



Politecnico di Torino

M.Sc. Mechatronic Engineering

A.a. 2025/2026

Graduation Session September 2026

Gesture-Based Follower-to-Leader Feedback for Bidirectional Implicit UAV Communication

Supervisor:

prof. Alessandro Rizzo

prof. Igor Paprotny

Candidate:

Tommaso Andina

Table of Contents

List of Figures v

Introduction xiii

1 From One-Way Wake Sensing to Bidirectional Interaction 1
- 1.1 Overview of the baseline leader-to-follower system 1
- 1.2 Why one-way implicit communication is insufficient 2
- 1.3 Feedback requirements: acknowledgement and error reporting . . . 4
- 1.4 Follower motion as an information-bearing signal 6
- 1.5 Design constraints for a return channel in close-proximity flight . . 7
- 1.6 System-level requirements for bidirectional operation 9

2 State of the Art on Motion-Based Communication and Robust Decoding 11
- 2.1 Implicit communication in multi-robot and multi-agent systems . . 11
- 2.2 Motion gestures as symbolic communication primitives 12
- 2.3 Signal detection in noisy dynamical systems 13
- 2.4 Matched filtering and energy-based detection 14
- 2.5 Communication protocols with acknowledgement and error handling 16
- 2.6 Positioning of this thesis with respect to existing work 17

3 Modeling the Follower-to-Leader Gesture Channel 20
- 3.1 Physical Principle of the Return Channel 20
- 3.2 Gesture Signal Model and Design Variables 22
- 3.3 Frequency, Amplitude and Duration Constraints 27
- 3.4 Noise Sources in Leader-Side Gesture Observation 30
- 3.5 Detectability of Sinusoidal Altitude Gestures 32
- 3.6 False Alarm and Missed Detection Mechanisms 34
- 3.7 Timing Constraints for Reliable Symbolic Communication 36

4 Design of the Gesture-Based Return Channel 40
4.1 Altitude Gesture Generator . . . 40
4.2 Sinusoidal Modulation and Selected Operating Frequency . . . 42
4.3 Symbol Alphabet: HEARTBEAT, ACK and ERROR . . . 43
4.4 Symbol Timing and Guard Intervals . . . 45
4.5 Encoding Strategy for Follower Status Feedback . . . 47
4.6 Constraints Imposed by Flight Stability and Safety . . . 50
4.7 Implementation of the Gesture Generator on the Follower UAV . . 52

5 Leader-Side Decoding and Protocol Logic 54
5.1 I/Q Matched Filter for Gesture Detection . . . 54
5.2 Energy Detector and Decision Threshold . . . 55
5.3 Windowing, Synchronization and Symbol Validation . . . 57
5.4 Finite-State Machine for Leader-Side Interpretation . . . 59
5.5 ACK and ERROR Handling Logic . . . 61
5.6 Protection Against False Transitions . . . 63
5.7 Recovery from Missed or Corrupted Symbols . . . 65
5.8 Computational Implementation of the Decoder . . . 66

6 Integration with the Forward Wake-Sensing Channel 68
6.1 Role of Leader-to-Follower Sensing in the Shared System . . . 68
6.2 Bidirectional Timing Between Forward Commands and Return Feedback . . . 70
6.3 Cross-Channel Interference Analysis . . . 73
6.4 Coordination Between Sensing Updates and Gesture Transmission . 74
6.5 System-Level Behavior During Nominal and Error Conditions . . . 76

7 Simulation-Based Validation of the Return Channel 79
7.1 Simulation Setup and Measurement Pipeline . . . 79
7.2 Validation of Gesture Generation . . . 81
7.3 Matched-Filter Response Under Controlled Conditions . . . 83
7.4 Decoder Accuracy in Simulated Flight Conditions . . . 85
7.5 Robustness to Disturbances and Geometric Variation . . . 87
7.6 FSM Behaviour During ACK, ERROR and Liveness Events . . . 90

8 Protocol Scalability in Multi-UAV Formations 93
8.1 Extending the Return Channel from One Follower to Many . . . 93
8.2 Addressing the Announcement . . . 97
8.3 Passive Identification Through Rotor Signatures . . . 100
8.4 The Grant Cycle and Its Timing Budget . . . 105
8.5 Feedback Topology of a Multi-Follower Formation . . . 108
8.6 Chains and Hierarchies . . . 111

8.7 Reliability of the Granted Protocol 113
8.8 Operating Envelope of the Multi-Follower Protocol 114

9 Conclusions 116
9.1 Summary of the Individual Contribution 116
9.2 Main Experimental Findings 117
9.3 Limitations of the Gesture-Based Feedback Channel 119
9.4 Future Work on Robust Bidirectional Implicit Communication . . . 120

Bibliography 122

List of Figures

1.1 Baseline leader-follower formation 2

1.2 The three gaps left open by a one-way channel. Information travels downward with the wake, so everything the Leader knows about the formation is something it produced itself. The three questions on the right are not independent shortcomings but three faces of the same limitation: a channel that moves in one direction only cannot carry the information that only the receiving end is in a position to generate. Each of them motivates one of the three symbols defined in Section 1.3. 4

1.3 Sensing and reply configurations 7

1.4 The complete bidirectional system. The formation alternates between two configurations: in the sensing configuration the Follower tracks the Leader through its wake, and in the reply configuration, entered through a role-swap, the vertical ordering is inverted for the length of one transmission so that the Follower's own downwash reaches the Leader. Three symbols travel in the return direction, distinguished by duration alone, and each maps to one of the three gaps of Section 1.2. This figure is the reference to which Chapters 3 to 6 add detail. 10

2.1 The principle of matched filtering. A waveform whose shape is known in advance can be recovered from noise that dominates it in amplitude by correlating the received signal against a copy of that shape: contributions that resemble the template accumulate coherently while the rest averages toward zero, so a peak appears at the instant of arrival even when nothing is visible in the received trace itself. This is the property the return channel relies on, and the reason its gesture is a waveform of fixed shape rather than an arbitrary displacement. 15

2.2 Positioning of the return channel with respect to the approaches surveyed in this chapter, along the two dimensions that separate them: whether information is carried by a dedicated signal or by the physics the vehicles already produce, and whether the resulting channel runs in one direction or two. The approaches that provide feedback do so over a dedicated link, and those that avoid a dedicated link provide no feedback. The upper-right quadrant is what this thesis occupies, and Table 2.2 states the same comparison in the four dimensions used throughout the chapter. 19

3.1 Geometry of the role-swap manoeuvre 22

3.2 Overview of the follower-to-leader signal chain 22

3.3 Axial decay of the Follower downwash. The fitted law $C/(z - z_0)$ is shown against the anemometric data (mean ± SD, $n = 3$), with the $1/z$ far-field law overlaid for comparison: the measured velocity falls by a factor of 2.1 while the axial distance increases by a factor of 5, which the $1/z$ law cannot reproduce. 26

3.4 Radial profiles of the Follower downwash at the five axial stations. The fitted self-similar shape uses a half-velocity radius $r_{1/2} = 9.6\,\mathrm{cm}$ held constant across all stations, the on-axis value at each station being taken from the axial fit of Figure 3.3. Open markers denote readings below the instrument's resolution floor. 27

