,\, \Delta\theta\Big). \tag{4.4}$$

This angle is calculated using a four-quadrant arctangent such that the full 360° domain of azimuths can be uniquely resolved with no ambiguity in the sign of the algebraic value In particular, the bearing estimate is not tracked continuously during a detection event, but rather peak-sampled: the bearing is only updated at the local maxima of the residual norm ρ within each event. The resultant disturbance in the estimate is then maximized, giving the highest possible Signal-to-Noise Ratio (SNR) and the best-conditioned vector, which strongly rejects the highly stochastic edge of the wake encounter.

It is precisely in the area of bearing estimation where the simulation validation fails to account for the idealization in the modeling. It is evident in Simulink that the ideal attitude dynamics severely mask the pitch response to the injected aerodynamic disturbance. As such, the pitch residual does not reach the true coupled magnitude as seen on the physical platform, and the estimated bearing collapses mathematically to the roll axis, independently of the true injected approach azimuth. The problem is rather the ideal rigid-body airframe model, which fails to capture the relevant complex cross-axis aerodynamic coupling pertinent to the true dynamics of the vehicle. As such, the bearing estimation pipeline is indeed validated by the empirical flight telemetry in Chapter 7, which does feature the natural cross-axis coupling, enabling the estimator to properly resolve the full circular domain of approach directions. However, the simulation environment is fully valid for benchmarking the magnitude-dependent stages, namely, detection and range estimation, which are decoupled from these specific directional artifacts.

4.5 Range reconstruction through downwash model inversion

Although the relative bearing limits the lead vehicle to an identifiable radial vector, the radial distance is crucial in recognizing the exact location on this vector. The amount of aerodynamic interference measured is not increasing monotonically even though the distance between vehicles is reduced toward zero. As shown in Section 3.6, the induced moment is dominated by the asymmetry of the downwash jet over the follower's footprint. Therefore, the sensor signal reaches a local peak and decays in both the direction towards the center of the wake, and radially outward. This function is non-monotonic, so a single value of magnitude corresponds to two distances, on either side of the extremum. As shown in Section 3.6, the induced moment is dominated by the asymmetry of the downwash jet over the follower's footprint. Therefore, the sensor signal reaches a local peak and decays in both the direction towards the center of the wake, and radially outward. This function is non-monotonic, so a single value of magnitude corresponds to two distances, on either side of the extremum. To eliminate uncertainty, the estimator confines the inversion technique to the outer part of the aerodynamic wake, which radiates out from its peak moment. The intensity of the signal decreases continuously with increasing radius, ensuring unique radial distance. When it comes to the virtual leader following a mathematically defined background trajectory in the simulation settings, this process makes it possible to assess the inverted distance against the true geometrical distance. By numerically inverting the wake distribution, the true distance can be obtained with the accuracy of less than 1 millimeter in theoretical terms. This leads us to the conclusion that the amount of the signal represents distance according to the monotonous outer part of the wake. For example, when the assumed diagonal flight is carried out with the real distance of about 0.057 m at the time of closest approach, the estimator managed to successfully identify the least distance.

It is strictly necessary to comply with both requirements. Therefore, the estimation epoch remains fixed at the estimation epoch number, and the internal ambiguous branch of the mathematical model is not taken into account. In addition, before executing the inversion algorithm, the original signal is processed through a median filter with a short memory, which allows eliminating high-frequency signals in one sample. In addition, as long as wake encounters are rare occurrences, the numerical inversion operation will cease when the last moment arrives in time when the residual norm of state is below OS-CFAR threshold, and from that moment the last known spatial estimate is kept until the end of one particular time period, when it is governed to decay towards the maximum of far boundary.

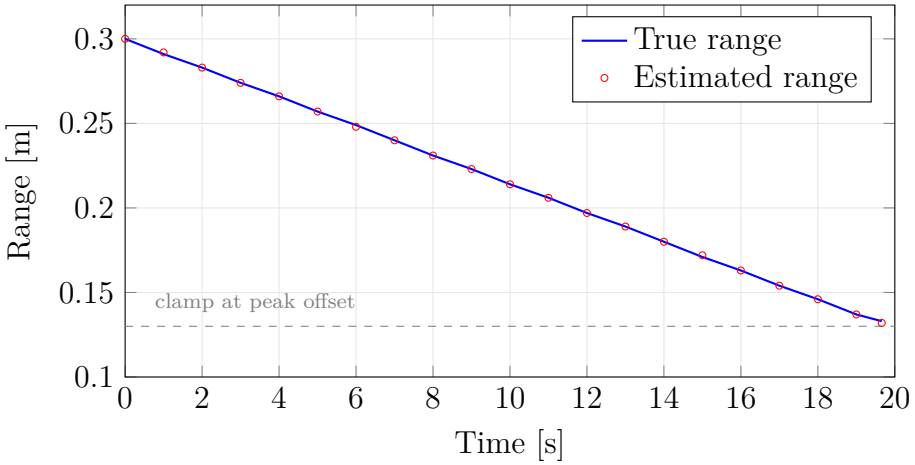

Figure 4.3: Estimated range against the true range in simulation, as the virtual leader approaches the follower. The virtual leader follows a prescribed trajectory, so the true range is known at every instant. The estimated range, recovered by inverting the wake profile, follows the true range with an error below one millimetre on the monotone side of the peak.

4.6 Relative position estimation in polar coordinates

The synthesis of relative bearing and radial range establishes the precise planar coordinates of the leading vehicle. The instantaneous spatial measurement, derived natively in polar coordinates $(\hat{r}, \hat{\theta})$, is kinematically transformed into a relative Cartesian position vector within the follower's body-fixed frame:

$$z_x = \hat{r}\cos\hat{\theta}, \qquad z_y = \hat{r}\sin\hat{\theta}, \tag{4.5}$$

the raw Cartesian pose is corrupted by significant high-frequency content introduced by the residuals, and is only available sporadically during wake periods. As such, the state estimator must deal with both residual noise in the raw Cartesian pose and the loss of signal during wake detection periods.
To achieve state smoothing, the architecture employs a discrete, first-order recursive exponential smoothing filter (an alpha filter) applied independently to the Cartesian planar axes:

$$\hat{x} \leftarrow \hat{x} + \alpha\,(z_x - \hat{x}), \qquad \hat{y} \leftarrow \hat{y} + \alpha\,(z_y - \hat{y}). \tag{4.6}$$

The filter is set to a fixed gain, which has been optimized to give the ideal trade-off between dynamic follow bandwidth, and effective noise filtering. This ensures that the estimate is following the leader exactly while suppressing any

jitter from one measurement to the next. It is noted that the recursive update only occurs at active detection instants. When signal is lost, the estimator will perform a zero-order hold of the last valid estimate, thus while there is no signal, the follower will be driving towards the last known point, rather than initialising a newly arbitrary origin, and executing uncontrolled kinematic trajectories. The zero-order hold is the most secure approach to this channel as the time duration of signal loss is greatly in excess of the time intervals between physical leader movements.

4.7 Filtering strategies: alpha filter, alpha-beta-gamma predictor and Kalman-based alternatives

The choice of the filtering architecture to meet the requirements of the relative position estimation task is not evident. The actual intermittency of the channel in the considered case leads to better performance of simplified structures in comparison with more complex ones. In this section, the first-order alpha filter will be evaluated relative to two standard options: discrete alpha-beta-gamma predictor and constant-velocity Kalman filter. The alpha filter itself is a purely positional estimator, which means that it uses a low-pass filtering process on the position measurements and a zero-order hold during intervals between samples. This filter does not assume any particular form of the target's motion and does not include parameters for the hold function during integration. The alpha-beta-gamma filter extends the state vector to include velocity and acceleration. It can be shown that a rigid constant-acceleration model corresponds to the propagation of the state estimator. Both these estimators are fixed-gain structures; thus, the stability regions for them are found by analyzing the closed-loop error equations using the Routh-Hurwitz criterion. Meanwhile, the constant-velocity Kalman filter itself selects a particular gain matrix (from an infinite number of stable ones) that minimizes the error covariance. However, its choice is still based on the specific model of the target's motion and corresponding noise characteristics.

Classical estimation theory provides a mathematical equivalence relating the Kalman optimum to fixed-gain structures. For linear time invariant systems operating in steady state, the Kalman gain approaches a constant matrix. In particular, a constant-velocity kinematic model is related to a discrete alpha-beta filter structure by:

$$\mathbf{K}_\infty = \begin{bmatrix} K_p \\ K_v \end{bmatrix} = \begin{bmatrix} \alpha \\ \beta/\Delta T \end{bmatrix}, \tag{4.7}$$

which have a fixed gain coefficient that are determined deterministically using a

single tracking index, which is defined as the ratio of the state uncertainty due to unmodeled target maneuvers to the covariance of the measurement noise [40]. As such, the Kalman filter does not propose a class of unique estimators, but rather automatically computes the theoretically optimal members of the fixed-gain family of estimators that are consistent with its statistical assumptions.

However, this mathematical optimality is strictly conditional. The Kalman filter is optimal only when operating with a linear kinematic model that correctly describes the target dynamics (with a zero-mean white Gaussian noise having covariance matrices known exactly) and deterministically regular measurements update. Its optimality is a mathematical theorem proven under the idealized model only; violating these strict constraints renders the Kalman filter no longer optimal in any mathematical sense. Inside the airframe channel, all the assumptions made in this work are violated. The constant velocity assumption is violated because the target is maneuvering in and out of a local region of turbulence. The noise covariance matrices are unknown at the start of the simulation, while the measurement noise is colored by the attitude dynamics of the closed-loop system. Moreover, the measurements arrive in bursts of high frequency separated by long time intervals. Overall, the conditions are not favorable to an optimal Kalman filter, and the result is a near-optimal (or fixed) gain filter. Its velocity state, tightly coupled to a violated kinematic assumption, erratically overreacts to the intermittent bursts of noise, rather than effectively providing temporal smoothing. The basic alpha filter, excluding the velocity state, deliberately avoids any predictive open-loop extrapolation, being structurally immune to such failure modes, simply retaining the last known spatial position across the measurement gaps. It is superior not because of a more sophisticated estimation algorithm, but because of fewer rigid parametric assumptions that may be violated in practice. While downstream, a derivative relative velocity is computed, it is strictly restricted to providing only transient damping feedback to the flight controller, explicitly avoiding its utilization as an open-loop predictor of the leader's spatial trajectory.

4.8 Detection-rate and update-rate tradeoffs

Although the estimation pipeline runs synchronously with the 200 Hz microprocessor attitude loop, it does not update the spatial position continuously in each computational step. A valid state update only occurs in the presence of a positive OS-CFAR detection, and this only physically happens when the trailing UAV flies through the local aerodynamic wake of the leader. As such, the effective state update rate depends on relative encounter geometry and kinematics, and not the underlying control loop frequency. During intervals in between encounters, the follower is out of the sensing volume of the leader and the estimator defaults to

the zero-order hold mechanism. As such, the effective spatial update rate is on the order of a few hertz, several orders of magnitude lower than the attitude loop.
The first fundamental trade-off, which was related to the inherent limitations of the system, is between the detection probability and the false-alarm rate determined by the threshold. The latter is controlled by the thresholding mechanism, which is responsible for deciding whether a specific event is a target or noise. Thus, reducing it results in increasing the number of positive detections and, consequently, the update rate, but at the same time, the false-alarm rate grows significantly since the threshold is lowered. In a closed-loop pursuit, the false alarm represents an inherent problem since it introduces a fictitious target which draws the follower aircraft toward a false position. Thus, the OS-CFAR architecture tackles this issue by keeping the false-alarm rate fixed at approximately one percent while the detection probability is adapted to the geometry of the encounter.

The second type of trade-off is between the system's responsiveness and estimation accuracy, dictated by exponential filter gain and debounce window size. High filter gain with short debounce interval tracks the trajectory closely, but preserves a much higher amount of high-frequency noise and discontinuities. On the opposite edge, a low filter gain combined with a large debounce interval smooths out the estimate significantly and introduces discontinuities into the state derivative, but helps to maintain continuity of events. The estimation pipeline described above deliberately favors stability, tending to the slow side of the spectrum. That is a physically correct approach for intermittent sensor signals where a responsive (but discontinuous) estimate would cause excessive control effort with no basis in the real physical state.
The ultimate limitation to the aggressiveness of the closed-loop pursuit is the kinematic constraint imposed by the system's overall dynamics. Due to the asynchronous nature of the updates (at the order of only a few Hz) and the static holding of these updates in the absence of incoming messages, the flight controller must regulate the follower relative velocity to avoid excessive overshoots of the narrow aerodynamic wake between successive measurements. This comprises the fundamental limitation behind the performance of the passive aerodynamic channel and emerges as the inevitable mathematical consequence of the annular, non-continuous spatial topology of the aerodynamic wake formally introduced in Chapter 3. A dedicated pursuit control loop designed to operate robustly in this restricted parameter space is then thoroughly developed in Chapter 5.

Chapter 5

Closed-Loop Pursuit Based on Implicit Sensing

Chapter 4 established the estimation pipeline required to reconstruct the relative spatial leader's position from the aerodynamic wake, whereas this chapter is dedicated to building the state estimate to synthesize a closed-loop pursuit architecture. The latter employs the former to autonomously regulate the follower's kinematics, such that it tracks the leader and maintains a prescribed inter-vehicle standoff distance, solely through the implicit aerodynamic channel and without utilizing any exteroceptive sensory data. The proposed control architecture relies on the classical pure pursuit guidance law. It is intentionally streamlined, as a lower-order control paradigm provides improved stability margins when operating on highly intermittent and stochastic measurement channels than more complex control topologies.

The chapter begins by situating the controller in the context of classical guidance laws before proceeding to define the control objective. A pursuit law expressed in polar coordinates is then developed, with the interaction with the state estimator investigated, together with the closed-loop stability properties of the resulting configuration. The robustness properties of the controller to dropouts are also evaluated. Finally, the implementation of the control law on the follower UAV is described. The controller is then experimentally validated using the dynamic flight tests described in Chapter 7, since a closed-loop pursuit requires true vehicle dynamics, unlike the kinematic simulation considered in Chapter 4

5.1 Control objectives for follower alignment and distance regulation

The Follower UAV has two main control objectives. These two objectives correspond to the two spatial states determined by the aerodynamic sensing channel. The first objective concerns angular alignment as the follower must regulate its heading so that its velocity vector aligns with the inter-vehicle line-of-sight (LOS) using the relative bearing estimate. The second objective concerns the regulation of the longitudinal distance as the follower must maintain a specified radial separation from the leading vehicle. This is to avoid collisions while keeping the target within the functional aerodynamic sensing volume using the radial range estimate.

These two objectives are the basis of classical guidance laws, which are a theoretical framework initially designed for missile interception and later adapted for the field of autonomous robotics [41]. A guidance law serves as a feedback mechanism which guides a pursuer towards a target according to their relative geometric kinematics. Classical guidance laws can be classified according to the geometric parameters they utilize; pure pursuit rigidly aligns the pursuer velocity vector along the instantaneous target LOS, and only makes use of the relative bearing. Meanwhile, proportional navigation (PN) commands lateral accelerations which cancel out the LOS rotation rate, and thus leads the target rather than pursuing it directly, requiring accurate measurement of the LOS rotation rate [41, 42]. Pure pursuit is the optimal guidance architecture for this application, since the aerodynamic channel ensures a reliable bearing estimate, while only providing a spatially limited intermittent range measurement. Pure pursuit only requires the instantaneous angle of the LOS to define the steering vector, reducing the computational burden on the embedded microprocessor.

These objectives have differing degrees of theoretical complexity due to their differing levels of observability on the underlying state variables. As relative bearing is globally observable at any position in the turbulent wake, angular alignment can be continually commanded. Meanwhile, as shown in Section 3.6, the range is definitely less observable where the signal drops as the leader moves away, past the distance of maximum moment. Distance regulation is therefore well posed there, and the set distance must be chosen in this region, so that distance regulation is mathematically well posed only there. The reference setpoint distance must therefore ensure that the closed-loop equilibrium lies within this well-posed region, thus creating loop convergence.

Furthermore, these control objectives must take into account the inherent intermittency of the sensing modality: because the follower perceives the leader only while actively engaged in wake interactions, both alignment and distance regulation are based on sporadic, high-frequency bursts. As a result, the controller must

inherently maintain its dynamic commands through these signal dropouts, rather than exhibiting erratic transients, a constraint that fundamentally dictates the following architectural design.

5.2 Polar-coordinate pursuit controller

Now that we know the state is measured in native polar coordinates (bearing and range), it makes sense to define the controller also in terms of the polar frame. The state estimator provides the leader's position in Cartesian coordinates ($(x_{\text{rel}}, y_{\text{rel}})$), expressed in the follower's body frame. We can then extract the radial position (range) and the unit vector pointing in the direction of the leader, given by the position vector:

$$r_{\text{est}} = \sqrt{x_{\text{rel}}^2 + y_{\text{rel}}^2}, \qquad (u_x, u_y) = \left(\frac{x_{\text{rel}}}{r_{\text{est}}}, \frac{y_{\text{rel}}}{r_{\text{est}}}\right), \tag{5.1}$$

where (u_x, u_y) represents the directional heading. This geometric decoupling encapsulates the essence of the pure pursuit law: the directional vector (u_x, u_y) dictates the trajectory heading, while the scalar magnitude governs the pursuit velocity. The follower's velocity remains strictly aligned with the LOS to the Leader, thus satisfying the foundational defining property of pure pursuit [41].
The proposed architecture diverges from textbook pure pursuit specifically regarding velocity regulation. While classical pure pursuit operates at a constant velocity to guarantee target interception, the follower UAV must maintain a discrete standoff distance; consequently, the command velocity is dynamically regulated by the range error. Defining the reference setpoint as r_{target}, the range error and the resulting radial control effort are formulated as:

$$e_r = r_{\text{est}} - r_{\text{target}}, \qquad u_r = K_P\, e_r + K_D\, \dot{e}_r, \tag{5.2}$$

The ultimate commanded planar velocity vector is derived by projecting this scalar radial effort along the directional unit vector:

$$v_x = u_r\, u_x, \qquad v_y = u_r\, u_y. \tag{5.3}$$

If the distance between the agent and the target is higher than the set point, a positive e_r prompts a restoring force to bring the vehicle towards the leader, while a negative one causes the vehicle to reduce or reverse. At the set point distance, the nominal instruction becomes zero. The controller comprises a pursuit law on the angle and a Proportional-Derivative (PD) regulator in velocity, the configuration of which is depicted in Figure 5.1. For the experimental setup, the control gains are set to $K_P = 0.08\,\text{s}^{-1}$ and $K_D = 0$, and the commanded velocity is aggressively

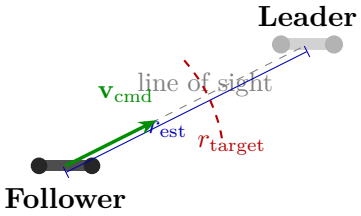

Figure 5.1: Geometry of the polar pursuit law. The follower commands a velocity $\mathbf{v}_{\text{cmd}}$ along the line of sight to the leader, enforcing the fundamental condition of pure pursuit. The scalar magnitude of that command is dictated by the differential between the estimated range r_{est} and the target setpoint r_{target}, ensuring the follower approaches when beyond the target distance and recedes when within it, ultimately stabilizing on the radial boundary defined by r_{target}.

saturated at $V_{\max} = 0.20\,\text{m}\,\text{s}^{-1}$ to make sure the vehicle's dynamics are not faster than the closed-loop bandwidth of the controller.

Although the derivative term is present, it does not add any benefit due to the noisy nature of the measurement channel (estimated range was intermittent and highly varying), as numerical differentiation amplifies the high frequency noises. While the controller applies a first order low pass filter to the error derivative to suppress this effect, the experiment has shown that differentiating an intermittent and noisy range reading adds more noise to the control effort, than suppressing the dynamics. Therefore, the controller acts as a purely proportional regulator, while retaining the derivative part for future use, in case the estimate range signal will be improved. Finally, to prevent the controller from inducing high-frequency actuator chatter due to imprecise estimation near the zero range, a dead zone of size 5 mm is introduced to the range error. The resulting control law, which includes both dead-zone and velocity saturation is depicted in Figure 5.2.

The angular alignment objective is inherently achieved by this very control formulation; there is no need to specify an independent heading loop. Since the scalar velocity command is continually projected along the updated directional unit vector, the follower's velocity will automatically be along the instantaneous relative bearing, actively turning as it updates. Multidirectional alignment is therefore implicitly achieved in the polar formulation. This is the primary computational economy that the pure pursuit provides; a single control law accomplishes both trajectory orientation, and, in this modified architecture, the control of distance.

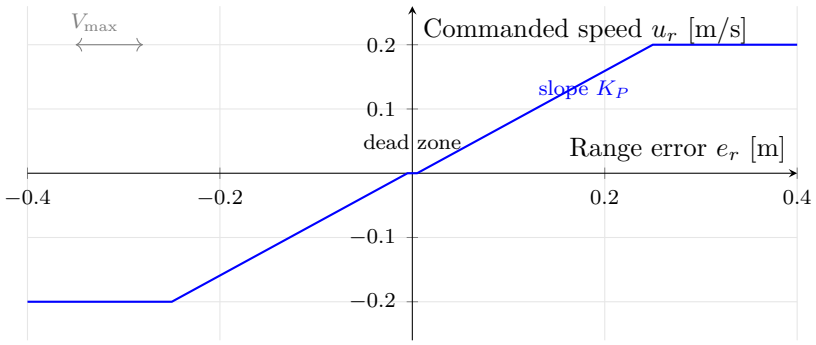

Figure 5.2: The radial velocity command parameterized as a function of the range error. A defined dead zone in the vicinity of zero error suppresses the command output, preventing the follower from erroneously reacting to stochastic estimation noise at the equilibrium setpoint. Outside this dead zone, the command scales proportionally with the error via the gain coefficient K_P, generating restorative effort toward the set distance. The command is strictly bounded at $\pm V_{\max}$, restricting the follower's kinematics to velocities safely within the tracking capabilities of the intermittent measurement updates.

5.3 Coupling between estimator dynamics and control response

The closed-loop controller is actuated by the estimated state resulting from the filtering process, not the true instantaneous relative position. Therefore, the inherent dynamics of the estimator dictate the total closed-loop response. The first order α filter introduced in Section 4.6 smooths the spatial position; thus trading-off tracking bandwidth for random noise rejection, and performs a zero order hold on the signal during temporal voids. Since the control law is acted upon this smoothed and latched estimate, the estimator and controller can be considered as a single highly coupled dynamic system as shown in Figure 5.3.

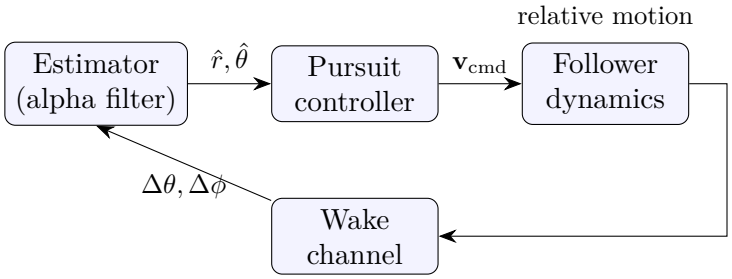

Figure 5.3: Block diagram of the sensing-control loop architecture. The aerodynamic wake induces attitude residuals on the follower airframe; the estimator maps these residuals into a relative spatial position; the pursuit controller subsequently generates a regulated velocity command; and the resulting kinematic motion alters the relative spatial geometry, thereby feeding back into the wake perturbation. Positioned in series within the closed loop, the phase lag and hold characteristics of the estimator directly modulate the resulting control response.

This dynamic coupling places an upper limit on the allowable control bandwidth. The state estimate always lags the true physical position by the settling time of the alpha filter, and is held fixed during detection dropouts. A too aggressive proportional control gain will cause the system to actively slew based upon the old, lagging state estimate, resulting in overshoot and oscillations in the kinematics. The gain of the proportional controller must be low enough that the resulting closed-loop dynamics are slow compared to the estimator phase lag and hold time. The selected gain of $K_P = 0.08\,\mathrm{s}^{-1}$ indicates a slow convergence to the reference distance spanning several seconds instead of high-bandwidth transient response; this keeps the kinematic commands coherent with an estimator which updates every few hertz.

The operational dead zone and the restrictive velocity saturation cap combine to this stabilizing effect on the situation. The dead zone eliminates the remainder of

the random noise that has survived the alpha filter, thereby preventing the airframe from responding to high-frequency fluctuations. Meanwhile, the velocity saturation restricts the maximum distance the follower can travel on a single held estimate when it is sliding on a single estimate, preventing such dangerous maneuvers during long tracking intervals. Ultimately, the estimator brings smoothing and latching, and the controller brings slow saturation, thus ensuring stability in closed-loop despite an intermittent sensor input.

5.4 Stability considerations for the sensing-control loop

Considered in isolation while its tracking loop is continually engaged, the range loop acts like a proportional controller operating on a range error. The analysis of the stability is then trivial. Since the follower's velocity is the first time derivative of its position, the range function responds to the radial velocity command as a first-order integrator along the line-of-sight; $\dot{r} = -u_r$, where the negative sign indicates that a positive restorative command to reduce the absolute value of the range requires a radial velocity command towards the leader. Incorporating the proportional control law $u_r = K_P\, e_r$ and defining the error as $e_r = r - r_{\text{target}}$, the closed-loop dynamics resolve to a first-order differential equation:

$$\dot{e}_r = -K_P\, e_r, \tag{5.4}$$

whose time-domain solution decays exponentially as $e_r(t) = e_r(0)\, e^{-K_P t}$, the range error therefore converges asymptotically to zero with a time constant of $\tau = 1/K_P$, ensuring the follower stabilizes at the target set-point without kinematic overshoot for any strictly positive gain coefficient. Parameterized at $K_P = 0.08\,\mathrm{s}^{-1}$, the system exhibits a conservatively slow time constant of approximately 12 s, intentionally tuned to accommodate a sensory measurement updating at merely a few hertz.
However, in two ways this idealized first-order response can be compromised, and the architecture explicitly addresses both. First, there is the issue of systemic phase delay: the phase lag of the alpha filter, combined with the asynchronous and intermittent nature of the measurements, results in a critical transport delay in the loop from true range to control input. In any feedback system, a large uncompensated transport delay will convert restorative exponential decay to divergent oscillation; the heavily attenuated proportional gain keeps the closed-loop time constant sufficiently large compared to this delay to preserve a consistent decaying response. This principle also governs the exclusion of the derivative gain: when applied to a time-lagged signal that is characterized by a lot of noise, the derivative term causes undesirable phase blurring at a much greater rate than it provides useful damping effects. The second effect that PC covers is the non-monotonicity of

the range signal in question. The first-principles reasoning assumes that lessening the range is supposed to result in lessening the error. This only applies when the signal decreases as the leader moves away from its initial position. For shorter distances from the leader, the effect is opposite: since the signal increases as the leader approaches, it may happen that a command supposed to mitigate the range will only cause an increase in the error. The clamp of Section 4.5 holds the estimate beyond the distance of maximum moment, so the controller never enters the region where the signal is ambiguous and the loop stays well behaved.
While a formal Lyapunov proof of stability for the combined, non-linear estimator-controller equations subject to stochastic intermittency, that is, a hybrid system composed of continuous decay punctuated by discrete hold intervals, is beyond the scope of this work, the engineering motivation behind this approach is strong and is grounded in classical results: pure pursuit is stable to slowly maneuvering targets and destabilizing against high-bandwidth evasions [41]. A low-gain first-order distance loop, clamped to the bounds of velocity saturation and operating within its monotonic regime, is structurally stable by design, a claim which is borne out by the flight data shown in Chapter 7.

5.5 Robustness to intermittent detections and temporary loss of signal

The critical property of the implicit aerodynamic channel is that the leading vehicle is visible only when the following UAV passes through the wake of the leading one. As a result, the measurements are lost during the time intervals when the follower is moving close to the wake boundary or the leader takes evasive maneuvers. The hold logic in the state estimator is used to keep the follower's dynamics bounded during these time intervals.
Upon loss of the aerodynamic signal, the estimator latches on its most recent valid position, not to arbitrarily reset the origin or extrapolate open-loop kinematics (as formalized in Section 4.6). As a consequence, the controller continues to steer the airframe towards the last known location of the leader during the time void, as opposed to simply ceasing pursuit or executing an uncommanded turn. Since most signal dropouts are of a transient nature, this zero-order hold keeps the follower on an intercept course, guaranteeing rapid regain of the wake boundary once the signal is back. Thus, flight dynamics of pursuit evolve as a continuous approach-and-hold limit cycle: the follower actively converges on the leader during positive detections, while coasting smoothly along the latched trajectory during transients, and immediately reverting to active regulation upon signal return. Two particular parameters of the design imposed limitations upon the kinetic danger the long signal voids presented. The aggressive speed saturation restricted the

maximum distance the follower could fall from the line on a blind run with a latched estimate, denying a stale position the opportunity to create a catastrophe. Meanwhile the dead zone ceased further control attempts as soon as the follower neared the target condition, leaving the airframe inert during the many micro-dropouts incident to approaching the setpoint. Together, these features guaranteed that an intermittent loss of signal would manifest as a mere interruption of active control, rather than a malfunction of the pursuit mechanism. Nevertheless, the theoretical limits of this robustness must be acknowledged, and this is in line with known vulnerabilities of pure pursuit in the guidance literature. Pure pursuit demonstrates effective tracking against low-bandwidth targets but is particularly poor against high-velocity maneuvers, fundamentally because it converges on the target's instantaneous coordinate and does not anticipate its future trajectory [41, 42]. In the context of implicit wake sensing, this is magnified into a severe limitation: latched state estimates would be entirely inappropriate if the leading vehicle permanently exited the wake for an extended period of time. Following this wrong state estimate, the follower vehicle would blindly steer towards coordinate data that are no longer valid, having no other sensor to detect that the target has already departed. Thus, closed-loop pursuit is only mathematically sound if signal dropouts are temporally brief compared to the characteristic maneuvering time of the leading vehicle; in fact, this is exactly the slowly-varying kinematic regime fundamentally assumed by the aerodynamic channel model in Chapter 3. Highly aggressive or distant target maneuvers necessarily exceed the theoretical tracking capability of this modality. Furthermore, implementing an advanced predictive law such as proportional navigation is mathematically impossible, since the necessary calculation of the LOS rate is not supported by a sensing channel that only gives intermittent bearing information and limited range data.

5.6 Practical implementation on the follower UAV

The pursuit controller is implemented directly on the follower's attitude-control and estimation embedded flight-control microprocessor, operating within the 200 Hz attitude-control and estimation loop. The algorithm has negligible computational overhead, requiring only a few simple operations: at each time step, the controller performs a vector norm operation, a four-quadrant arctangent, a proportional scaling, and a velocity saturation. All calculations are performed in single-precision floating point to respect the embedded processor's real-time capabilities. The controller also has minimal memory requirements, storing only the previous range error and its filtered estimate.

The output of the pursuit controller is a commanded planar velocity in the follower's local body frame, which is injected as a setpoint into the vehicle's preexisting

cascaded velocity control loops. As a result, the pursuit controller acts purely as a high-level guidance supervisor, constantly determining the optimal spatial trajectory, while the lower level inner loops take care of the involved multivariable rotor actuation to actually get there, see Figure 5.3. The modular layering thus decouples the pursuit logic from the specific high-frequency airframe dynamics, allowing the architecture to be readily portable onto any rotary-wing platform that is able to accept planar velocity setpoints.

The set distance is the only parameter that must be calibrated against the channel. It is located in the region where the wake signal decreases with distance, beyond the clamp defined in Section 4.5, so that the closed-loop equilibrium lies where the range can be recovered unambiguously. Below the offset of maximum moment the signal maps to two distances, and the range cannot be resolved; beyond that offset each signal level corresponds to a single distance. The set distance is therefore placed a few centimetres beyond the offset of maximum moment, near enough that the follower remains within the wake and far enough that each measured signal corresponds to one distance. This choice gives the controller a well-posed equilibrium, and the pursuit behavior the system produces in flight is analysed in Chapter 7.

Chapter 6

Interface with the Return Communication Channel

The preceding chapters describe only the forward sensing channel: how the following quadrotor senses the wake of the leading one, estimates its position, and subsequently follows along an autopilot-like closed loop. This analysis is entirely self-contained and focuses on the problem of sensing and control in isolation. That is, however, only one half of a bidirectional communication process: a topic that is the main focus of this chapter. In particular, this paper details how the leader-to-follower sensing channel described in this work interfaces with the follower-to-leader gesture-based return channel developed in the companion thesis. The inner workings of the return channel, such as the symbolic alphabet, decoding process, and communication protocol, are described in more detail in the companion work, while this paper details the intersections between the two channels in terms of the mechanics of the forward estimation algorithm.