3.5 The admissible region for the two gesture parameters. The frequency is bounded below by the band occupied by ordinary formation-keeping motion and above by the closed-loop bandwidth of the Follower's altitude loop, both identified in Section 7.1. The amplitude is bounded above by the vertical clearance that must remain during a role-swap and below by the disturbance the Leader must be able to detect; both bounds rise with frequency, because the same loop that limits f_0 also attenuates what it executes, so a given excursion costs more commanded amplitude the closer f_0 sits to the bandwidth. The chosen pair lies inside the region with margin on every side. The lower boundary is indicative, since the noise floor on dz is an estimate rather than a quantity identified from a hover log. 30

3.6 Illustrative effect of the high-pass stage on the Leader's vertical velocity estimate. Above, a simulated dz in which a one-second gesture at f_0 is superimposed on a sensor bias and a slow drift an order of magnitude larger than itself: the gesture is present but not visible, because the quantity is dominated by components that carry no information. Below, the same signal after the filter, which removes the bias and the drift while passing f_0 essentially unchanged, leaving the gesture standing clear of the residual noise. The traces are generated from a standalone simulation for illustration; filtered signals from the full model are reported in Chapter 7. 32

3.7 Why detection uses a quadrature pair. The same gesture is presented to the correlator at every arrival phase, a quantity the Leader cannot know in advance because it depends on when the Follower happened to begin transmitting. A single real correlation returns an energy that collapses entirely at the phases where the gesture is orthogonal to the reference, so a detector built on it would miss a transmission that was present and undistorted. The quadrature pair holds the energy within roughly $\pm 11\,\%$ of its mean. The residual ripple is not a property of the method but of the window length: the correlation spans two and a half cycles of f_0 rather than an integer number, and an integer-cycle window would remove it. That margin is small against the detection margin of Section 3.5, whereas the collapse of the real correlation is not. 34

3.8 Placement of the detection threshold between the two failure mechanisms. The threshold cannot be derived from the statistics of the noise-floor energy alone, because consecutive matched-filter outputs share all but one sample of their correlation window and are therefore strongly correlated; it was calibrated instead against the two levels shown, both observed in simulation. The margins above and below are deliberately asymmetric: the smaller one guards against reporting a gesture that was never sent, the larger against failing to report one that was, and the two cannot be widened at the same time. 36

3.9 Nominal gesture duration against observed peak width 38

4.1 Block diagram of the altitude gesture generator: a trigger and symbol code select a nominal duration at onset, driving a shared duration and phase state that determines both the tapered sinusoid produced at each sample and the gesture-busy flag consumed by the arbitration logic of Section 4.5. The resulting offset is added to the Follower's nominal altitude reference before being sent to the altitude controller. 41

4.2 Commanded altitude offset for a single ACK gesture, generated at $F_s = 200\,\text{Hz}$. The Tukey envelope reaches full amplitude only after the first 20 samples and returns to zero over the last 20, so that the transition into and out of the gesture is continuous and the altitude controller is never asked to track a discontinuity. 43

4.3 The three symbols of the alphabet on a common time axis. Amplitude, frequency and taper are identical in all three cases; only the duration for which the sinusoid is sustained distinguishes one symbol from another. 45

4.4 Why the guard interval is 300 samples. Each pair of blue blocks is two ACK symbols transmitted with the spacing shown; the bars beneath are the energy excursions they produce at the receiver, each wider than the symbol that generated it because the correlation window slides onto and off the signal. At 100 samples the two excursions overlap and the receiver reports a single event of the wrong duration class. At 150 they separate and both symbols are recovered, which locates the boundary empirically. The value adopted is twice that minimum, so that the spacing continues to clear both the spreading and the debounce interval as the excursions narrow at larger role-swap separations. 47

4.5 Arbitration of the three transmission conditions into a single stream of symbols. Only the forward-channel signal passes through a latch: an anomaly arriving while the generator is busy is not queued. The busy flag returned by the generator is what gates the whole selection. 49

4.6 Vertical clearance during a role-swap. The gesture oscillates the Follower about the nominal separation, so the instantaneous gap closes by the excursion actually executed rather than by the amplitude commanded. At the lower end of the recommended range the gap reaches 11.3 cm at the bottom of every cycle, of which the two airframes occupy 8.0 cm. The remaining clearance is real but not generous, and it is exhausted at almost the same separation at which the classification margin is, which is why the recommended range stops short of the interval over which classification alone would still be correct. 52

5.1 Structure of the in-phase/quadrature matched filter on the Leader. The high-pass-filtered vertical velocity estimate is held in a sliding window matching the correlation length, and correlated against two fixed unit-energy references at f_0. Squaring and summing the two correlation values yields an energy measure that does not depend on the phase at which the gesture arrives, which is the property established in Section 3.5 and the reason the quadrature pair is used rather than a single real correlation. The whole chain is evaluated once per control step; what becomes of E is the subject of Section 5.2. 55

5.2 State machine internal to the energy detector, distinct from the protocol machine of Figure 5.4. Guards are shown in square brackets: nothing is reported while energy stays below threshold, the width accumulates for as long as a crossing persists, and the symbol is decided only once the crossing closes. Both protections act inside this loop, so a discarded crossing never reaches the protocol layer. . 57

5.3 Why the blanking period is 400 samples rather than 200. At start-up the high-pass stage of Section 3.4 sees the Leader's standing vertical bias as a step, and its response decays with a time constant of $\tau \approx 99$ samples. The matched filter, correlating against that decaying transient, produces an energy that exceeds the detection threshold by orders of magnitude and remains above it well after the correlation buffer has filled. Sizing the blanking interval to the buffer alone would therefore leave the Leader acting on its own start-up transient; sizing it to four time constants of the filter that precedes the buffer brings the energy back below threshold before any output is accepted. 59

5.4 Leader-side protocol state machine, drawn as a statechart: transitions are labelled *event / action*, and the dashed state is transient, occupied for a single sample. Only acknowledgements are interpreted through the states, since an ACK is meaningful solely as the answer to an outstanding command. Error reporting and liveness are handled outside the state graph, because neither is a reply to anything the Leader has done. 61

5.5 Reaction stage on the Leader. Loss of liveness takes precedence over both other conditions. The two halts differ in kind: error recovery is bounded and can be cut short by an acknowledgement, whereas a lost Follower halts the formation with no timeout, since silence carries no evidence that advancing has become safe. 63

6.1 Augmented graph of the formation 69

6.2 Latency of the bidirectional exchange 71

6.3 The two operating modes and their alternation 74

6.4 One period of the sensing/reply cycle. The formation spends most of its time in the sensing configuration, where the Follower tracks the Leader and updates its range estimate; a heartbeat is initiated after 5.0 s of channel silence, and for the length of the reply window the vertical ordering is inverted, so the Follower cannot sense the Leader and holds its last valid estimate rather than acquiring an incorrect one. Because the interval between consecutive heartbeats is half the window within which the Leader expects one, a single lost symbol does not clear the liveness flag. 78