6.1 Role of follower-to-leader feedback in the shared bidirectional system

The forward and return channels carry a fundamental asymmetry that underpins their integration. The macroscopic formation navigates in relation to an absolute spatial reference coordinate that it has predefined before takeoff and that only the global leader knows. By having the leader independently navigate towards this reference, the follower can only approach the global objective by constantly referring to the leader's condition. As a result, the dynamic path from the absolute reference to the follower goes through the aerodynamic wake between the two vehicles, the forward channel being the critical second link in this chain of dependencies.

This topological dependency makes the forward sensing channel structurally indispensable and the only means by which the follower can integrate into the formation architecture. If the forward link fails, the follower is not only denied real-time position updates but is completely disconnected from the overall formation objective as there is no alternate kinematic path to access it. In other words, the forward sensing channel dictates the physical reachability of the global objective and, by virtue of being the critical load-bearing element of the bidirectional architecture, dictates the success of every formation maneuver.
Conversely, the return channel is a strictly complementary supervisory channel due to its nature of disseminating discrete state information from the follower to the leader, which cannot affect the global convergence as the reference is independent. Instead, it controls the behavior of fault-tolerance mechanisms in the presence of nominal tracking violation. In case of a critical fault or a follower drifting outside the bounds of a wake, an uninformed leader would continue to perform towards a global objective, causing the formation to diverge. The return channel reduces this risk by allowing the follower to broadcast its state explicitly allowing the leader to stop advancing rather than blindly advancing. Thus, while the forward channel dictates the theoretical possibility of a mission, the return channel guarantees the practical safety. Neither channel overrides the other, but their interface ensures that the failure in the passive sensing pipeline is safe, rather than silent.

6.2 Minimal description of gesture-based acknowledgements

Defining the operational interface requires only a cursory discussion of the return channel; a detailed exposition appears in the companion dissertation. The follower communicates with the leader by means of gestures, devoid of RF links, using the same forward sensing mechanism: it superimposes a high-frequency vertical motion, a periodic "bobbing," onto its forward speed. The leader tracks this maneuver by virtue of the perturbations in the follower's wake field, thereby allowing the follower to send the leader a symbolic message in the form of a periodic sequence of motions differing in timing and presence; the message includes a heartbeat to signify health, acknowledgments, and so on.
Just two aspects of this return architecture's design have a direct impact on the properties of the bidirectional interface. First, note that the return architecture encodes a gesture as a vertical motion, while the forward channel encodes gestures as planar horizontal motions (see Section 6.4 for details). Second, since the aerodynamic disturbance travels exclusively downward, the trailing vehicle must perform maneuvers above the leader in order to perform a gesture. Meanwhile, the follower must stay below the leader in order to perform forward sensing. Thus, the

forward channel and return architecture require the follower to occupy opposite vertical volumes relative to the leader.

6.3 Interaction between leader sensing and return-channel activity

The forward sensing channel is active while the follower is in its nominal position following the leader. In this sensing configuration, the follower manages to intersect the wake, determine the leader's position in space relative to itself, and actively pursue it. In the return channel, however, the follower only needs to issue a temporary reply configuration, in which it performs a rapid vertical roll to position itself above the leader to send its gestural message. Governed by the unidirectional nature of the turbulent wake, these two configurations are mutually exclusive. Only the lower vehicle could implicitly sense the upper one.

This mutual exclusion has an implication for the forward estimation pipeline: during a gesture from the superior position, the follower loses its sensing modality. As such, the forward channel provides zero update in position for the duration of the reply maneuver; the follower's estimator is forced to do a zero-order hold on the state estimate. This results in the target's relative coordinates essentially being frozen at their final valid values, just as described in Section 4.6 for the case of stochastic signal dropouts. The reply configuration merely represents a deterministic extension of this zero-order hold period during which the state estimate is latched. Since the baseline estimator pipeline was already synthesized to deal with higher-level signal intermittency, the deliberate reply gaps are simply absorbed by the same hold logic that is used to manage natural wake dropouts.

The fundamental difference is that a reply-induced dropout is deterministically triggered by the follower's high-level logic rather than stochastically imposed by the environment wake drift, and its temporal extent is strictly bounded by the time of the gesture encoding and role-swap kinematics. From the viewpoint of the forward channel, this interaction is benign: active sensing is suspended for the brief period, the kinematic estimate is safely held, and the closed-loop sensing instantly resumes as soon as the follower drops back to the inferior position. The associated tracking costs and the system limitations on reply frequency are thoroughly discussed in Section 6.5.

6.4 Avoiding interference between forward sensing and feedback gestures

Given that both communication channels simultaneously exploit the same physical medium, namely the local aerodynamic flow field, a consideration of potential cross-channel interference is of utmost importance. Thus, by design, formation flight under nominal conditions avoids systematic interference by keeping the respective actuation axes strictly orthogonal. The forward channel encodes spatial information in the planar, horizontal motion of the leader, such that an arbitrary variation in the horizontal planar velocity physically alters the resulting downwash, which the follower decodes in the orthogonal attitude residuals. Meanwhile, the return channel encodes information solely in the vertical motion of the follower, superimposing a high-frequency oscillation on the altitude vector, which the leader perceives in vertical velocity perturbations only. Therefore, the state variables, whose variation encodes information in one channel, are orthogonal to the state variables, whose variation is monitored in the other channel. A vertical gesture does not induce significant variations in the horizontal planar velocity, and a horizontal pursuit command will not generate vertical oscillations at the designated gesture frequency.

This geometric orthogonality is further assisted by spectral separation, since the return gestures are designed to work in a narrow, high-frequency band specifically located above the low-bandwidth dynamics that make up standard formation keeping. The forward pursuit commands, however, are low-frequency motions. So, if cross-axial aerodynamic coupling did occur, the resulting interference would be outside the narrow frequency band tracked by the return-channel decoder, but would be completely indistinguishable from low-frequency drifts in the forward channel. Thus, the two channels are completely transparent to each other during steady-state flight.

The only exception to this otherwise transparent operation is the physical role-swap maneuver itself, which is a mutual exclusion rather than a signal interference. During the inversion, the physical ordering of the UAVs is flipped: while the follower is on top of the leader, the forward channel is not just attenuated, but physically impossible to follow, as the wake cannot propagate upwards. The two sensing channels are never concurrent: the instantaneous vertical geometry dictates which one has physical validity. As established in Section 6.3, this total preemption of signal is acceptable to the forward channel, as the interruption is deterministic and self-terminating. The critical sensing link of the follower to the global reference is re-established, cyclically, after every discrete gesture transmission.

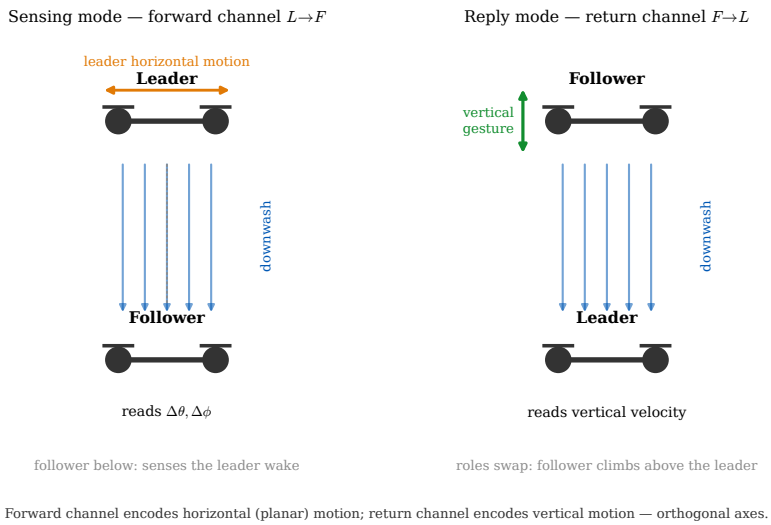

Figure 6.1: Geometry of the vertical role-swap between the two channels. In sensing mode (left) the follower stays below the leader and reads the wake-induced attitude residuals, while the leader encodes information through its horizontal motion. In reply mode (right) the follower climbs above the leader and superimposes a vertical oscillation, which the leader decodes as a vertical-velocity perturbation. The two channels act on orthogonal axes and never occupy the same vertical configuration.

6.5 System-level timing constraints from the follower perspective

The inevitable alternation between active sensing and gesture transmission dictates a rigid temporal budget at the level of systems, mostly from the perspective of the forward channel. The guiding principle is a zero-sum temporal trade-off; the time employed to transmit a reply is time directly subtracted from time spent in active state estimation. Since any gesture necessitates a full vertical role-swap that completely preempts the channel of forward sensing, the overall formation controller must strictly ration reply events against the consequent degradation in spatial tracking precision.

This macroscopic temporal constraint is completely separate and much more stringent than the high-frequency timing tolerances necessary simply to be able to decode symbols at the leader's decoder. Even perfect decoding and zero errors impose a huge penalty in lost sensing time; a follower programmed to send continuous telemetry would spend the vast majority of its flight envelope above the leader and thus be incapable of tracking the global objective. Thus, the forward

channel determines the absolute theoretical maximum for the allowable reply frequency: transmissions are sparse enough to ensure that the formation stays in the nominal sensing configuration for the majority of the mission, and that the target's position is frequently enough sampled so that pursuit errors remain bounded. The allocated multi-second time slot for the return channel's periodic heartbeat is deliberately chosen to meet this constraint, maximizing active sensing time and treating transmission of gestures as a rare and temporary anomaly.
Furthermore, the reply maneuver itself introduces an appreciable amount of its own non-determinism into the forward channel. For the baseline aerodynamic wake, Chapter 7 established that detection is in bursts with micro-dropouts on the scale of tens of milliseconds in between, where tracking is held, in order to maintain an estimate. A reply maneuver will involve a macro-dropout of significantly greater duration, because the state estimator is held throughout. Unmeasured kinematics of the target during the reply are thus only known upon return of the sensing topology. The duration of a reply maneuver must therefore be kept short enough, so that the leader cannot acquire a spatial deviation that is destabilizing while the follower flies blind. This is precisely the constraint of slowly-varying kinematics, used to design the forward channel. Pursuit remains stable, as long as the leader is maneuvering slowly in comparison to the effective spatial sampling rate, and a reply maneuver is always tolerable, as long as it lasts negligibly in comparison to the characteristic correlation time of the leader. The return channel operates well in conjunction with the forward channel at the only operating band of low-bandwidth kinematics, and falls away in either case.

Chapter 7

Experimental Validation of Leader Sensing and Pursuit

The previous chapters were focused on the design and modeling of the implicit communication channel, the sensing pipeline, the pursuit controller and the validation of the pipeline on simulation, considering therefore a known and predictable Leader's trajectory. Contrary, this chapter will address the testing of the whole pipeline described in the previous chapters on a real vehicle: the Parrot Mambo Minidrone. The main purpose of this work is to analyze if an implicit communication channel can be effectively implemented on a real drone, and therefore consider the hardware limitations. However, since real flights do not provide any external frame of reference to estimate the true relative position, the focus of this chapter will be on the internal validation and analysis of the results.
The chapter follows the pipeline from the ground up. Section 7.1 describes the setup and the flight campaign. Section 7.2 validates detection against the noise floor. Section 7.3 tests range and bearing while the leader hovers. Section 7.4 covers diagonal and moving-leader approaches. Section 7.5 characterizes the detection region and its blind zones from the flight data and the anemometric field. Section 7.6 reports the closed-loop pursuit. Section 7.7 concerns the interaction with the return channel. Section 7.8 discusses the limits of the experiments.

7.1 Experimental setup and instrumentation

The trailing vehicle is a Parrot Mambo nano-quadrotor weighing approximately 63 g with an X-frame design, half-wheelbase of 0.06 m and rotor radius of 0.016,5 m that uses only the proprioceptive inertial measurement unit (IMU) of its stock flight controller, without any additional exteroceptive sensors. As a result, the state of the leader is inferred purely from the attitude estimate that is continuously calculated

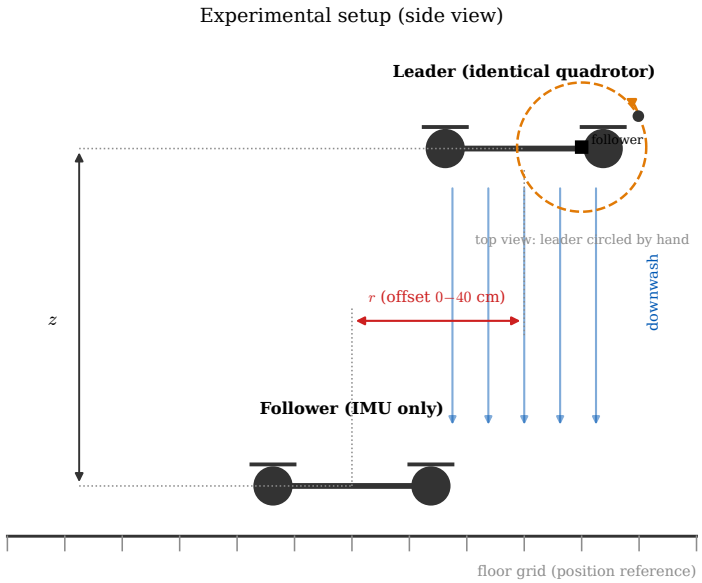

Figure 7.1: Experimental setup. The follower nano-quadrotor relies only on its onboard IMU; an identical leader hovers above it over a floor grid used as position reference. For the static tests the leader is placed at fixed lateral offsets r from 0 cm to 40 cm; for the dynamic tests it is moved by hand along a low-speed circular path around the follower (top-view inset).

for the purpose of baseline flight stabilization by the flight-control algorithm running on the vehicle. The detection, estimation (described in Chapter 4) and pursuit control (described in Chapter 5) algorithms therefore run on the embedded flight-control microprocessor at 200 Hz and are compiled in single-precision floating-point arithmetic for the target real-time embedded processor, flashing the resulting binary over Bluetooth.

The leading vehicle was a quadrotor, which was identical in structure to the follower. For the purpose of testing the algorithms at fixed offsets, the leader was kept stationary and located at specific coordinates above the floor grid. The leader was placed at different positions above the grid so that it hovered laterally offset from the follower. For the purpose of testing the algorithms at varying distances, the leader was moved by hand along a low-speed circular path around the follower. The rotors of the leader were kept running at all times so that it could induce the necessary disturbances on the follower. The experiments were conducted indoors using the floor grid, which served as a reference for localization. The overall arrangement is shown in Figure 7.1.

High-frequency telemetry is logged locally on board using the vehicle's native

post-flight logging infrastructure: this consists of the orthogonal pitch and roll attitude residuals, their computed Euclidean norm, the dynamic adaptive threshold, both the raw and temporally debounced detection flags, the estimated relative bearing, the reconstructed radial range, and the synthesized relative position of the leader. Thus, these logs constitute the basis for the experimental analysis in this chapter; furthermore, the validation is performed using two kinds of flight data: one is composed of eleven undisturbed leader-free hover flights, where the follower is operating in isolation without any aerodynamic wake; these are crucial to establish a baseline for the stochastic noise against which all wake detections are compared, and the other consists of a set of interaction flights with the leader present. A subset of these tests involves placing the leader at fixed, discrete, spatial offsets of nominally 0, 10, 20, 30, and 40 cm in order to characterize the signal amplitude as a function of geometry, whereas the rest are also made up of a set of extended dynamic flights, lasting on average 50 s, where the leader is circling the follower while the whole estimation and control are in a closed loop. These controlled-offset flights clearly quantify the signal response as a function of geometry, whereas the extended tests are used to verify the full-circle bearing observability, the monotonic range reconstruction, and the overall closed-loop stability of the pursuit.

7.2 Ground-level validation of wake detection

The main goal of this test is to check the follower's ability to sense the leader's wake, and to evaluate if the follower has picked up a significant signal beyond just the noise. To do this, the baseline of random noise is established by analyzing the flights without encounters; the eleven free flights produced a median of about 14 mrad and had a 99th percentile near 58 mrad. So that means that any significant deviation beyond this value is likely to be caused by the leader plane.

During active leader-present flights, the residual norm undergoes extreme transient excursions well above the noise floor: analyzing the controlled-offset flight telemetry shows the upper decile (the ninetieth percentile) of the norm magnitude is routinely between 60 and 90 milliradians. Even during moment of maximum proximity, peaks increase by an order of magnitude, reaching hundreds of milliradians. This implies the deterministic wake signal has a good Signal-to-Noise Ratio over the baseline. It is this difference between the tail distribution of the background distribution and the wake distribution the detection algorithm actively takes advantage of.

The dynamic adaptive threshold tracks the time varying noise floor and only allows detection of high amplitude events. Examination of the physical flight data shows that the Ordered-Statistic Constant False Alarm Rate thresholding, as described in Section 4.2, safely maintains the false-alarm probability below approximately one percent during leader-free flights, while still detecting the leader during all active

interaction flights. Thus, it effectively addresses the critical issue with the original fixed buffer thresholding algorithm. Examination of the same telemetry with the original gated four-sigma threshold shows that, during active wake following, the number of samples contributing to the statistics was limited to a few percent of the buffer during active wake following. This resulted in the buffer essentially freezing in place and causing the threshold to be permanently stuck at its lowest possible value for the duration of the flight. Replacing the conditional standard deviation estimate with an order statistic estimator eliminates this effect, since the extreme values no longer bias the mean and standard deviation estimates. The flight data serve as an excellent reference for the sanity check of any algorithm that tracks the statistics of the received signal. In particular, they demonstrate the failure mode of the original gated standard deviation approach, as well as the properties of the estimate with the conditionally gated standard deviation. Ultimately, the analysis confirms that the norm-based detector with the order-statistics-based dynamic threshold correctly and reliably detects the leading vehicle based on implicit wake sensing regardless of the approach angle and without the need for calibration due to the battery voltage effects.

7.3 Hover tests for range and angle estimation

Having empirically established the detection criteria, the subsequent experimental evaluations focus on the reconstruction of the relative bearing and radial range while the leading vehicle maintains prescribed hovering coordinates relative to the follower.

The empirical telemetry confirms successful recovery of the relative bearing over the entire azimuth domain. Across two long dynamic flights, the estimated bearing covers all twelve 30° sectors around the trailing UAV, each containing hundreds or thousands of valid samples. Figure 7.2 presents this coverage as a polar histogram; every sector is populated, with approximately equal distribution of samples. This means that the follower was able to successfully localize the leader from arbitrary approach directions; not only along the principle axes. This is an important improvement on the single-axis baseline proposed by Cerruti [5] which caused significant sensing voids at diagonal bearings due to the implicit use of an OR logic. The results therefore constitute empirical evidence that the two-axis vector norm and four-quadrant arctangent formulation from Section 4.4 successfully disambiguate the target bearing while completely avoiding geometric blind spots.

Similarly to the range angle, the radial distance is reconstructed as a highly repeatable, monotonic function of the spatial separation: analysis on the long range pursuit flights yields a tight negative correlation between the estimated range and the residual norm, with correlation coefficients of approximately -0.96 and -0.95

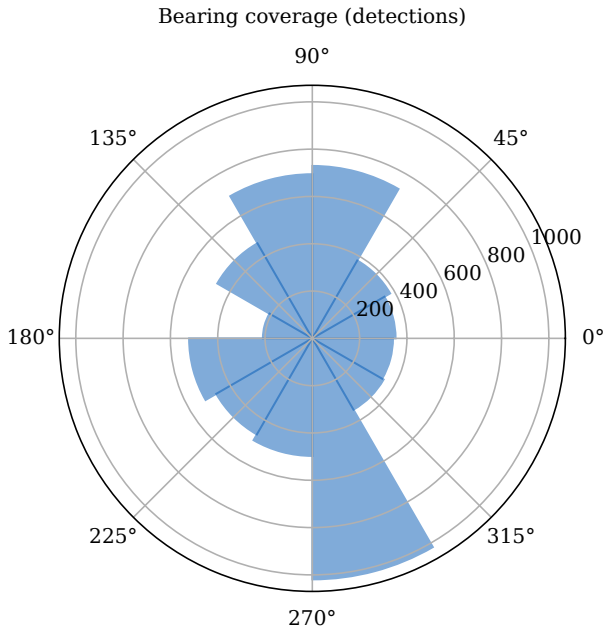

Figure 7.2: Angular coverage of the bearing estimate over a pursuit flight. Each wedge is one 30° sector, and its radius is the number of detections with the leader in that sector. The leader is localized in all twelve sectors around the follower, including the diagonal sectors near 45° where the single-axis baseline was blind. The bearing is therefore observable across the full circle.

across 1,703 and 4,320 individual detection samples, respectively. This confirms deterministically that an increased aerodynamic signal corresponds to a decreased radial range in a strictly monotonic manner. Moreover, the functional relationship adheres to a high repeatability across independent flight tests, thus establishing itself as a basic aerodynamic property. The state estimate, throughout the time of the maneuvers, remains within the defined operating interval; clamped below by the peak moment offset and above by the maximal sensing limit—ensuring the controller never reaches into the ambiguous inner space. As was formally established in Section 4.5, the strong correlation between the two is in fact a confirmation of internal algorithmic consistency rather than a validation of the range accuracy, due to the lack of an exogenous reference system to which the state estimation could be compared. Having said that, the determined range fully conforms to the theoretical expectations of the aerodynamic model: providing a stable, monotonous, and repeatable measure of the spatial proximity in any situation where the target enters the turbulent wake.

7.4 Diagonal approach and moving-leader scenarios

While the static hover cases are designed to restrict the lead aircraft to a constant bearing, the following dynamic tests incorporate a moving target. By circling the follower during the extended pursuit flights, the lead is able to scan the full range of bearings while varying the range, thus testing the estimation pipeline with a target that is operating as designed. These dynamic flights specifically include the diagonal approaches, which were the root cause of the failure in the single-axis baseline.

Subjected to a moving target, the estimation pipeline was able to simultaneously track relative bearing and range, enabling the full reconstruction of the leader's planar location. Figure 7.3 shows the estimated leader position in the follower body frame over one extended flight: the estimates cover the surrounding plane up to about 0.3 m from the follower on each axis, and none falls inside the clamp radius at the offset of maximum moment. As such, the trailing UAV reconstructs the planar dynamics of the leader, rather than simply detecting its presence, using only its internal inertial reference frame. The diagonal crossing paths, where the target was approaching the UAV from the sides of the field, were tracked without issue, as the distance's Euclidean norm retained the full value of the aerodynamic forcing regardless of the angle of approach.

The dynamic evaluations also highlight the fundamental intermittency which dictates the overall system architecture. Active detection is inherently transient, occurring only when the follower is in the wake of the leader. Over the course of the flight tests, the leading vehicle is detectable for around 42 % of the total flight time. These detections are also highly fragmented in time, being punctuated by large drops to zero signal; during the most dynamic flights, the median detection length is 70 ms, with a similar median period of time between detections. As a result, the follower UAV is passing through the downwash of the leader at a rate of several times per second in the most dynamic cases. The state estimator is capable of executing a zero-order hold without difficulty over these time scales. This intermittency is not a concern for the estimator, and represents the fundamental expected behavior of an implicit aerodynamic channel, which only has an effect within the confines of the downwash.

7.5 Spatial characterization of detection regions and blind zones

The topology of the region around the model was governed by the aerodynamic wake, with a central blind zone due to the free stream and an outer edge defined

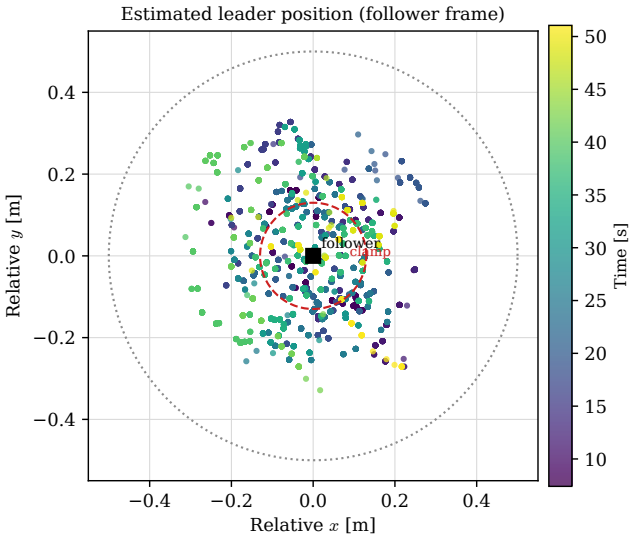

Figure 7.3: Estimated leader position in the follower body frame during one extended pursuit flight, coloured by time. The follower sits at the origin. Points populate the plane around it while respecting the two operational bounds: the inner dashed circle is the clamp at the offset of maximum moment (0.13 m), the outer dotted circle is the far-field edge (0.50 m). No estimate falls inside the clamp, confirming the follower stays in the observable annulus while tracking the leader around the full circle.

by the collapse of the flow field. This section characterizes this geometry using telemetry from the controlled-offset flights, alongside the empirical anemometric field characterized in Chapter 3.

The sensor signal's relationship to the radial offset is illustrated in Figure 7.4. At the evaluated controlled offsets, the top decile is about 68 mrad on the axis, peaks at around 90 mrad in an intermediate offset between 10 and 20 cm, and drops to 66 mrad at further distances. Compared to the established noise floor of 14 mrad, the top decile signal-to-noise ratio, therefore, varies between roughly 4.3 on the axis and 6.5 at the intermediate peak, equivalent to about 13 dB to 16 dB. Notably, the signal has a minimum on the centerline, not a peak. This is the tell-tale sign of a moment due to asymmetric aerodynamic loading on the follower by the downwash jet, rather than a scalar flow quantity such as speed. Indeed, on the centerline the aerodynamic load is distributed symmetrically around the follower, leading to little net moment and, thus, minimal sensitivity directly below the leader.
The detection boundary is inherently dictated by the anemometric flow field, as the

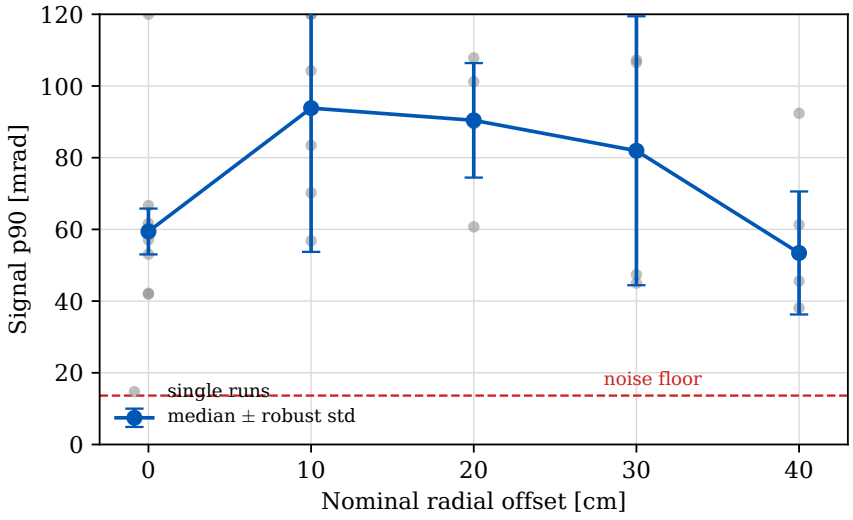

Figure 7.4: Wake signal magnitude parameterized against nominal radial offset, derived from the controlled-offset flight telemetry. The signal amplitude is defined as the ninetieth percentile of the residual norm; error bars denote the statistical spread across independent flight trials at each discrete offset. The signal is attenuated along the central axis, escalates to a spatial maximum at an intermediate offset near 10 cm to 20 cm, and subsequently decays at larger distances, while remaining robustly above the ambient noise floor throughout the evaluated domain. Because the prescribed offsets are nominal, this distribution visualizes the macroscopic topological shape of the signal rather than a metrologically exact calibration curve.

empirical measurements reported in Section 3.2 show that the structural collapse of the turbulent jet near a radial offset of 15 cm leads to the aerodynamic flow being attenuated below the instrumentation noise floor and no disturbance reaching the follower. Considering both phenomena, the effective detection region is found to be an annular volume, whose inner boundary is limited by the weak-signal zone around the central axis while its outer boundary is dictated by the structural collapse of the wake near 15 cm. The follower is able to extract a trackable signal solely within the annulus. In such manner, the sensing architecture features two blind zones stemming from different physical mechanisms. The inner blind zone, located along the centerline, is a consequence of geometry resulting in the aerodynamic loading being entirely balanced out with no induced moment. The outer blind zone, beyond the wake collapse, is a result of the aerodynamic flow failing to bridge the inter-vehicle separation. Both are direct consequences of the channel model derived in Chapter 3 and the empirical flight telemetry is found to agree with the theoretical expectations.

The statistical variance in the signal across independent flight trials, depicted by the error bars in Figure 7.4, provides additional physical intuition: this variance peaks at smaller nominal offsets and attenuates at progressively larger radial distances, in a manner that is geometrically explicable relative to the signal's spatial gradient. Near the steeper moment peak, minor manual positioning errors during testing introduce major deviations into the measured signal, while equivalent positioning errors along the lesser, farther-out flank produce lesser signal variations. Thus, this statistical variance is intrinsically a measure of the nominal precision of the manual spatial offset as much as it is of any intrinsic stochastic variance of the aerodynamic channel itself. The geometric sensitivity motivates an approach of interpreting the discrete measurements as a characterization of the signal's macroscopic shape, as opposed to attempting to extract a metrologically exact signal-versus-distance transfer function.

7.6 Closed-loop pursuit performance

The final stage of the experiment forms the closed-loop control system, and the follower UAV is able to pursue the leading aircraft autonomously. The extended flights are the result of the tracking controller, described in Chapter 5, which directly takes the estimated relative position in the plane and sets a restoring planar velocity towards the leader, and thus regulates the standoff distance. These extended flights therefore represent a full end-to-end test of the whole system, from initial wake detection to the kinematic response.

Strictly according to the theory, the pursuit in the closed loop leads to the continuous approach-and-hold limit cycle. The figure 7.5 shows the time course of the estimated radial range and the detecting state in one of the dynamic flights. While actively tracking the leader, the range estimates are continuously updated, and the follower actively flies to reach the set-point. On the loss of the signal, the instantaneous range estimate is latched with the zero-order hold and asymptotically relaxed to the maximum range; during this time, the follower glides on the latched range until the signal is acquired again. During the whole maneuver, the range estimate stays strictly within the theoretical boundary limits. The lower limit of the range clamp corresponds to the maximal induced moment near the 0.13 m, and the upper bound corresponds to the far-field edge around the 0.50 m. Moreover, the range estimate never goes beyond the inner clamp boundary, thus assuring that the follower never leaves the observable space even at the maximal approach. Both the 49 and 51 s long flights ended with the trailing UAV containing the estimated position of the leader within roughly the 0.3 m range from the center, continuously.

The overall system performance is fundamentally limited by the characteristics of the aerodynamic channel, as described in the previous chapters. The effective

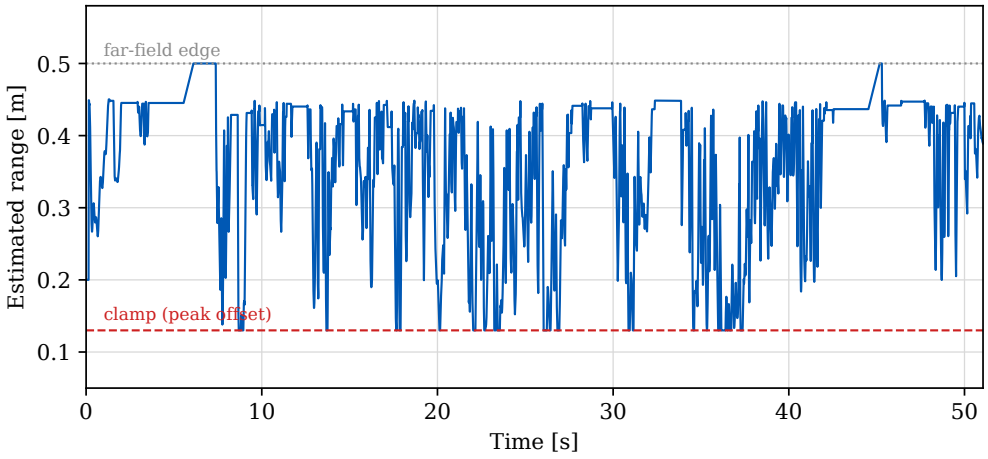

Figure 7.5: Estimated radial range spanning a continuous closed-loop pursuit flight. The dynamic range smoothly fluctuates in response to the leader's relative kinematics, asymptotically increasing toward the far-field boundary during signal dropouts and actively decreasing during proximity encounters. The estimate strictly respects the defined operational bounds, never violating the numerical clamp located at the peak moment offset (dashed line). This confirms that the trailing vehicle consistently maintains a spatial separation where the radial range remains fully observable.

spatial update rate is dictated by the encounter geometry; since a position update requires successful wake detections, the operating localization rate is on the order of several hertz, far below the 200 Hz microprocessor cycle frequency. As a result, the system's dynamic response is sluggish and critically damped rather than exhibiting high-bandwidth transients. This is the expected behavior for a system operating with sparse state updates, and matches the heavily attenuated, first-order controller used in Chapter 5. In the end, the follower is able to successfully maintain a bounded pursuit of the leader using only implicit wake sensing, without requiring additional radio links or exteroceptive sensing channels.