7.1 Simulation pipeline and recording points 81

7.2 Commanded and executed altitude offset for a single ACK gesture. The altitude loop tracks the reference with attenuation and lag rather than reproducing it: the peak excursion falls from 12.0 cm to 8.7 cm, and the executed waveform trails the commanded one by 69°. The taper survives the loop, so the executed trajectory is still continuous at both ends, and the residual tail after the symbol ends is the loop settling rather than commanded motion. 83

7.3 Matched-filter energy for the three symbols, aligned at symbol onset and drawn on a logarithmic scale. The peak reached differs by more than a factor of five across the alphabet although amplitude and frequency are identical: HEARTBEAT clears the threshold by 2.96 times, ACK by 14.4 and ERROR by 16.3. The width of each excursion, not its height, is what the classifier reads. 85

7.4 Detection and correct-classification rate against role-swap separation. Detection is complete across the shaded window and well beyond its lower edge; classification is the binding quantity at both ends, falling below 18 cm as the HEARTBEAT width grows past the 235-sample boundary and above 30 cm as the HEARTBEAT energy approaches λ and its excursion fragments. The same symbol sets both edges, through opposite mechanisms. 89

7.5 Observed peak width of the three symbols against role-swap separation, with the 235- and 400-sample classification boundaries. Shaded bands span one standard deviation over the noise realisations. Every width decreases monotonically with separation; the HEARTBEAT trace is both the lowest and the first to lose its margin at either end of the range, which is what sets both edges of the operating window. 90

7.6 State and event traces for the four scenarios on a common time axis. Shaded bands mark detected symbols. In the nominal exchange the machine leaves Waiting Ack on the acknowledgement; in the silent case it leaves at the timeout with no output; the spontaneous error is acted on outside the state graph; and the liveness flag clears 2,000 samples after the last heartbeat, independently of whether a command is outstanding. 92

8.1 The reply configuration as a resource of unit capacity. The shaded region is the set of positions from which a Follower's gesture reaches the Leader with the margin established in Section 7.5: bounded axially by the recommended separation range and radially by the half-velocity radius of Section 3.2. Its extent is of the order of a single airframe, so a second Follower cannot be placed inside it, and one placed outside delivers a fraction of the signal the detector was dimensioned for. 95

8.2 An addressed announcement in a three-vehicle formation. The transmission is broadcast by construction, since the wake descends onto every Follower beneath the Leader, and the address is carried by its duration. Each Follower classifies the width of a detection it already produces; only the vehicle whose address matches interrupts its pursuit, while the others continue tracking. Nothing about the transmission other than its duration differs from the announcement already sent in the two-vehicle system. 99

8.3 Acoustic spectra of the three six-inch propellers at the same throttle command, and therefore at very nearly the same shaft speed. The shaft component stays where it is while the blade-passing component moves with the blade count, appearing at 139, 208 and 272 Hz for two, three and four blades against a shaft frequency of approximately 69 Hz: in each case the predicted integer multiple, to within one per cent. Because the three propellers come from the same parametric model and differ only in blade count, the displacement cannot be attributed to profile, chord or material. 102

8.4 Two signatures present at once. Each panel is the spectrum of the superposition of two recordings taken separately at the same throttle command; since the bench carries a single motor, two propellers could not be run simultaneously, and the superposition is formed numerically, which is admissible because the transform is linear. Both blade-passing components survive and remain individually resolvable, including in the closer case where they lie 65 Hz apart. The result is a statement about spectral separability, and not about how two physically adjacent jets would interact. 104
8.5 The grant cycle against the liveness window. With the manoeuvre duration estimated in Section 6.3, a cycle carrying one heartbeat occupies 5.3 s, of which the role-swap alone accounts for more than half and the symbol the Leader is waiting for accounts for less than a tenth. Two such cycles do not fit inside the window within which every Follower must be heard from, so the liveness budget inherited from the two-vehicle system admits one Follower and no more. What the protocol costs is not channel time but manoeuvre time. 107
8.6 Why a reply is not a local event in a chain. An intermediate vehicle senses the tier ahead and is sensed by the tier behind, so when it is granted the reply configuration and climbs, it simultaneously leaves the position from which the vehicle below was sensing it. Both suspensions are of the kind established in Section 6.4, in which a deprived Follower holds its last valid estimate rather than acquiring an incorrect one, but two links are affected rather than one. Propagating grants downward, so that a vehicle granted a slot does not itself issue one until it has returned to sensing, keeps at most one tier in the reply configuration at any instant. 112

Introduction

Implicit Communication in Close-Proximity Multi-UAV Systems

When several aerial vehicles operate close enough to one another that their wakes overlap, they are already exchanging information whether or not anyone has arranged for them to. A rotor displaces air, that displacement propagates, and a vehicle flying within it registers a disturbance on the same inertial sensors it uses to hold its own attitude. The usual engineering response is to treat this as noise and reject it. A smaller body of work treats it as a signal instead, and asks what a formation could do if the aerodynamic coupling that close flight makes unavoidable were read rather than suppressed.

The appeal of that position is not that it outperforms a radio link, because it does not. It is that it removes a dependency. A formation that coordinates through a wireless channel inherits everything that channel inherits: spectrum that may be contested or denied, a protocol stack, a power budget, and a mode of failure in which the vehicles remain perfectly capable of flying while becoming unable to fly together. A formation that coordinates through the physics it already produces has none of those, and the reason is simply that the carrier is not something the system provides but something the system cannot avoid emitting.

The price is severe and should be stated at the outset. The channel exists only where the wake reaches, which for a nano-quadrotor is a region measured in centimetres; it carries very little; it is slow; and it is shared with the flight dynamics it is embedded in, so that every transmission is also a manoeuvre. Work in this area is therefore not concerned with capacity but with whether anything useful can be carried at all, and with what a formation can be built to do with the little that can.

The Need for Feedback in One-Way Aerodynamic Architectures

The system this thesis extends belongs to that line of work. In it, a Leader flies ahead and a Follower beneath, and the Follower localises the Leader by sensing its downwash, tracking it without any exchange of messages and without any sensor beyond the inertial unit it already carries. The arrangement works, and its asymmetry is not an oversight but a consequence of the physics: a wake travels downward, so information travels downward with it.

That asymmetry is also the system's limitation, and it is total rather than partial. Everything the Leader knows about the formation is something the Leader itself produced. It cannot know whether a manoeuvre it commanded was perceived, because perception happens at the other end of a one-way channel. It cannot know whether the Follower is still behind it, because a Follower that has failed, lost stability or drifted out of the sensing region produces exactly the same signal as one that never existed, namely none. And it cannot be told that something has gone wrong, because an anomaly originating at the Follower has no direction in which to travel. These are not three separate shortcomings but three faces of one: a channel that moves in a single direction cannot carry the information that only the receiving end is in a position to generate.

The consequence is not merely that the Leader is uninformed. It is that the Leader continues. Lacking any report from the vehicle behind it, it remains committed to its own reference and advances toward it, and the distance between a formation and a member that has stopped following grows for as long as nothing interrupts it. In a system whose entire coupling mechanism has a range of a few centimetres, that separation becomes irrecoverable quickly.

Research Problem

The problem addressed here is therefore to establish a return path, from Follower to Leader, under two constraints that between them exclude almost every conventional answer.

The first is that no radio link may be used. This is not a matter of convenience: if an explicit wireless channel were acceptable, the problem would not exist, since closing the loop over one is a solved engineering task. The premise of implicit communication is that coordination should not depend on a channel that can be denied, and a return path that reintroduces that dependency would defeat the purpose of the system it is meant to complete.

The second is that no sensor or actuator may be added to either vehicle. The platform is a nano-quadrotor of a few tens of grams, with an inertial unit, a

downward rangefinder and a control loop that runs on a microcontroller. Anything the return channel needs must already be on board.

Together these leave one candidate. The Follower cannot send a message, but it can move, and its motion displaces air that reaches the Leader. If that motion is shaped deliberately rather than incidentally, the disturbance it produces carries information, and the Leader can recover it from the state estimate it already computes. The research question is whether a channel built this way can be made reliable enough to be worth having: whether a deliberate motion can be distinguished from ordinary flight, whether several distinct messages can be distinguished from one another, whether the resulting channel can coexist with the sensing channel it shares a formation with, and what a formation can usefully be built to do with it.

Individual Contribution of the Thesis

The thesis contributes a complete return channel, from physical carrier to protocol behaviour, together with the analysis that dimensions it. Five elements are specific to this work.

A channel model fitted from measurement rather than adopted. A dedicated anemometric campaign characterises the wake of the platform actually used, and the resulting decay law and radial profile are what every subsequent design parameter is dimensioned against.

A symbolic alphabet carried by duration alone. Three symbols, reporting liveness, acknowledgement and anomaly, are distinguished by how long the detected disturbance lasts rather than by its amplitude or its frequency, which makes the receiver insensitive to the attenuation the channel imposes.

A receiver designed for the specific difficulty this signal presents. A quadrature matched filter recovers a gesture whose arrival phase is unknown, a threshold placed between two measured levels replaces a noise model that does not apply, and a state machine converts symbols into decisions about whether the formation may continue to advance.

An analysis of how two channels share one formation. The return channel requires the vertical ordering of the vehicles to be inverted and the forward channel requires it not to be, so the two cannot operate at once; the thesis establishes the coordination that lets them alternate without either corrupting the other.

An extension to formations of more than two vehicles, with its limits quantified. The protocol extends without any change to its physical layer, detector, alphabet or state machine, and the analysis shows that what bounds the size of the formation is not the channel but the manoeuvre each reply requires.

Underlying all five is a claim about what the return channel is for, which the

thesis argues rather than assumes. It does not make the formation converge faster, and it does not change what the formation converges to. It makes convergence interruptible on evidence, which a one-way system cannot do.

Structure of the Thesis

Chapter 1 describes the baseline one-way system, establishes why it cannot be completed without a return path, and derives the requirements and constraints any such path must satisfy. Chapter 2 surveys the relevant literature along four narrowing levels, from implicit coordination in multi-agent systems to matched filtering and acknowledgement protocols, and positions this work with respect to it.

Chapter 3 models the channel. It characterises the wake from measurement, defines the gesture waveform, and derives the constraints that bound its frequency, amplitude and duration, together with the detectability and timing limits that follow. Chapter 4 designs what the channel carries: the symbol alphabet, the guard intervals that separate transmissions, the arbitration that resolves competing conditions, and the generator that produces the gesture on the Follower.

Chapter 5 builds the receiver, from the filtering and correlation that recover a gesture from the Leader's own velocity estimate through to the state machine that interprets the result and the reactions it commands. Chapter 6 places the two channels in one system and establishes the timing under which they alternate.

Chapter 7 validates the design, reporting the operating window, the classification accuracy across it, and the behaviour of the protocol under nominal and fault conditions. Chapter 8 extends the protocol to several Followers and establishes what bounds the size of the formation. Chapter 9 summarises what was built, what the measurements establish, where the limits lie, and what would have to be measured next.

Chapter 1

From One-Way Wake Sensing to Bidirectional Interaction

1.1 Overview of the baseline leader-to-follower system

The work presented in this thesis is part of a two-drone formation consisting of a Leader and a Follower flying at a fixed nominal vertical separation, with the Leader positioned above the Follower. Coordination between the two vehicles is entrusted entirely to their aerodynamic interaction, rather than to an explicit communication channel: no radio link, Bluetooth pairing, or GNSS-based coordination is used to exchange information between the two drones. The formation instead relies on the fact that the Leader's rotors generate a downwash field that the Follower, flying underneath, can sense through its own attitude dynamics.

At the level of detail relevant to this thesis, the mechanism can be summarized as follows: the Follower observes the disturbance induced by the Leader's downwash and translates it into a command following action. The specific signal processing pipeline that makes this possible, from the aerodynamic disturbance to a validated relative position estimate, is not the subject of this work and is developed in full in the companion thesis by Francesco Pettinato [1], which addresses the forward sensing channel from the Leader to the Follower. For the purposes of this chapter, it is sufficient to treat that channel as a given: the Follower perceives a disturbance and reacts to it, and this reaction is what allows the formation to hold together without any explicit exchange of messages.

The system described so far is the result of a parallel and coordinated effort with

a colleague, split along the two directions of information flow that a leader-follower formation requires. The companion thesis addresses the direction from the Leader to the Follower, that is, how the Follower senses and interprets the Leader's presence and motion. This thesis addresses the direction that remains open once that forward channel is in place: how the Follower can, in turn, communicate back to the Leader, again without introducing any explicit radio link or additional hardware. The two contributions are designed to fit together into a single bidirectional system, and the forward sensing channel summarized in this section is developed in detail in the companion thesis by Francesco Pettinato, to which the reader is referred for the underlying detection and estimation logic.

Precisely because this forward channel exists but travels in only one direction, the formation as described here has a structural limitation: the Leader has no means of knowing whether its commands are being received, whether the Follower is still operating correctly, or whether it is still present at all. The next section examines this limitation in detail and motivates the need for a return channel from the Follower to the Leader.

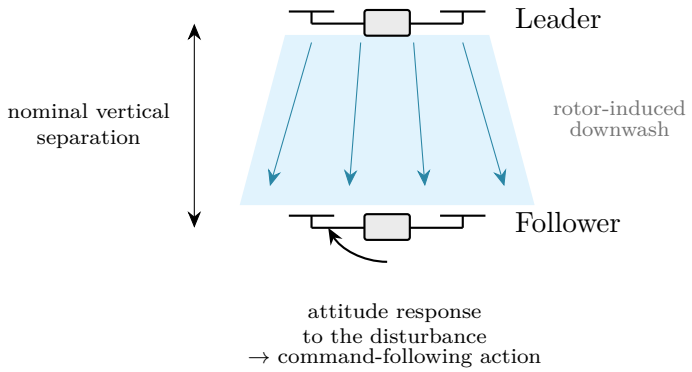


Figure 1.1: Baseline formation. The Leader flies above the Follower at a fixed nominal separation, and its rotors generate a downwash field that the Follower senses through its own attitude dynamics. No radio link, pairing or absolute positioning is exchanged between the two vehicles: the aerodynamic interaction is the only path along which information travels, and it travels in one direction only.