The importance of such results becomes evident if they are compared to the baseline architecture they were compared to. Earlier runs of the same experiment by Cerruti [5] achieved no fewer than one successful pursuit attempt per several dozens of attempts with the same robot platform due to the single-axis detector and its fixed detection threshold. In contrast, the combination of the two-axis vector norm detector, threshold logic and the described above estimation and control loop allow the pursuer to maintain a lock onto the target during both of the long-duration flights, throughout their entire duration. Not only does the proposed solution offer a higher chance of achieving a successful pursuit, but it also achieves a stable

closed-loop control from the pursuer around the entire circular field of view of the target detector, reconstructs a monotonous radial range towards the target and uses these signals to form a stable pursuit cycle without interruption throughout the entirety of the flight of the pursuer drone. This constitutes the main practical result and validation of this thesis.

7.7 Impact of follower-to-leader gestures on leader sensing

In a bidirectional architecture, the trailing UAV is both transmitting on the return gesture channel and receiving messages on the forward wake channel. This raises an important question regarding the experiments conducted in this chapter, namely, if the active transmission on the return gesture channel corrupts the simultaneous aerodynamic sensing taking place on the estimation pipeline. Both analytical and experimental results confirm this concern to be unwarranted, with the companion research showing that cross-channel interference is completely eliminated during nominal flight, and that the single topological configuration affecting active sensing is a deliberate, deterministic suspension.

During steady formation flight, there is no coupling between the two actuation axes. The implicit sensing pipeline we have developed uses only the follower's orthogonal attitude residuals to decode the leader's planar, horizontal motions. The return channel encodes telemetry information in a high frequency vertical oscillation superimposed on the follower's altitude reference, which is sensed by the leader only in its vertical velocity channel. Thus, a follower sending a return gesture does not contaminate its own attitude residual measurements in the horizontal direction. The orthogonality of the approach is ensured since the return gestures are encoded in a narrow high frequency band isolated from the low bandwidth forward channel dynamics. The flight experiments confirm that the sensing pipeline works irrespective of any return channel gestures.The single exception to this orthogonality is physical role-swap, where the entire sensing channel is suspended for the duration of a temporary altitude switch. In order to transmit a discernable gesture, the trailing UAV must rapidly ascend above the leader, since the downwash only affects the follower's local volume in the vertical axis. For the duration of the inverted order, the forward channel is completely erased, since the follower's wake cannot reach the leader (who is now above the follower). This constitutes a true mutual exclusion (as described strictly in Section 6.4), and not a corruption of the signal itself. Active spatial sensing is then deterministically paused and immediately reacquired by the follower as it descends to its nominal lower coordinate.

Consequently, in the experimental sensing evaluation, a role-swap maneuver is a conventional detection void, identical in nature to the stochastic micro-dropouts

natively processed by the estimation pipeline (see Section 7.4 for details). The spatial state estimator executes a zero-order hold throughout the inversion maneuver. The distinguishing characteristic is that this reply-induced gap is commanded and bounded by the gesture encoding sequence, rather than stochastically imposed by environmental wake drift. Given the total duration of the reply maneuver is short compared to the leader's characteristic kinematic correlation time, the latched state estimate is valid upon reacquisition of signals, according to slow varying kinematic assumptions that form the pursuit architecture. Ultimately, the active return gestures are compatible with the passive wake sensing proposed in this thesis: they induce zero cross-channel corruption during steady-state flight, and the sensing voids they require are absorbed by the pipeline's natural intermittency compensation mechanisms.

7.8 Discussion of experimental limitations

The experimental campaign characterizes both the capabilities and the fundamental limitations of obtaining spatial state information from the aerodynamic wake. This section describes limitations imposed on the system by the observed flight and bench telemetry.

The main limitation is imposed by the control authority of the experimental platform. The follower is a 63 g nano-quadrotor, and the downwash of the leader is an immense disturbance compared to the forces it can generate for attitude stabilization. Hence, the follower cannot stay inside the wake for an extended amount of time, since such a disturbance would render it unable to maintain controlled flight. The intermittent nature of the entries in the log files, characterized by transients punctuated by large drops of signal, is a direct result of this inability. The follower samples the wake in short bursts and then retreats to cleaner regions of the workspace to avoid losing control. It is not a failure of the estimator, but a direct result of the impossibility of a nano-quadrotor to withstand the forces in the wake of a larger drone. A more capable platform could have been used to stay in the wake and make spatial sensing possible at all times, but the nano-quadrotor is not capable of withstanding the forces in the wake of a larger drone. The sensing architecture has been designed with this limitation in mind and relies on intermittent access to the information contained in the wake. The secondary limitation is the microprocessor sampling frequency, which limits an observable spectral bandwidth of the downwash. An onboard inertial measurement unit samples at 200 Hz; according to the Nyquist sampling theorem, this constrains the resolved frequency space to a maximum of 100 Hz. However, a dedicated propeller bench test reveals that the tonal signature of the wake, the blade-passing frequency at which the downwash pulses, lies between roughly 140 and 510 Hz, depending on the rotor speed and blade count. This

entire range far exceeds the basic 100 Hz Nyquist limit of the follower's sensory hardware. The characteristic signature of a blade-passing frequency cannot be sampled by the platform and would experience severe aliasing if a spectral analysis were attempted onboard. The trailing vehicle is mathematically isolated from accessing the high-frequency spectral content it could potentially use to uniquely identify a specific rotor signature. Further, within the resolved bandwidth below 100 Hz, the aerodynamic disturbance does not manifest any coherent tonal content. Spectral analysis of the recorded flight data shows that while the downwash from the leader increases the broadband energy by an order of magnitude at 10 cm compared to the isolated follower, it does not introduce any significant spectral lines. The prominent narrow-band peaks in the spectra of the induced moment at the follower do not exceed the local noise level by more than a factor of two to three, and they are confined to the 2 Hz to 20 Hz range, being associated with the rigid-body dynamics and control loops of the vehicles rather than any inherent aerodynamic periodicity. The induced moment at the follower does not contain any significant tones at the frequencies corresponding to the propeller rotation rate. The direct implication of this observation is that the aerodynamic channel does not pass on the *frequency* information about the leader, although it allows recovering the *amplitude* of the disturbance, and hence the range, from which the spatial coordinates can be inferred. Therefore, while the follower can localize the leader, it cannot identify it from the aerodynamic cues and track any frequency-modulated information sent by it. In particular, demodulating the fundamental tone from the blade passing would require an inertial measurement chain with a sampling rate above 1 kHz, and it is not evident whether such a signal would be present above the broadband noise at these frequencies. The tertiary limitation is the complete loss of radial range observability in the central axial domain. Namely, the empirical flight telemetry demonstrates that the reconstructed range estimate is never invalid below the numerical clamp parameterized near 0.13 m, corresponding to the spatial offset of maximum induced moment. Specifically, the mathematical inversion algorithm inherently rejects the inner spatial region in which a singular signal magnitude ambiguously maps to multiple discrete distances. Physically, directly beneath the leading vehicle, the aerodynamic loading is evenly distributed on the follower's airframe, thus inducing a negligible net moment. In this region, the follower is able to accurately track the relative bearing from the leader, but loses all information regarding the deterministic radial distance. Namely, the sensing architecture is blind to range estimates within a specific disc surrounding the central wake axis; the closed-loop pursuit controller is designed to maintain a radial standoff beyond this unobservable disc, for this exact reason. With these limitations in mind, the empirical experiments demonstrate that the trailing vehicle successfully detects the leader, dynamically resolving its relative bearing across the entire circular azimuth, and reconstructing a monotonic range estimate beyond the near-axis blind zone by

using only the implicit aerodynamic wake, as observed by its endogenous inertial sensors.

Chapter 8

Scalability of Wake-Based Sensing in Multi-UAV Formations

Thus far, this dissertation has exclusively investigated the dynamics and control architecture of a singular leader-follower pair. The present chapter broadens this scope to explore the theoretical implications of extending this paradigm into a decentralized multi-UAV formation, wherein each constituent vehicle derives its relative spatial awareness implicitly from the aerodynamic wake of its immediate predecessor. Given that the empirical flight campaigns detailed in Chapter 7 were strictly confined to bipartite configurations, this extension is approached from a purely analytical and topological perspective, deferring macroscopic experimental validation to future research endeavors. The primary objective is to formalize the mathematical framework governing decentralized wake-coupled swarms, to isolate the systemic phenomena that intrinsically bound the formation's scalability, and to analytically derive both the governing directed graph and the theoretical convergence time for the minimal multi-agent topology. This chapter sequentially constructs this theoretical extension. Section 8.1 formulates the topological transition from a discrete pair to a continuous cascaded chain. Section 8.2 models the spatial accumulation of state-estimation variance along the sequence. Section 8.3 establishes the rigorous criteria for string stability. Subsequently, Section 8.4 constructs the Laplacian representation of the network's sensing dependencies, thus incorporating a virtual reference node, to extract the eigenspectrum and exact convergence temporal bounds for the foundational case. Finally, Section 8.5 delineates the rigid geometric constraints dictated by aerodynamic physics, culminating in Section 8.6, which synthesizes these factors to establish the absolute theoretical limits on scalability.

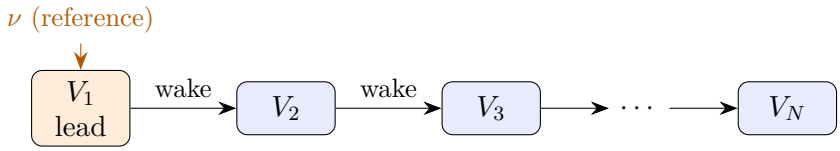

Figure 8.1: Topological representation of the formation as a sequential chain. Each UAV derives its pursuit trajectory by sensing the downwash of its immediate predecessor, thereby rendering each link a discrete, unidirectional wake-sensing edge. The terminal lead vehicle, V_1, strictly regulates its state toward a mathematically invariant virtual node, ν, which deterministically anchors the entire kinematic chain to a prescribed spatial coordinate.

8.1 Extension from one leader-follower pair to UAV chains

The most intuitive topological extrapolation of the leader-follower paradigm is a sequential, cascaded chain. Within this framework, every intermediate UAV simultaneously acts as a follower to the agent directly ahead and as a wake-generator for the agent immediately behind. Consequently, the formation manifests as a linear array wherein the implicit wake-sensing mechanism is iterated link by link. Specifically, the primary agent (V_1) acts as the global leader; the secondary agent (V_2) autonomously pursues V_1 by localizing within its turbulent wake; the tertiary agent (V_3) replicates this process utilizing the wake of V_2, and this recursive architecture propagates indefinitely. Each localized link operates identically to the bipartite wake-sensing edge empirically validated in preceding chapters, rendering the macroscopic chain a direct mathematical composition of these discrete subsystems. Figure 8.1 provides a topological abstraction of this structure.

This cascaded topology intrinsically inherits all the spatial limitations of a singular aerodynamic link, imposing geometric boundary conditions that are markedly more severe than those governing conventional RF-linked swarms. Sustained spatial awareness mandates that every vehicle remains strictly confined within the aerodynamic envelope of its predecessor. Therefore, the permissible inter-vehicle spacing is rigidly constrained by the maximum functional range of the downwash, empirically bounded at approximately 15 cm of radial offset for the current nano-quadrotor platform (as characterized in Chapter 3). Should any UAV drift beyond this physical threshold, it suffers complete sensory deprivation and is categorically severed from the network. Consequently, arbitrary spatial distributions are physically unrealizable, and the ultimate scalability of the chain is fundamentally dictated by the system's capacity to contain cumulative tracking errors within this narrow geometric margin.

The global reference for the entire formation is not an active physical agent, but rather a mathematically static spatial coordinate, denoted as the virtual node (ν). Originally conceptualized in parallel research concerning return-channel dynamics, this node acts as the absolute rendezvous datum for the formation. It possesses an invariant state, exhibits zero internal dynamics, receives no feedback, and does not represent a physical airframe. It functions as the ultimate hierarchical head of the chain; the primary physical leader V_1 regulates its tracking error exclusively relative to ν, thereby tethering the macroscopic swarm to a deterministically prescribed location. The necessity of this virtual node stems from a fundamental property of unanchored consensus networks: in a purely decentralized topology lacking a rigid global reference, constituent agents asymptotically converge to a consensus state that is highly sensitive to arbitrary initial conditions, thereby meaning the formation will stabilize, but at a non-deterministic spatial coordinate. The virtual node mathematically eliminates this unbounded degree of freedom, since ν unconditionally projects its state without being subjected to reciprocal influence, it strictly forces the steady-state configuration of the entire network to converge to its exact defined value, an operational guarantee formally proven via Laplacian analysis in Section 8.4.

8.2 Accumulation of estimation error along the formation

Each discrete link within the sequence intrinsically generates a localized state estimation error, an unavoidable consequence of each follower attempting to reconstruct its target's relative position through the highly stochastic and intermittent aerodynamic channel. The critical analytical question governing the feasibility of long chains is how these localized, per-link errors mathematically compound as they propagate down the formation. Assuming the spatial tracking error introduced by a singular arbitrary link is characterized by a standard deviation σ_1,then the absolute positional error of the n-th vehicle relative to the global hierarchical head represents the mathematical accumulation of the errors across all n intervening links. Under the assumption that these per-link estimation errors are statistically independent, their respective variances undergo algebraic summation. Consequently, the standard deviation of the terminal accumulated error scales proportionally with the square root of the sequence length:

$$\sigma_n = \sqrt{\sum_{i=1}^{n} \sigma_1^2} = \sqrt{n}\,\sigma_1. \tag{8.1}$$

This formulation corresponds exactly to the dispersion profile of an uncorrelated random walk. Conversely, if the per-link errors exhibit strong statistical correlation

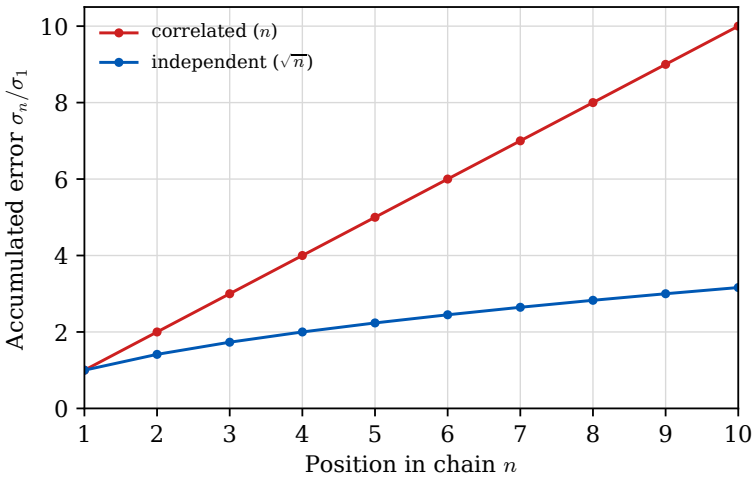

Figure 8.2: Theoretical bounding of accumulated state estimation error along the formation chain, normalized against the unitary per-link standard deviation σ_1. Assuming strict statistical independence, the terminal error profile expands proportionally to $\sqrt{n}$. Conversely, total statistical correlation induces linear expansion scaling with n. Given the stochastic heterogeneity of the wake channel, the localized errors fundamentally tend toward independence; thus, the lower $\sqrt{n}$ trajectory characterizes the expected scaling behavior, while the upper linear trajectory defines the absolute conservative limit.

for instance, due to a ubiquitous exogenous wind disturbance biasing every link uniformly, then the errors sum coherently rather than in quadrature. In this scenario, the accumulation profile degenerates into strict linear growth:

$$\sigma_n = \sum_{i=1}^{n} \sigma_1 = n\,\sigma_1. \tag{8.2}$$

This linear domain defines the absolute worst-case theoretical bound. Figure 8.2 graphically contrasts these two divergent trajectories: while the empirical error accumulation inherently resides between these mathematical limits, the highly localized and heterogeneous physical nature of the aerodynamic channel, where each follower intersects a geometrically distinct cross-section of its predecessor's turbulent wake, strongly biases the stochastic per-link errors toward statistical independence. Therefore, the square-root progression model (8.1) represents the most probable systemic behavior, while the linear model (8.2) serves as an ultra-conservative upper boundary.
Irrespective of the precise correlation coefficient, the absolute spatial variance monotonically increases as a direct function of the chain's length, establishing

the primary physical limitation on macroscopic scalability: an agent positioned deeply within the sequence calculates its global coordinates relying on a cascaded measurement heavily degraded by the aggregated variance of all preceding links. Because aerodynamic stability necessitates that every UAV remain strictly within the finite geometric boundary of the downwash, systemic failure becomes inevitable when the localized accumulated error σ_n approaches the physical magnitude of the wake radius. At this threshold, the n-th vehicle can no longer probabilistically guarantee spatial adherence to the signal source, precipitating immediate disconnection. Given that the empirical per-link error quantified in Chapter 7 is structurally comparable to the narrow 15 cm effective wake radius, this critical failure point is mathematically breached after merely a handful of cascaded links. This severe geometric tolerance constitutes the fundamental physical reason why purely passive, aerodynamically coupled chains cannot scale indefinitely.

8.3 String stability from the sensing and pursuit perspective

Whereas error accumulation governs the steady-state geometric dispersion of the chain, string stability dictates the transient dynamic propagation of kinematic disturbances along the sequence. A vehicular formation is rigorously defined as string stable if, and only if, a localized high-frequency perturbation applied to a specific agent is dynamically attenuated as it propagates through subsequent trailing agents [43]. If the cascaded closed-loop control architecture induces even marginal amplification at each successive link, this magnitude compounds geometrically, forcing a negligible transient at the leader to evolve into a catastrophic, unbound oscillation at the tail.

The mathematical prerequisite for string stability depends entirely upon the localized per-link transfer function: let γ denote the steady-state magnitude of the transfer function mapping the spatial tracking error of an arbitrary vehicle to the corresponding tracking error of its immediate follower. Given that the sequential links represent fundamentally identical, cascaded dynamic plants, the terminal kinematic error of the n-th vehicle resulting from a perturbation at the global head scales proportionally as the product of these localized gains:

$$\frac{e_n}{e_0} = \gamma^n. \tag{8.3}$$

Consequently, the kinematic disturbance is mathematically guaranteed to asymptotically decay along the formation sequence if and strictly if the per-link gain satisfies:

$$\gamma \leq 1. \tag{8.4}$$

When the condition $\gamma < 1$ is strictly maintained, any transient anomaly is sequentially suppressed, ensuring the chain remains unconditionally string stable regardless of its absolute length. Conversely, if $\gamma > 1$, the perturbation undergoes geometric magnification, rendering the network violently unstable beyond a critical sequence threshold. The precise magnitude of γ is deterministically established by the specific tuning of the pursuit controller. A high-bandwidth controller designed for aggressive, highly responsive tracking inevitably induces closed-loop phase lag and kinematic overshoot, thereby amplifying the predecessor's spatial transient into the follower's trajectory and driving γ definitively above unity. In contrast, a critically damped, low-bandwidth controller configured to track the target monotonically without inducing overshoot systematically acts as a low-pass filter, maintaining γ safely below unity. The first-order controller synthesized in Chapter 5 was intentionally parameterized to guarantee monotonic convergence without overshoot; therefore, the identical tuning paradigm required to stabilize the isolated bipartite pair simultaneously enforces the condition $\gamma \leq 1$. Although this theoretical framework strongly implies string stability, explicitly verifying this dynamic attenuation across a physical multi-agent chain remains a critical objective for future empirical validation.

8.4 Graph representation of sensing dependencies

The directed spatial dependencies governing the multi-agent formation are mathematically formalized via a directed graph, where the spectral properties of the associated graph Laplacian rigorously dictate the asymptotic convergence dynamics of the formation state errors. This section models the directed graph for the minimal non-trivial configuration, hence the bipartite pair coupled with the virtual reference node, calculates its corresponding eigenspectrum, and analytically extracts the explicit temporal convergence constant. This formulation adheres strictly to the continuous-time consensus framework established by Olfati-Saber et al. [44] and Olfati-Saber and Murray [45], which proves that the convergence velocity of a linear consensus protocol is exclusively parameterized by the eigenvalues of the graph Laplacian, specifically bounded by the network's algebraic connectivity.
This foundational topological configuration comprises three distinct elements: the virtual reference node ν, the physical leader UAV L, and the trailing physical follower UAV F. Because aerodynamic downwash propagates exclusively along the downward vertical axis, the primary sensory channel is strictly unidirectional. Consequently, the graph topology transitions between two distinct, mutually exclusive operational modes, mathematically detailed in Figure 8.3. During the dominant *sensing mode*, the follower is spatially positioned beneath the leader; the

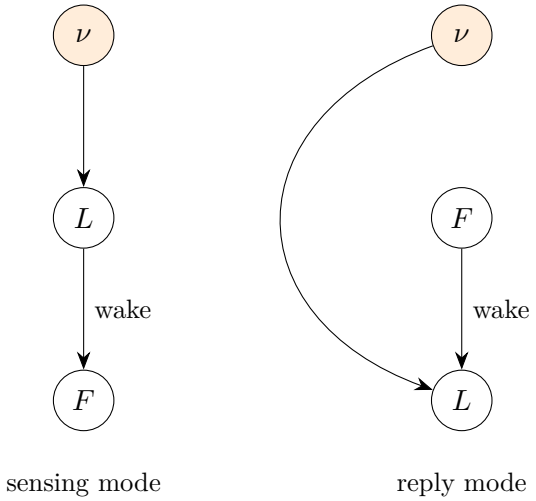

Figure 8.3: The augmented graph in its two modes. In sensing mode the wake carries from the leader to the follower below it, and the leader regulates toward the reference ν. In reply mode the leader descends below the follower and the sensing edge reverses, while the reference still drives the leader. No edge enters ν in either mode, so the reference is never influenced by the vehicles, and the path from ν to the follower is restored after each reply.

aerodynamic wake establishes a unidirectional sensing edge from L to F, while the leader actively anchors its state toward the global reference, generating an edge from ν to L. During the transient *reply mode*, the leader initiates a controlled descent beneath the follower to reestablish the bidirectional communication channel, forcing the physical sensing edge to strictly invert its orientation. Crucially, across both discrete modes, the in-degree of the virtual node ν remains absolute zero: the global reference is mathematically immutable and completely isolated from the dynamic state feedback of the physical vehicles. A continuous directed path $\nu \to L \to F$ is fully established exclusively during the sensing mode; because the system deterministically reverts to this stable topology following every reply sequence, the terminal follower remains asymptotically reachable from the absolute reference throughout the duration of the operational cycle.
Isolating the continuous sensing mode and assigning the proportional pursuit gain k as the uniform scalar weight for each active directed edge, the continuous-time graph Laplacian is formalized as follows:

$$\mathcal{L} = \begin{bmatrix} 0 & 0 & 0 \\ -k & k & 0 \\ 0 & -k & k \end{bmatrix}, \tag{8.5}$$

where the matrix rows and columns are ordered according to the state vector

sequence ν, L, F. The initial row is entirely null, mathematically verifying that no incoming edges influence the absolute reference node. The global error dynamics of the formation evolve according to the continuous linear differential equation $\dot{\mathbf{x}} = -\mathcal{L}\,\mathbf{x}$; therefore, the spectral decomposition of $\mathcal{L}$ fundamentally dictates the asymptotic convergence profile of the formation error.
The characteristic eigenvalues of the Laplacian $\mathcal{L}$ are calculated as:

$$\lambda_1 = 0, \qquad \lambda_2 = \lambda_3 = k. \tag{8.6}$$

The corresponding eigenvectors provide a comprehensive mathematical abstraction of the formation's terminal behavior. The singular zero eigenvalue $\lambda_1 = 0$ corresponds to the right eigenvector $\mathbf{v}_1 = [1,1,1]^\top$, defining the consensus subspace wherein all nodes achieve perfectly identical state values. This represents the absolute convergence configuration of the formation, occurring when both physical vehicles successfully regulate their tracking errors to zero relative to the reference. Because the state of ν is statically invariant, this shared terminal coordinate is rigidly constrained to the prescribed state of ν, mathematically proving that the virtual node deterministically anchors the final configuration of the network. The strictly positive, repeated eigenvalue $\lambda_2 = \lambda_3 = k$ characterizes the transient decaying modes: any arbitrary initial state deviation of the physical vehicles from the absolute reference resides exclusively within the subspace spanned by these eigenvectors, and must exponentially decay at a uniform rate exactly equal to k. Because the zero eigenvalue is simple, the reference node remains globally reachable across the active network graph during the primary sensing mode, thereby satisfying the fundamental graph-theoretic requirement for consensus convergence.
The non-zero eigenvalue k mathematically defines the algebraic connectivity of this specific directed graph topology, exclusively establishing the system's absolute convergence velocity. The aggregate formation error decays proportionally to e^{-kt}, yielding a systemic convergence time constant of $\tau = 1/k$. Implementing the empirically validated proportional pursuit gain $k = K_P = 0.08\,\mathrm{s}^{-1}$ from Chapter 5, the resultant time constant evaluates to $\tau = 12.5\,\mathrm{s}$. Notably, this value is mathematically identical to the isolated time constant of a singular control loop, as the two sensing links operate strictly as cascaded first-order systems sharing a uniform gain. The absolute convergence timeout, $t^\star$, defined as the specific temporal duration required for the macroscopic formation error to transition from an arbitrary initial magnitude e_0 into a strictly bounded terminal convergence envelope defined by ε, is analytically derived from this exponential decay equation:

$$t^\star = \frac{1}{\lambda_2}\,\ln\frac{e_0}{\varepsilon} = \frac{1}{k}\,\ln\frac{e_0}{\varepsilon}. \tag{8.7}$$

This derivation establishes the requisite operational timeout that a high-level mission planner must allocate to ensure the formation fully settles before executing

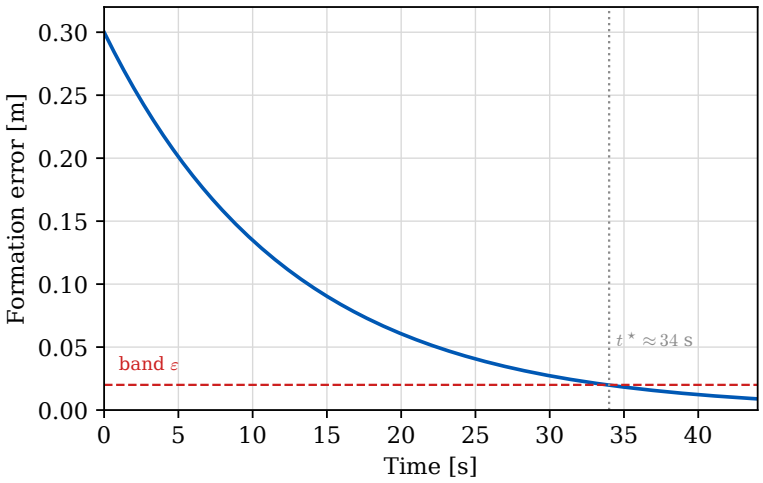

Figure 8.4: Transient exponential decay profile of the aggregate formation error converging toward the terminal boundary band, mathematically modeled for the minimal pair coupled with the virtual reference node. The error magnitude diminishes exponentially at a rate governed exclusively by the non-zero Laplacian eigenvalue k. The timeout parameter $t^\star$ defines the precise temporal coordinate at which the decaying error intersects the specified tracking tolerance ε. Employing the established parameters $k = 0.08\,\mathrm{s}^{-1}$, initial error $e_0 = 0.30\,\mathrm{m}$, and an aggressive tolerance bound $\varepsilon = 0.02\,\mathrm{m}$, the resultant analytical timeout is bounded at approximately 34 s.

subsequent maneuvers, a trajectory graphically depicted in Figure 8.4. Parameterizing this model for the bipartite pair with the established gain $k = 0.08\,\mathrm{s}^{-1}$, assuming an initial aggregate spatial error of $e_0 = 0.30\,\mathrm{m}$, and targeting a stringent convergence tolerance of $\varepsilon = 0.02\,\mathrm{m}$, the resultant convergence timeout is calculated as:

$$t^\star = \frac{1}{0.08} \ln \frac{0.30}{0.02} \approx 34\,\mathrm{s}. \tag{8.8}$$

Relaxing the terminal convergence boundary to a slightly wider tolerance of $\varepsilon = 0.05\,\mathrm{m}$ correspondingly reduces the requisite timeout to approximately 22 s.
This constitutes the fully explicit analytical derivation of the convergence dynamics for the minimal two-vehicle configuration tethered to a virtual node, extracted directly from the spectral properties of the associated graph Laplacian. Mathematically generalizing this Laplacian framework to accommodate an extended sequence of N cascaded vehicles is structurally straightforward, necessitating only the incorporation of one supplementary state node and one corresponding directed edge per additional agent. However, extracting the generalized eigenvalue spectrum

for these macroscopic sequences mandates independent spectral computation for each arbitrary formation length. Because this generalized spectrum is not strictly required to validate the foundational bipartite architecture introduced within this thesis, the comprehensive eigenvalue analysis for highly extended topologies is relegated to future analytical development.

8.5 Constraints on formation geometry and inter-vehicle spacing

The overarching spatial geometry of any wake-coupled formation is fundamentally constrained by the rigid aerodynamic topology of the downwash, imposing operational restrictions that are significantly more severe than those governing generic, radio-linked swarms. Because every trailing vehicle must maintain continuous spatial localization strictly within the specific turbulent volume generated by its immediate predecessor and given that the downwash manifests physically as a highly concentrated, unidirectionally downward jet, the permissible inter-vehicle spacing is simultaneously bounded across multiple spatial dimensions.

The absolute vertical separation is structurally dictated by the axial dissipation of the wake's kinetic energy: the follower must operate at a sufficiently lower altitude than the leader to ensure the turbulent signal is not attenuated below the detector's ambient noise floor, a limit empirically bounded within a few tens of centimeters for this specific platform (as established in Chapter 3). Simultaneously, the horizontal planar separation is dictated entirely by the annular topology of the reliable sensing region, since the follower must maintain a discrete radial offset localized explicitly within the spatial band where the induced moment is maximized—avoiding both the central geometric axis, where aerodynamic symmetry renders the signal negligible, and the far-field boundary, where the flow field structurally collapses. Consequently, the mathematically optimal formation geometry mandates a target radial offset situated firmly within the peak-signal annulus (empirically spanning approximately 15 to 30 cm) coupled with an active bearing controller to maintain continuous alignment within this band. The resultant stable formation must therefore adopt the geometric form of a staggered, descending staircase, where each successive vehicle is operating below and radially offset from its immediate predecessor, rather than a generic horizontal plane. A strictly horizontal planar configuration would erroneously force each trailing vehicle precisely onto the central aerodynamic axis of its leader, systematically eliminating the requisite induced moment and causing catastrophic, network-wide signal loss.

Critically, this staggered spatial geometry compounds linearly with the absolute length of the chain, because every individual UAV must physically operate below and radially offset from the agent preceding it, a cascaded formation comprising

N vehicles systematically necessitates N discrete downward altitude steps and N accumulated horizontal offsets. The maximum vertical descent is inherently bounded by the UAV's operational altitude ceiling or its proximity to the ground, while the systematic horizontal stagger forces the tail of the formation to sweep across a massive, unconstrained spatial footprint. Consequently, a high-population chain cannot be geometrically condensed into a highly localized volume. This fundamental geometric compounding independently bounds the practical scalability of the formation, completely decoupled from the statistical error accumulation or dynamic string-stability constraints analyzed previously. Furthermore, these limits are heavily interconnected: an extended chain demands significantly more operational airspace, which inherently necessitates increasing the nominal inter-vehicle spacing to prevent physical collisions. This expanded spacing drastically attenuates the aerodynamic signal-to-noise ratio at every structural link, fundamentally exacerbating the localized per-link tracking error and accelerating the overall statistical failure of the chain.

8.6 Scalability limits of wake-based leader tracking

The comprehensive theoretical analysis detailed in the preceding sections formally isolates three distinct, fundamental physical constraints that systematically bound the maximum sustainable length of a wake-coupled sequential chain, thereby defining the absolute scalability limits of this purely passive sensing modality.

The primary constraint is the progressive accumulation of statistical state estimation error: the global localization variance expands monotonically along the formation sequence, thus scaling proportionally with $\sqrt{n}$ assuming statistical independence, and severely deteriorating to linear n scaling under highly correlated disturbance conditions. The absolute maximum chain length is definitively bounded at the structural index where this compounding tracking variance exceeds the finite physical boundary of the maximum wake reach. Beyond this critical threshold, the trailing vehicle can no longer probabilistically guarantee sustained spatial adherence within the turbulent downwash of its predecessor, leading to inevitable signal loss and structural fracture. Given the high-variance per-link error characteristics and the narrow 15 cm effective wake radius empirically established for this specific platform, this statistical failure threshold is mathematically breached after merely a small sequence of cascaded links.