1.2 Why one-way implicit communication is insufficient

The formation described in the previous section works as long as the Follower correctly senses the Leader and reacts to it. From the Leader's point of view, however, the entire interaction is invisible: the Leader emits its motion, and

whatever happens to that motion once it reaches the Follower is, by construction, outside the Leader's reach. This asymmetry is not a minor implementation gap but a structural property of a one-way channel, and it is worth following its consequences through rather than simply stating that it exists, because each of them independently motivates a different requirement for the return channel developed in this thesis.

The first consequence concerns command execution. When the Leader performs a maneuver intended to convey information, it has no way of verifying that the maneuver was interpreted correctly, or interpreted at all. The Follower might fail to detect the pattern because of unfavorable relative geometry, might decode it into the wrong symbol because of noise, or might simply be too far from the sensing region for the disturbance to be perceptible. In every one of these cases, the Leader proceeds as if the command had been received, because nothing in the system tells it otherwise. A formation that depends on implicit commands for anything beyond simple following therefore needs some notion of acknowledgement; without it, the Leader is issuing instructions into a channel whose delivery it cannot confirm, which is a fragile position to be in as soon as the commands carry any operational weight.

The second consequence is more severe, because it does not depend on any single command having been sent. Even if the Leader never issues an explicit instruction, it still needs to know, on an ongoing basis, that the Follower is present, powered, and behaving as expected. A one-way sensing channel offers no such guarantee: the Follower could suffer a fault, lose stability, or drift out of the sensing region entirely, and the Leader would have no signal indicating that anything has changed. The formation would continue to behave, from the Leader's perspective, exactly as it did when everything was working, right up until the moment the two vehicles are far enough apart that the absence becomes obvious through other means, by which point recovery options are considerably more limited. A return channel that carries no information beyond command acknowledgement would still leave this gap open, since liveness needs to be established independently of whether a command happens to be in flight at a given moment.

The third consequence concerns situations that originate on the Follower's side rather than the Leader's. The Follower may encounter conditions worth reporting even when no command was issued and no fault has yet compromised its liveness: an anomaly in its own control loop, an unexpected disturbance, a state that calls for the Leader to adjust its behavior in response. In a one-way architecture, the Follower has no means of raising this kind of flag, since the physical channel available to it exists only in the direction the Leader can sense, and here the direction is reversed. The Leader remains unaware of the problem until its effects become visible through the formation's overall behavior, at which point the anomaly has already had time to propagate into a larger disturbance of the formation geometry.

Taken together, these three gaps, unconfirmed commands, undetected loss of

the Follower, and unreported anomalies, are not independent design choices to be addressed piecemeal but three faces of the same underlying limitation: a channel that only ever moves in one direction cannot, by definition, carry any of the information that only the receiving end is in a position to generate. Closing this gap requires a channel running from the Follower back to the Leader, built without reintroducing the explicit radio link that the rest of the system was designed to avoid. The requirements that such a channel must satisfy, acknowledgement, liveness, and anomaly reporting, are developed into concrete design constraints in the remainder of this chapter.

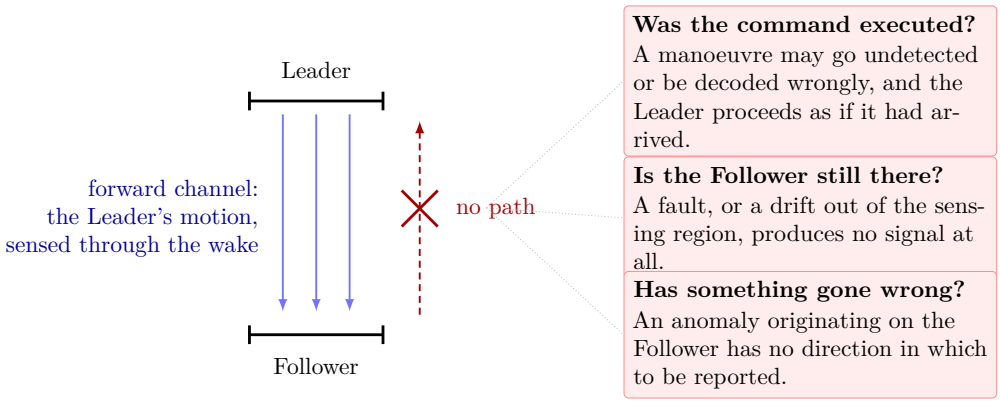


Figure 1.2: The three gaps left open by a one-way channel. Information travels downward with the wake, so everything the Leader knows about the formation is something it produced itself. The three questions on the right are not independent shortcomings but three faces of the same limitation: a channel that moves in one direction only cannot carry the information that only the receiving end is in a position to generate. Each of them motivates one of the three symbols defined in Section 1.3.

1.3 Feedback requirements: acknowledgement and error reporting

The three gaps identified in the previous section, each tied to a different blind spot in the Leader's knowledge of the Follower's state, translate directly into three functional requirements that any return channel must satisfy. It is useful to state them explicitly and independently of one another before turning, in later chapters, to how they are implemented, because each requirement constrains the design in a different way and none of them can be dropped in favor of the others without reopening one of the gaps discussed above. The first requirement is an acknowledgement of received commands: the Follower must be able to confirm, after

the fact, that a specific instruction sent by the Leader was correctly detected and is being acted upon. Without this requirement, the Leader has no way to distinguish a successfully delivered command from one that was missed, misinterpreted, or never within sensing range in the first place, and it is forced to assume success by default, which is precisely the fragile situation described in Section 1.2. The second requirement is the reporting of anomalies originating on the Follower's side, independently of whether a command was involved. The Follower must have a way to flag conditions that the Leader could not otherwise infer, so that the Leader is not left to discover a developing problem only once it has already affected the formation's overall behavior. The third requirement is an indication of liveness that does not depend on any command or anomaly being present. The Follower must periodically signal that it is still operating normally, even when there is nothing specific to report, so that the Leader can distinguish an uneventful, well-functioning Follower from one that has become silent because it is no longer able to communicate at all. Table 1.1 summarizes the three requirements together with their motivation and the consequence of leaving each one unmet. The same three-way structure reappears, in a more concrete form, when the symbol alphabet of the return channel is defined in Chapter 4; at this stage, however, the requirements are stated purely at the functional level, without committing to any particular signal or encoding.

Table 1.1: Functional requirements for the return channel.

Requirement	**Motivation**	**Consequence if missing**
Command acknowledgement	The Leader needs confirmation that a specific instruction was detected and is being executed by the Follower.	The Leader cannot distinguish a delivered command from one that was missed or misinterpreted, and must assume success without evidence.
Anomaly reporting	The Follower needs a way to flag conditions worth the Leader's attention, even when no command is in flight.	Problems originating on the Follower remain invisible to the Leader until they manifest as a disturbance of the formation itself.
Liveness indication	The Leader needs periodic evidence that the Follower is still operating, independently of commands or anomalies.	A silent but well-functioning Follower becomes indistinguishable from one that has lost the ability to communicate, delaying recognition of an actual loss.

1.4 Follower motion as an information-bearing signal

The three requirements introduced in the previous section point toward a single underlying idea, which forms the conceptual core of this thesis: if the Follower's motion can be shaped in a controlled way, it can itself become a signal, in exactly the same sense that the Leader's motion already is one for the forward channel. The forward channel exists because the Follower is positioned to sense the aerodynamic consequences of the Leader's presence; the return channel proposed here is built by placing the Follower, temporarily, where the same kind of aerodynamic consequence becomes observable to the Leader.