The secondary constraint is the rigorous requirement for dynamic string stability: the macroscopic chain guarantees stability independent of its physical length strictly if every individual link systematically attenuates transient kinematic disturbances, mandating an absolute per-link transfer gain of $\gamma \leq 1$. While the critically

damped, low-bandwidth proportional pursuit controller is theoretically synthesized to explicitly satisfy this mathematical bound, formal empirical validation across a physical multi-agent chain remains a critical necessity. Should any specific controller tuning inadvertently induce dynamic overshoot, the associated per-link gain would instantly exceed unity, precipitating uncontrollable geometric amplification of high-frequency disturbances and inducing catastrophic instability beyond a defined formation length, entirely irrespective of the static statistical error budget.

The tertiary constraint is the rigid geometric compounding of the formation structure. Because the underlying physics of the unidirectional downwash fundamentally mandate that each trailing vehicle must operate strictly below and spatially offset from its predecessor, an extended sequence manifests inherently as a continuously descending and horizontally staggering structure. Consequently, the sheer spatial volume mandated to safely house the formation drastically expands with the addition of each consecutive agent. This geometric requirement inherently bounds the macroscopic length of the chain based exclusively on the finite availability of operational airspace, completely independent of the previously defined algorithmic or statistical limitations. Crucially, these three identified bounds are highly interdependent: the mandated spatial geometry dictates the permissible inter-vehicle spacing; this spacing directly modulates the aerodynamic signal strength, defining the baseline per-link statistical error; and the compounding integration of this statistical error, coupled with the dynamic string stability gain, ultimately determines the absolute threshold where the extended chain fractures.

These systemic limits represent fundamental physical properties of the aerodynamic wake channel itself, completely decoupled from specific computational estimation architectures or arbitrary control tuning paradigms. They represent the inescapable engineering compromise mandated by a strictly passive sensing modality that elegantly eliminates the necessity for heavy RF communication links or vulnerable optical tracking markers. A conventional radio-linked formation remains structurally unconstrained by narrow wake boundaries, rigid staggered descent geometries, or heavily cascading per-link sensing variances, and can subsequently be deployed into arbitrary, highly dense macroscopic configurations. The implicit wake channel systematically trades this immense spatial flexibility for rigorous systemic passivity: it ingeniously leverages the endogenous proprioceptive sensors every UAV naturally possesses, but strictly confines its operational validity to the narrow, localized geometric regions where aerodynamic disturbances robustly couple between adjacent vehicles. While this chapter has comprehensively established the theoretical scalability limits and computed the explicit dynamic equations for the minimal foundational configuration, the logical continuation of this research mandates a dedicated empirical multi-vehicle flight campaign coupled with a generalized eigenvalue analysis of macroscopic, arbitrary-length chain topologies.

Chapter 9

Conclusions

This thesis developed a fully passive sensing channel enabling a quadrotor to track a leading vehicle exclusively through aerodynamic wake interactions, eliminating the need for radio communication links or artificial markers by relying solely on the endogenous inertial measurement unit. The follower perceives the leader via the aerodynamic moment induced by the turbulent wake upon its attitude dynamics; from this disturbance, the estimation pipeline reconstructs the leader's spatial coordinates and closes a dynamic pursuit loop Thus, this concluding chapter summarizes the primary theoretical and practical contributions, aggregates the definitive experimental findings, explicitly delineates the operational limitations revealed by the empirical data, and outlines the subsequent research trajectories established by this work.

9.1 Summary of the individual contribution

The foundational contribution of this dissertation is the synthesis and validation of the forward channel within an implicit aerodynamic communication architecture: the leader-to-follower physical link, wherein the trailing UAV perceives the leader by isolating the disturbance its aerodynamic wake induces on the follower's attitude state. This research operates in parallel with the complementary return-channel, which formalizes the follower-to-leader gesture link; together, these components constitute a unified bidirectional system. The channel formulation is fundamentally predicated on aerodynamic physics: Chapter 3 theoretically modeled the leader's turbulent wake and empirically characterized it via an anemometric campaign on the physical platform; therefore, this analysis conclusively demonstrated that the induced aerodynamic moment originates from the asymmetric pressure loading of the downwash jet across the follower's airframe, rather than scaling directly with the absolute flow velocity. This critical physical distinction dictates the entire

subsequent sensing topology: it mathematically explains why the sensory signal is negligible directly beneath the leader, where aerodynamic loading is perfectly symmetric, and reaches a spatial maximum at an intermediate radial offset, where the aerodynamic imbalance is maximized. Leveraging this aerodynamic model, Chapter 4 synthesized the spatial estimation pipeline: a robust detector was formulated utilizing the Euclidean norm of the two orthogonal attitude residuals, rendering the detection logic strictly invariant to the leader's approach azimuth and successfully eliminating the diagonal blind spots inherent to the single-axis baseline. The detection threshold was engineered to dynamically track the stochastic ambient noise floor, supplanting a static threshold susceptible to calibration drift caused by battery voltage degradation. Ultimately, an Ordered-Statistic Constant False Alarm Rate (OS-CFAR) detector was implemented, resolving a severe buffer freezing artifact exposed during empirical flight tests of an earlier variance gated design. Following temporal signal debouncing, deterministic estimators successfully reconstructed the relative bearing, which wash extracted from the directional phase of the residual vector and the radial range, derived via numerical inversion of the monotonic outer branch of the aerodynamic moment profile. Chapter 5 subsequently closed the dynamic control loop by implementing a polar-coordinate pursuit controller, predicated on the classical pure-pursuit guidance law [41], which actively steers the trailing UAV along the instantaneous line-of-sight vector while regulating its scalar velocity to maintain a predefined radial standoff distance, explicitly avoiding target interception. The controller architecture is intentionally parameterized with heavily attenuated gains and first-order dynamics, deliberately optimized to guarantee stability when operating on highly intermittent, high-variance measurement data, constraining the operational equilibrium safely beyond the region of range ambiguity. Chapter 7 executed the comprehensive empirical validation of this integrated system on physical flight hardware. Finally, Chapter 8 analytically extrapolated this architecture to macroscopic formations, explicitly deriving the governing directed graph and the theoretical convergence time for the minimal configuration coupled to a virtual reference node [44].

9.2 Main experimental findings

The empirical flight campaign yielded four definitive findings, each rigorously corroborated by high-frequency onboard telemetry: first, the aerodynamic wake signal maintains a robust statistical separation from the ambient noise floor, thus being validated across a baseline dataset of eleven leader-free hover flights, where the residual norm exhibited a median magnitude of approximately 14 mrad and a ninety-ninth percentile bounded near 58 mrad. During active wake interactions, the signal amplitude escalated significantly, registering ninetieth percentiles between

60 mrad to 90 mrad, with transient peaks surging several times higher, yielding a Signal-to-Noise Ratio (SNR) of approximately 13 dB to 16 dB. Active detection remains highly reliable whenever the follower intersects the wake, and the OS-CFAR thresholding architecture consistently bounds the false-alarm rate to approximately one percent during baseline flights, representing a definitive functional improvement over the legacy static design. Second, the relative bearing is reliably resolved across the complete azimuthal domain, since throughout the dynamic pursuit flights, the estimated direction to the leader successfully populated all twelve discrete 30° sectors surrounding the follower, conclusively demonstrating the elimination of the diagonal blind spots inherent to the single-axis baseline. Third, the radial range is successfully reconstructed as a highly consistent and monotonic indicator of spatial separation: analysis of the flight logs revealed a rigorous negative correlation (approximately -0.95) between the sensory signal amplitude and the estimated range, thus demonstrating strong repeatability across independent trials, validating the algorithm's internal consistency (acknowledging the deliberate absence of an exogenous absolute metrological reference). Fourth, the integrated architecture successfully sustains stable, closed-loop pursuit: over the extended dynamic evaluations lasting approximately 50 s, the follower autonomously maintained a bounded pursuit cycle, continuously constraining the estimated leader position to within an approximate 0.3 m radius envelope while strictly preserving range observability throughout the maneuvers. Collectively, these empirical results represent a profound capability leap beyond the baseline architecture established by Cerruti [5], which yielded an exceptionally low operational success rate, approximately one successful pursuit out of several dozen attempts, due to the structural limitations of its single-axis detector and static threshold. The synergistic integration of the two-axis Euclidean norm, OS-CFAR thresholding, and continuous polar state estimation successfully elevates transient, binary wake detection into a robust, sustained spatial track and autonomous pursuit. The fundamental significance of this research lies in demonstrating that a 63 g nano-quadrotor, relying entirely on endogenous inertial sensors, can autonomously localize and track an identical vehicle exclusively via aerodynamic downwash. The channel that prior literature treated merely as an amplitude scalar for binary proximity detection is formally proven herein to encode complete directional and distance kinematics sufficient to close a pursuit loop. This establishes the turbulent wake,historically treated merely as an exogenous disturbance to be rejected,as a rich, multidimensional relative-state sensing modality in its own right.

9.3 Limitations of the proposed sensing and estimation approach

The empirical validation simultaneously illuminated the fundamental systemic limitations of the proposed architecture, establishing strict boundary conditions on its operational applicability, since a 63 g trailing vehicle possesses insufficient actuation authority to indefinitely stabilize against the massive aerodynamic disturbance generated by the leader's downwash. Consequently, the follower cannot physically sustain prolonged continuous flight within the wake; it must execute transient sensory intersections and rapidly retreat to recover stability. This physical constraint mandates that the sensory channel is inherently intermittent, active only for a fraction of the flight envelope, rendering the pursuit a continuous approach-and-hold limit cycle rather than an uninterrupted track. This intermittency is a direct physical consequence of the specific airframe scale and actuator saturation limits, not an algorithmic deficiency; a macroscopic platform with superior control margins would theoretically permit continuous wake dwell and uninterrupted sensing.

The secondary limitation is imposed by the microprocessor sampling frequency, which strictly bounds the observable spectral bandwidth of the aerodynamic flow field, since the endogenous inertial measurement unit operates at 200 Hz, thus imposing a rigid Nyquist limit of 100 Hz on resolvable frequency content. However, empirical bench tests confirm that the fundamental blade-passing tonal signature of the wake resides significantly higher, spanning approximately 140 to 510 Hz. Consequently, the high-frequency tonal signature is mathematically inaccessible to the platform and would suffer severe aliasing. Within the resolvable bandwidth below 100 Hz, the induced moment manifests exclusively as broadband turbulence devoid of coherent spectral lines. Thus, the current sensing channel strictly propagates the scalar amplitude of the disturbance (facilitating spatial localization) but provides zero usable frequency modulation. The trailing vehicle is therefore physically precluded from utilizing spectral analysis to categorically identify specific target vehicles or decode frequency-modulated communication streams.

The tertiary limitation is the fundamental loss of radial range observability within the central axial domain, since directly beneath the leading vehicle, the aerodynamic loading distributes symmetrically across the follower's airframe, resulting in a negligible net induced moment. While the relative bearing remains resolvable, the follower completely loses deterministic radial distance information within this localized zone. Therefore, the sensing architecture is mathematically blind to range estimation within a concentrated spatial disc surrounding the central wake axis, dictating that the pursuit controller must be explicitly parameterized to maintain a radial standoff strictly beyond this unobservable region. These constraints do not

represent failures of the estimation algorithms or control logic, but are inherent properties of the passive sensing modality and the highly constrained physical platform; each limitation directly motivates a trajectory for future research.

9.4 Future work on robust wake-based formation control

The theoretical and empirical conclusions of this dissertation establish three distinct, high-impact trajectories for subsequent research.

The first trajectory focuses on the empirical realization of macroscopic, multi-agent formations: Chapter 8 formalized the theoretical scalability of a cascaded wake-sensing chain, modeling the compounding estimation variance, defining the strict prerequisites for dynamic string stability [43], and analytically extracting the Laplacian eigenspectrum and convergence timeout for the minimal virtual-reference configuration [44, 45]. However, this rigorous analysis currently terminates at the bipartite pair validated empirically. The immediate continuation of this research mandates deploying an extended chain of multiple physical UAVs to empirically validate the modeled error accumulation, quantify the per-link string stability transfer gain, and benchmark the absolute convergence times against the theoretical Laplacian bounds. Furthermore, the rigid geometric requirement for a staggered, descending staircase formation, dictated by the wake's topology, must be practically evaluated against finite operational airspace constraints. The absolute limit where cumulative per-link statistical error intercepts the finite wake boundary will definitively establish the maximum physical length of the passive chain.

The second trajectory targets the exploitation of the wake's high-frequency spectral signature. Empirical bench testing conclusively demonstrated that individual propellers generate distinct, coherent blade-passing frequencies, ranging from 140 to 510 Hz, and that heterogeneous signatures can be cleanly separated via spectral analysis, successfully isolating dual distinct propellers from a single composite recording. Because this tonal data is fundamentally inaccessible to the baseline 200 Hz inertial unit, the current channel operates exclusively in the amplitude domain. Integrating specialized, high-bandwidth inertial sensors (e.g., kilohertz-range gyroscopes prevalent in First-Person-View platforms) would unlock the frequency dimension of the aerodynamic channel. This capability would permit the follower not merely to localize the leader, but to categorically identify it by its unique spectral footprint. Furthermore, it would enable discrete multi-bit data transmission via active modulation of the leader's rotor speed between defined frequency tiers, facilitating separable frequency-division multiplexing across a dense swarm. Validating whether these tonal signatures maintain a sufficient SNR above the broadband turbulent floor at elevated frequencies constitutes the critical

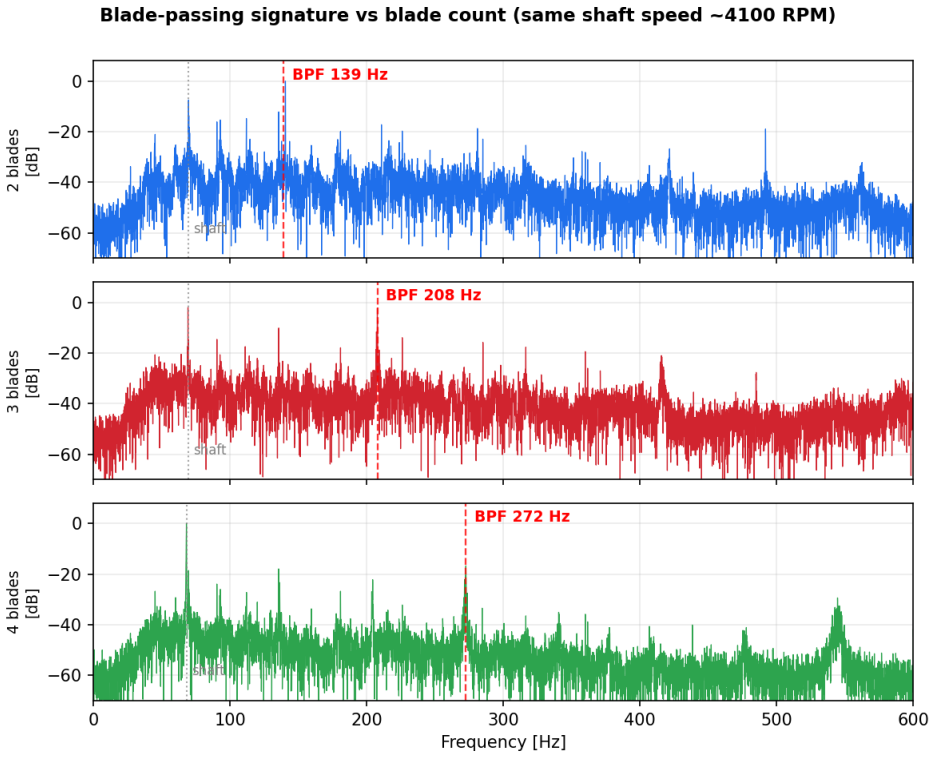

Figure 9.1: Blade-passing signature against blade count, measured on the propeller bench at near-equal shaft speed. The three propellers turn at about the same rate, so the shaft line sits near 69 Hz in all three, but the blade-passing frequency moves with the blade count: 139 Hz for two blades, 208 Hz for three, and 272 Hz for four. Different propellers therefore carry distinct signatures at the same rotation speed, which is the basis for identifying a drone by its wake. All three frequencies lie above the 100 Hz limit of the 200 Hz inertial unit, so none is accessible to the current sensor.

empirical hurdle for this direction.

The final trajectory proposes synthesizing these amplitude and frequency modalities into a singular, highly specialized bio-inspired sensor architecture: this concept envisions a sensor mathematically modeled as a heavily damped mechanical oscillator, analogous to biological hair-cell flow sensors [46]. The low-frequency, steady-state mean deflection of the oscillator would provide the scalar amplitude data currently utilized for spatial pursuit localization, while its high-frequency, underdamped resonant response,precisely tuned to the kilohertz domain, would isolate the high-bandwidth blade-passing spectral signatures. Thus, the distinct functionalities currently requiring disparate low-speed and high-speed inertial units

would be seamlessly integrated into the dual dynamic regimes (mean deflection and resonance) of a single physical element. While analogous multimodal sensors have previously aggregated airflow, acceleration, and angular velocity metrics [47, 48], this proposal pioneers their application for relative-state spatial estimation and high-frequency target identification via an adjacent vehicle's wake, fundamentally diverging from existing flight-by-feel applications that strictly measure a vehicle's endogenous airflow. Developing this theoretical oscillator model, verifying the emergence of the tonal signature at the sensor level, and optimizing the array topology across the airframe define the novel, high-impact research program inaugurated by this thesis.

Bibliography

[1] Leonard Bauersfeld, Koen Muller, Dominic Ziegler, Filippo Coletti, and Davide Scaramuzza. *Robotics meets Fluid Dynamics: A Characterization of the Induced Airflow below a Quadrotor as a Turbulent Jet*. 2024. arXiv: 2403.13321 [cs.RO]. URL: https://arxiv.org/abs/2403.13321 (cit. on pp. 1, 5, 10–12, 18, 20, 30–32, 43, 44).

[2] Henry Smith, Ajay Shankar, Jennifer Gielis, Jan Blumenkamp, and Amanda Prorok. «SO(2)-Equivariant Downwash Models for Close Proximity Flight». In: *IEEE Robotics and Automation Letters* 9.2 (Feb. 2024), pp. 1174–1181. ISSN: 2377-3774. DOI: 10.1109/lra.2023.3337701. URL: http://dx.doi.org/10.1109/LRA.2023.3337701 (cit. on pp. 2, 18, 22).

[3] Guanya Shi, Wolfgang Hönig, Xichen Shi, Yisong Yue, and Soon-Jo Chung. «Neural-Swarm2: Planning and Control of Heterogeneous Multirotor Swarms using Learned Interactions». In: *IEEE Transactions on Robotics* 38.2 (2022), pp. 1063–1079. DOI: 10.1109/TRO.2021.3098436 (cit. on pp. 2, 18, 21, 22).

[4] Jennifer Gielis, Ajay Shankar, and Amanda Prorok. «A Critical Review of Communications in Multi-robot Systems». In: *Current Robotics Reports* 3.4 (2022), pp. 213–225. DOI: 10.1007/s43154-022-00090-9 (cit. on pp. 2, 3).

[5] Leonardo Cerruti. «Bio-inspired Implicit Communication and Control for Quadcopter Drone Fault-Tolerant Formation Flight». Relatore: Alessandro Rizzo. Corso di laurea magistrale in Ingegneria Elettronica. Master's thesis. Politecnico di Torino, 2024. URL: https://webthesis.biblio.polito.it/31710/1/tesi.pdf (cit. on pp. 4, 13, 15, 26, 27, 45, 75, 81, 100).

[6] Kwang-Kyo Oh, Myoung-Chul Park, and Hyo-Sung Ahn. «A survey of multi-agent formation control». In: *Automatica* 53 (2015), pp. 424–440. DOI: 10.1016/j.automatica.2014.10.022 (cit. on p. 8).

[7] Steven van der Helm, Mario Coppola, Kimberly McGuire, and Guido de Croon. «On-board range-based relative localization for micro air vehicles in indoor leader–follower flight». In: *Autonomous Robots* 44 (2020). DOI: 10.1007/s10514-019-09843-6 (cit. on pp. 8, 24).

[8] Tobias Nägeli, Christian Conte, Alexander Domahidi, Manfred Morari, and Otmar Hilliges. «Environment-independent formation flight for micro aerial vehicles». In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2014, pp. 1141–1146. DOI: 10.1109/IROS.2014.6942701 (cit. on p. 8).

[9] Zhen Wang, Dianxi Shi, Chunping Qiu, Songchang Jin, Tongyue Li, Ziteng Qiao, and Yang Chen. «VecMapLocNet: Vision-based UAV localization using vector maps in GNSS-denied environments». In: *ISPRS Journal of Photogrammetry and Remote Sensing* 225 (2025), pp. 362–381. ISSN: 0924-2716. DOI: 10.1016/j.isprsjprs.2025.04.009. URL: https://www.sciencedirect.com/science/article/pii/S0924271625001455 (cit. on pp. 8, 24).

[10] Louiza Hamada, Pascal Lorenz, and Aicha Djerouni. «Light-Guided Autonomous Drone Navigation for Indoor GPS-Denied Environments». In: *Drones and Autonomous Vehicles* 3.2 (2026), p. 10007. DOI: 10.70322/dav.2026.10007 (cit. on p. 8).

[11] J. Gordon Leishman. *Principles of Helicopter Aerodynamics*. 2nd ed. Cambridge Aerospace Series. Cambridge, UK: Cambridge University Press, 2006 (cit. on pp. 9, 10).

[12] P. B. S. Lissaman and Carl A. Shollenberger. «Formation Flight of Birds». In: *Science* 168.3934 (1970), pp. 1003–1005 (cit. on p. 10).

[13] Henri Weimerskirch, Julien Martin, Yannick Clerquin, Peggy Alexandre, and Sarka Jiraskova. «Energy saving in flight formation». In: *Nature* 413 (2001), pp. 697–698 (cit. on p. 10).

[14] Steven J. Portugal, Tatjana Y. Hubel, Johannes Fritz, S. Heese, D. Trobe, B. Voelkl, S. Hailes, A. M. Wilson, and J. R. Usherwood. «Upwash exploitation and downwash avoidance by flap phasing in ibis formation flight». In: *Nature* 505 (2014), pp. 399–402 (cit. on p. 10).

[15] Anoop Kiran, Narek Harutyunyan, Nora Ayanian, and Kenneth Breuer. «Downwash Dynamics: Impact of Separation on Forces, Moments, and Velocities for Dense Quadrotor Flight». In: *AIAA AVIATION Forum and ASCEND*. 2024, p. 4145. DOI: 10.2514/6.2024-4145 (cit. on pp. 12, 35, 38, 42).

[16] Jennifer Gielis, Ajay Shankar, Ryan Kortvelesy, and Amanda Prorok. «Modeling Aggregate Downwash Forces for Dense Multirotor Flight». In: *Experimental Robotics (ISER)*. 2024. DOI: 10.1007/978-3-031-63596-0_35 (cit. on pp. 17, 18, 22).

[17] Karan P. Jain, Trey Fortmuller, Jaeseung Byun, Simo A. Mäkiharju, and Mark W. Mueller. «Modeling of aerodynamic disturbances for proximity flight of multirotors». In: *International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2019, pp. 1261–1269. DOI: 10.1109/ICUAS.2019.8798116 (cit. on pp. 18, 22, 23, 26, 27).

[18] Derrick Yeo, Elena Shrestha, Derek A. Paley, and Ella M. Atkins. «An empirical model of rotorcraft UAV downwash for disturbance localization and avoidance». In: *AIAA Atmospheric Flight Mechanics Conference*. 2015, p. 1685. DOI: 10.2514/6.2015-1685 (cit. on p. 18).

[19] Ajay Shankar, Heedo Woo, and Amanda Prorok. «Docking Multirotors in Close Proximity using Learnt Downwash Models». In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2024. DOI: 10.1109/ICRA48891.2023.10161567 (cit. on pp. 18, 22).

[20] Iris David Du Mutel de Pierrepont Franzetti, Riccardo Parin, Elisa Capello, Matthew J. Rutherford, and Kimon P. Valavanis. «Wall effect evaluation of small quadcopters in pressure-controlled environments». In: *Journal of Intelligent & Robotic Systems* 110.125 (2024). DOI: 10.1007/s10846-024-02155-7 (cit. on p. 19).

[21] Darius J. Carter, Lauren Bouchard, and Daniel B. Quinn. «Influence of the Ground, Ceiling, and Sidewall on Micro-Quadrotors». In: *AIAA Journal* 59.4 (2021), pp. 1398–1405. DOI: 10.2514/1.J059787 (cit. on p. 19).

[22] Takuzumi Nishio, Moju Zhao, Fan Shi, Tomoki Anzai, Kento Kawaharazuka, Kei Okada, and Masayuki Inaba. «Stable Control in Climbing and Descending Flight under Upper Walls using Ceiling Effect Model based on Aerodynamics». In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2020, pp. 172–178 (cit. on p. 19).

[23] Chunyue Zhu, Yujia Tian, Preben Buchhave, and Clara M. Velte. «Similarity scaling of the axisymmetric turbulent jet». In: *arXiv preprint arXiv:2206.03979* (2022) (cit. on p. 20).

[24] R. Kapulla, K. H. Manohar, S. Paranjape, and D. Paladino. «Self-similarity for statistical properties in low-order representations of a large-scale turbulent round jet based on the proper orthogonal decomposition». In: *Experimental Thermal and Fluid Science* 123 (2021), p. 110320. DOI: 10.1016/j.exptherm flusci.2020.110320 (cit. on p. 20).

[25] Stephen B. Pope. *Turbulent Flows*. Cambridge, UK: Cambridge University Press, 2000. DOI: 10.1017/CBO9780511840531 (cit. on p. 20).

[26] Anoop Kiran, Narek Harutyunyan, Nora Ayanian, and Kenneth Breuer. «Downwash Dynamics: Impact of Separation on Forces, Moments, and Velocities for Dense Quadrotor Flight». In: *arXiv preprint arXiv:2507.09463* (2025) (cit. on p. 20).

[27] Igor Geles, Leonard Bauersfeld, Angel Romero, J. Xing, and Davide Scaramuzza. «Superhuman Safe and Agile Racing through Multi-Agent Reinforcement Learning». In: *arXiv preprint arXiv:2605.22748* (2026) (cit. on p. 21).

[28] Ruiqi Zhang, Dingqi Zhang, and Mark W. Mueller. «ProxFly: Robust Control for Close Proximity Quadcopter Flight via Residual Reinforcement Learning». In: *arXiv preprint arXiv:2409.13193* (2024) (cit. on p. 22).

[29] Pei-An Hsieh, Kong Yao Chee, and M. Ani Hsieh. «Online Adaptation for Flying Quadrotors in Tight Formations». In: *arXiv preprint arXiv:2506.17488* (2025) (cit. on p. 23).

[30] Jorge Peña Queralta, Carmen Martınez Almansa, Fabrizio Schiano, Dario Floreano, and Tomi Westerlund. «UWB-based System for UAV Localization in GNSS-Denied Environments: Characterization and Dataset». In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. arXiv:2003.04380 2020 (cit. on p. 23).

[31] Peihan Zhang, Gang Chen, Yuzhu Li, and Wei Dong. «Agile Formation Control of Drone Flocking Enhanced with Active Vision-based Relative Localization». In: *arXiv preprint arXiv:2108.05505* (2021) (cit. on p. 24).

[32] Shiwen Lei, Andreas Jakobsson, and Zhiqin Zhao. «A CFAR Adaptive Matched Detector for Target Detection in Non-Gaussian Noise with Inverse Gamma Texture». In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. arXiv:1705.04557. 2017 (cit. on p. 25).

[33] Sushila Dhaka, Jane-Hwa Huang, Chih-Min Yu, and Li-Chun Wang. «SNR-Adaptive Optimal Threshold Design for Energy Detection in Dynamic Spectrum Access». In: *arXiv preprint arXiv:2607.00754* (2026) (cit. on pp. 25, 26).

[34] Abdelmadjid Maali, Ammar Mesloub, Mustapha Djeddou, Hassane Mimoun, Genevieve Baudoin, and Abdelaziz Ouldali. «Adaptive CA-CFAR Threshold for Non-Coherent IR-UWB Energy Detector Receivers». In: *IEEE Communications Letters* 13.12 (2009), pp. 959–961. DOI: 10.1109/LCOMM.2009.12.091579 (cit. on p. 26).

[35] Alexandru Martian, Mahmood Jalal Ahmad Al Sammarraie, Călin Vlădeanu, and Dimitrie C. Popescu. «Three-Event Energy Detection with Adaptive Threshold for Spectrum Sensing in Cognitive Radio Systems». In: *Sensors* 20.13 (2020), p. 3614. DOI: 10.3390/s20133614 (cit. on p. 26).

[36] Janne J. Lehtomäki, Johanna Vartiainen, Risto Vuohtoniemi, and Harri Saarnisaari. «Adaptive FCME-based threshold setting for energy detectors». In: *4th International Conference on Cognitive Radio and Advanced Spectrum Management (CogART)*. 2011. DOI: 10.1145/2093256.2093289 (cit. on p. 26).

[37] Pavel Petráček, Viktor Walter, Tomáš Báča, and Martin Saska. «Bio-inspired compact swarms of unmanned aerial vehicles without communication and external localization». In: *Bioinspiration & Biomimetics* 16.2 (2021), p. 026009. DOI: 10.1088/1748-3190/abc6b3 (cit. on pp. 27, 28).

[38] Naser El-Sheimy, Haiying Hou, and Xiaoji Niu. «Analysis and Modeling of Inertial Sensors Using Allan Variance». In: *IEEE Transactions on Instrumentation and Measurement* 57.1 (2008), pp. 140–149. DOI: 10.1109/TIM.2007.908635 (cit. on p. 37).

[39] Hermann Rohling. «Radar CFAR Thresholding in Clutter and Multiple Target Situations». In: *IEEE Transactions on Aerospace and Electronic Systems* AES-19.4 (1983), pp. 608–621. DOI: 10.1109/TAES.1983.309350 (cit. on p. 48).

[40] Paul R. Kalata. «The Tracking Index: A Generalized Parameter for α-β and α-β-γ Target Trackers». In: *IEEE Transactions on Aerospace and Electronic Systems* AES-20.2 (1984), pp. 174–182. DOI: 10.1109/TAES.1984.310438 (cit. on p. 54).

[41] N. A. Shneydor. *Missile Guidance and Pursuit: Kinematics, Dynamics and Control*. Chichester, England: Horwood Publishing, 1998 (cit. on pp. 57, 58, 63, 64, 99).

[42] Paul Zarchan. *Tactical and Strategic Missile Guidance*. 6th ed. Progress in Astronautics and Aeronautics. Reston, VA: AIAA, 2012 (cit. on pp. 57, 64).

[43] D. Swaroop and J. K. Hedrick. «String Stability of Interconnected Systems». In: *IEEE Transactions on Automatic Control* 41.3 (1996), pp. 349–357. DOI: 10.1109/9.486636 (cit. on pp. 90, 102).

[44] Reza Olfati-Saber, J. Alex Fax, and Richard M. Murray. «Consensus and Cooperation in Networked Multi-Agent Systems». In: *Proceedings of the IEEE* 95.1 (2007), pp. 215–233. DOI: 10.1109/JPROC.2006.887293 (cit. on pp. 91, 99, 102).

[45] Reza Olfati-Saber and Richard M. Murray. «Consensus Problems in Networks of Agents with Switching Topology and Time-Delays». In: *IEEE Transactions on Automatic Control* 49.9 (2004), pp. 1520–1533. DOI: 10.1109/TAC.2004.834113 (cit. on pp. 91, 102).

[46] Alex C. Hollenbeck, Ramana Grandhi, John H. Hansen, and Alexander M. Pankonien. «Bioinspired Artificial Hair Sensors for Flight-by-Feel of Unmanned Aerial Vehicles: A Review». In: *AIAA Journal* 61.12 (2023), pp. 5206–5231. DOI: 10.2514/1.J062931 (cit. on p. 103).

[47] Bo Yang, Di Hu, and Lei Wu. «Design and Analysis of a New Hair Sensor for Multi-Physical Signal Measurement». In: *Sensors* 16.7 (2016), p. 1056. DOI: 10.3390/s16071056 (cit. on p. 104).

[48] Lansheng Zhang, Zheyi Hang, and Huan Hu. «Bio-inspired artificial hair flow sensors: a comprehensive review of design, fabrication, enhancements, and applications». In: *Microsystems & Nanoengineering* 11.1 (2025), p. 88. DOI: 10.1038/s41378-025-00895-6 (cit. on p. 104).