Concretely, the Follower is capable of modulating its own altitude in a controlled and repeatable way, superimposed on its ordinary flight behavior. Such a modulation is not, by itself, a communication channel: for it to carry information back to the Leader, the disturbance it produces must actually reach the Leader in an observable form. This is where the role-swap idea introduced conceptually in Chapter 3 becomes relevant already at this stage of the argument: by briefly exchanging vertical positions, with the Leader descending and the Follower rising above it, the same physical mechanism that lets the Follower sense the Leader in the nominal formation lets the Leader sense the Follower during the exchange. The situation is, in this sense, symmetric: the roles of transmitter and receiver are exchanged along with the roles of upper and lower vehicle, while the underlying physical principle, an aerodynamic disturbance sensed through onboard motion, remains the same on both legs of the interaction.

Figure 1.3 anticipates this idea in simplified form, showing the two configurations the formation alternates between: a sensing configuration, in which the Leader flies above and the Follower senses its downwash as in the baseline system, and a reply configuration, in which the two vehicles temporarily invert their relative altitude so that the Follower's own motion becomes observable to the Leader. No parameters of the underlying gesture, its frequency, amplitude, or duration, are introduced here; these choices belong to the detailed channel model developed in Chapter 3 and to the gesture design of Chapter 4. What matters at this stage is only the conceptual shift: the Follower is not merely a passive receiver of the Leader's implicit commands, but is also capable, when needed, of becoming a transmitter within the same aerodynamic medium.

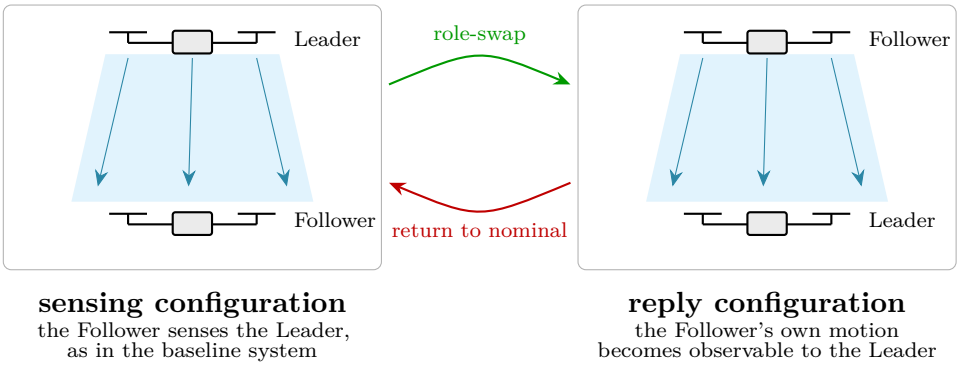


Figure 1.3: The two configurations the formation alternates between. In the sensing configuration the Leader flies above and its downwash reaches the Follower, as in the baseline system. By briefly exchanging vertical positions, the same physical mechanism operates in the opposite direction: the Follower's wake now propagates onto the Leader, so that a controlled modulation of the Follower's altitude becomes a signal the Leader can observe. Transmitter and receiver are exchanged along with upper and lower vehicle, while the underlying carrier remains the same.

This shift in perspective, from the Follower as a sensor to the Follower as an occasional transmitter, is what allows all three requirements of Section 1.3 to be met without introducing any component beyond what the formation already possesses: a controllable vertical channel on the Follower's side, and a sensing capability on the Leader's side that mirrors the one the Follower already relies on.

1.5 Design constraints for a return channel in close-proximity flight

The idea introduced in the previous section, letting the Follower's own motion carry information back to the Leader, only becomes a viable design once it is checked against the practical constraints that the rest of the system already operates under. These constraints are not specific to the return channel; they are inherited from the platform and from the formation architecture as a whole, and they delimit the space of solutions within which the channel developed in this thesis had to be designed. Table 1.2 summarizes them together with their motivation; the same constraints are invoked again in Chapter 3 to justify the specific choice of gesture frequency and amplitude, so it is worth stating them once, clearly, at this stage. The first constraint is that no additional hardware may be introduced. The entire premise of the formation, on both the forward and the return channel, is that implicit aerodynamic sensing removes the need for dedicated communication equipment; adding a sensor or actuator specifically for the return channel would undermine

the motivation for building it in the first place. The second constraint follows directly from the first and rules out radio links of any kind, including short-range options such as Bluetooth, as a means of closing the loop. If an explicit wireless channel were acceptable, the return channel described in this thesis would not be necessary; the entire design problem exists precisely because such a channel is excluded by assumption. The third constraint concerns coexistence with the forward channel. Whatever signal the Follower uses to communicate with the Leader must not interfere with the Leader's own downwash disturbance as sensed by the Follower during normal formation flight, and must not be mistaken by either vehicle's control loop for a disturbance requiring correction. This constraint has direct implications for the choice of gesture frequency and amplitude in Chapter 3, since these parameters must be selected to keep the return-channel signal spectrally and temporally distinguishable from the forward-channel dynamics. The fourth constraint is computational. The onboard controller has limited processing power, and every block implementing the return channel must be realizable within that budget: fixed-size data structures, no dynamic memory allocation, and algorithms simple enough to execute at the control loop's sampling rate without introducing additional latency. This constraint shapes not only the detection logic developed in Chapter 5 but also the gesture generator itself, which must be implementable as a lightweight function running alongside the existing flight control blocks.

Table 1.2: Design constraints for the return channel.

Constraint	Motivation
No additional hardware	The formation's entire rationale rests on implicit aerodynamic sensing; introducing dedicated sensors or actuators for the return channel would contradict that premise.
No explicit radio link	An explicit wireless channel, including short-range options, would remove the need for the implicit return channel altogether; it is excluded by the same design assumption that motivates the rest of the system.
Non-interference with the forward channel	The return-channel signal must remain distinguishable, in frequency and in time, from the Leader's downwash disturbance already exploited by the forward channel, and must not be misread by either vehicle's control loop as an unrelated disturbance.
Limited onboard computation	The return channel must run within the same computational budget as the rest of the flight controller: fixed-size structures, no dynamic allocation, and algorithms compatible with the control loop's sampling rate.

1.6 System-level requirements for bidirectional operation

The discussion so far has moved from the limitations of a one-way channel, through the functional requirements they impose, to the conceptual mechanism, role-swap and motion-based signaling, and the practical constraints within which that mechanism must operate. Taken together, these elements can now be collected into a small set of system-level requirements that the completed bidirectional system must satisfy as a whole, rather than requirement by requirement in isolation.

The symbol set used by the return channel must be mutually distinguishable under the noise conditions actually present in close-proximity flight, so that the Leader can tell one message apart from another with adequate reliability. The response times involved, from the moment the Follower initiates a gesture to the moment the Leader acts on the corresponding decision, must be compatible with the timescales of the formation's own flight dynamics, so that the return channel is useful rather than merely present. Finally, the forward and return channels must

coexist without interfering with one another, in the sense already introduced in Section 1.5: activity on one channel must not corrupt or be mistaken for activity on the other. These three requirements, distinguishability, timeliness, and non-interference, are stated here in general terms; verifying that a specific design satisfies them requires comparing candidate solutions against existing approaches to motion-based and implicit communication in multi-robot systems. This is the subject of the next chapter.

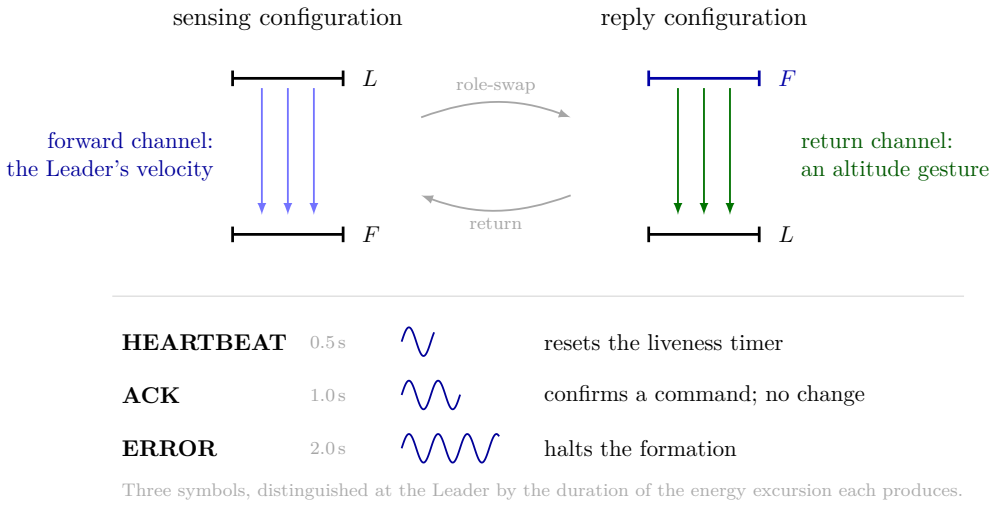


Figure 1.4: The complete bidirectional system. The formation alternates between two configurations: in the sensing configuration the Follower tracks the Leader through its wake, and in the reply configuration, entered through a role-swap, the vertical ordering is inverted for the length of one transmission so that the Follower's own downwash reaches the Leader. Three symbols travel in the return direction, distinguished by duration alone, and each maps to one of the three gaps of Section 1.2. This figure is the reference to which Chapters 3 to 6 add detail.

Chapter 2

State of the Art on Motion-Based Communication and Robust Decoding

2.1 Implicit communication in multi-robot and multi-agent systems

Cooperation among robots need not rely on an explicit exchange of messages. In multi-robot and multi-agent systems, coordination can instead emerge from the observation of physically grounded signals that are produced incidentally by the agents themselves, such as their motion, their relative position, or their effect on the surrounding environment. This mode of interaction, generally referred to as implicit communication, has been the subject of extensive study in swarm robotics and multi-agent coordination [2, 3], where the absence of a dedicated communication channel is treated less as a limitation and more as a design opportunity: the same physical quantities that already govern an agent's interaction with its neighbors can be repurposed to carry information, without the need for additional hardware or dedicated bandwidth.

The conceptual roots of this idea can be traced back to the observation of natural collectives. Flocks of birds and schools of fish are able to coordinate complex group behaviors while relying exclusively on locally sensed cues, such as the relative motion of neighboring individuals, rather than on any centralized or explicit form of communication [4, 5]. The present work does not attempt to reproduce the mechanisms underlying these biological systems in any quantitative sense; the

analogy is invoked here only to situate the broader research direction pursued in this thesis, in which a dedicated radio link between leader and follower is replaced by a channel that exploits an already-present physical interaction.

Framed in this way, implicit communication is not a peculiarity specific to the formation architecture considered here, but an instance of a broader and well-established research theme. The remainder of this chapter positions the aerodynamic and gesture-based channels developed in this and the companion thesis within that theme, before narrowing, in the following chapters, to the specific modeling and detection problems they raise.

2.2 Motion gestures as symbolic communication primitives

Within the broader landscape of implicit communication, a more specific and directly relevant line of work concerns the deliberate encoding of discrete symbols into patterns of motion, so that the trajectory or the kinematic profile of an agent becomes the carrier of a message rather than a mere by-product of its task execution. This idea reframes movement itself as a symbolic primitive: a gesture, in this sense, is a purposefully shaped motion segment designed to be recognized and decoded by an observer, much as a waveform is designed to be recognized by a receiver in conventional communication systems.

Two strands can be distinguished within this research direction. The first treats motion as a carrier of intent for a human observer, formalizing the notion that a trajectory can be shaped to make an agent's goal inferable rather than merely predictable [6, 7], and extending it from a single manipulator to multi-robot systems [8]. Applied to aerial platforms, the same idea has produced flight motions deliberately designed to convey direction, state, or intended action to a collocated user [9, 10, 11]. The second strand, to which this thesis belongs, differs in both the receiver and the encoding. The message is not addressed to a human observer who interprets it perceptually, but to another vehicle that recovers it through a signal-processing receiver, and the carrier is reduced from the full trajectory to the controlled variation of a single degree of freedom, leaving the remaining axes free to serve the formation task. This reduction turns a problem of perceptual clarity into one of detection, and it is in these terms that the channel is developed in the chapters that follow.

The gesture channel presented in the following chapters belongs to this second, narrower family: altitude is chosen as the sole modulated degree of freedom, and a sinusoidal variation of the follower's vertical position is used as the elementary waveform onto which the symbol alphabet is mapped. The specific design of this signal, together with the reasons behind the choice of a sinusoidal rather than a

stepped or pulsed profile, is discussed in detail in Chapter 3.

2.3 Signal detection in noisy dynamical systems

Once a symbol has been encoded as a controlled variation of a single motion axis, the companion problem is that of recovering it reliably at the receiving end. This is fundamentally a signal detection problem, but one posed in a setting that is less benign than a typical communication channel: the platform carrying the sensor is itself a dynamical system with its own oscillatory behavior, arising from attitude corrections, actuator noise, and the ordinary disturbances of flight, all of which act as a background against which the intentional gesture must be distinguished. The detection problem is therefore not simply one of separating a signal from additive noise, but of separating a deliberately shaped disturbance from the platform's own dynamics, which are neither stationary nor, in general, small compared to the signal of interest.

The theoretical foundations for this class of problems come from classical detection theory, where a signal known in advance is sought within a noisy observation and the trade-off between missed detections and false alarms is formalized through a decision threshold set according to the desired operating point [12]. Central to this framework is the notion of signal-to-noise ratio as the quantity governing detectability: a weak but well-characterized periodic signal can still be recovered reliably provided that the receiver structure is matched to its shape and that enough observation time, or equivalently enough periods of the signal, is available for the noise contribution to average out [13, 14].

What distinguishes the present setting from the classical formulation is the origin of the noise floor. In the standard treatment the disturbance is external to the receiver and can be characterized independently of it, whereas here it is produced by the platform carrying the sensor: attitude corrections, actuator activity and the response of the control loops themselves all contribute to the background against which the gesture must be detected, and none of them is stationary or independent of the vehicle's current state. This is precisely the situation faced by the return channel developed in this thesis, where the vertical disturbance produced by the gesture must be extracted from a leader platform that is simultaneously executing its own flight control. Rather than relying on a general characterization from the literature, the noise floor is measured directly for this configuration in Chapter 3, and the detection architecture that follows from it is developed in Chapters 3 and 5.

2.4 Matched filtering and energy-based detection

Among the techniques surveyed in the previous section, the matched filter occupies a central role in the design of the return channel and therefore deserves separate treatment. Originating in classical radar and communications theory, the matched filter is the linear filter that maximizes the output signal-to-noise ratio at the sampling instant when the shape of the signal to be detected is known in advance and the corrupting noise is white [15, 14]. Its optimality follows directly from the Cauchy-Schwarz inequality applied to the correlation between the received waveform and a template matched to the expected signal, and it is this property, rather than any heuristic argument, that motivates its adoption whenever the transmitted waveform is fixed and known, as is the case for the sinusoidal gesture used in this work.

A complication specific to the present application is that, while the shape and duration of the gesture are known by design, its phase at the receiver is not: the leader has no a priori knowledge of when, within a period of the sinusoid, a gesture will begin, since the follower initiates transmission asynchronously with respect to the leader's sampling. A matched filter built for a single fixed phase would therefore suffer significant performance loss whenever the received signal happens to be out of phase with the template, even though the underlying waveform is correctly identified in shape and frequency. This is the classical motivation for quadrature, or in-phase/quadrature (I/Q), detection: by correlating the received signal against two templates in phase quadrature with one another and combining the two outputs, typically through their sum of squares, it is possible to obtain a decision statistic that is insensitive to the unknown phase of the incoming signal [12, 16]. This property is precisely what makes the approach suitable for the follower-to-leader channel, and the architecture adopted for the leader-side decoder in Chapter 3 and Chapter 5 builds directly on this principle.

An alternative, non-coherent family of techniques dispenses with any notion of waveform template altogether and instead bases the detection decision purely on the energy of the received signal within an observation window, comparing it against a threshold set according to the estimated noise floor [13]. Energy detectors are attractive for their simplicity and for requiring no assumption on the signal's phase or even its exact shape, but this generality comes at the cost of reduced sensitivity compared to a matched filter when the signal shape is in fact known, since no coherent gain is extracted from the correlation with a template.

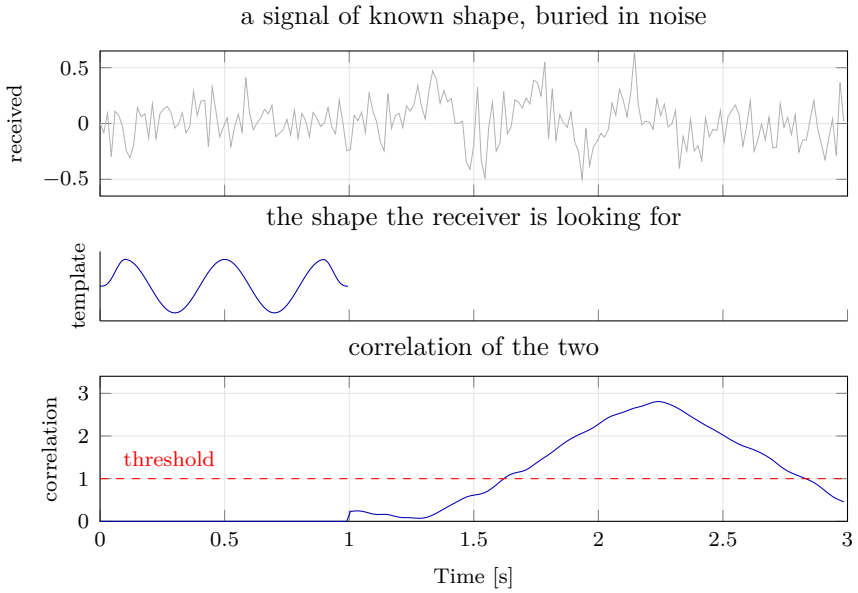


Figure 2.1: The principle of matched filtering. A waveform whose shape is known in advance can be recovered from noise that dominates it in amplitude by correlating the received signal against a copy of that shape: contributions that resemble the template accumulate coherently while the rest averages toward zero, so a peak appears at the instant of arrival even when nothing is visible in the received trace itself. This is the property the return channel relies on, and the reason its gesture is a waveform of fixed shape rather than an arbitrary displacement.

Table 2.1 summarizes the main trade-offs among the detection strategies considered.

Table 2.1: Comparison of candidate detection strategies for known periodic waveforms.

Technique	Advantages	Limitations
Fixed-phase matched filter	Optimal SNR gain for known, phase-aligned signals; simple implementation	Performance degrades sharply with unknown or drifting phase
I/Q (quadrature) matched filter	Retains matched-filter SNR gain while remaining insensitive to unknown phase	Requires two parallel correlations and a combining stage; slightly higher computational cost
Energy detection	No assumption on signal shape or phase; minimal design complexity	Lower sensitivity than matched filtering; more susceptible to non-white or bursty noise

This comparison anticipates the choice made in Section 2.6 and formalized in Chapter 3: since the gesture waveform is known precisely by design but its phase at the leader is not, the I/Q matched filter offers the best available compromise, combining the sensitivity of a matched approach with robustness to the asynchronous nature of the follower-to-leader transmission.

2.5 Communication protocols with acknowledgement and error handling

Beyond the physical layer of detection, the design of a reliable feedback channel also raises questions that belong to the classical theory of communication protocols: how does a receiver confirm that a transmission has been correctly received, how is the absence of an expected response interpreted, and what course of action should be taken when a message is lost or corrupted. These questions are addressed in the networking literature largely independently of the physical medium involved, through the well-established mechanisms of acknowledgement, timeout, and retransmission that underlie most reliable communication protocols [17, 18]. In an acknowledgement-based scheme, the receipt of a message is confirmed by a dedicated response; the absence of such a response within a bounded time window is treated not as silence but as evidence of failure, prompting the sender to retransmit or to take a predefined recovery action [19].

This family of mechanisms is particularly relevant to channels that are inherently unreliable or bandwidth-limited, where the cost of a lost message cannot be neglected

and where the feedback path itself may be as constrained as the forward one. Error-handling strategies developed for such settings typically distinguish between the detection of a failure, established by the absence of a timely acknowledgement or by an explicit error indication, and the recovery action taken once a failure has been detected, which may range from a simple retransmission to a coarser fallback behavior when the fault is judged severe enough that a normal retry would not suffice [18]. This distinction between an explicit positive confirmation and an explicit negative indication, rather than relying on silence alone to signal either outcome, is a recurring design principle in protocols intended to operate over channels where false negatives, that is, failures that go unnoticed, would be costly.

These concepts, developed here independently of any particular physical channel, provide the theoretical grounding for the protocol logic adopted at the leader side of the system studied in this thesis. The gesture-based return channel is built around a small alphabet of symbols that plays a role directly analogous to acknowledgement and error signaling in conventional protocols: a periodic liveness indication confirms that the link is active even in the absence of a specific event to report, a positive symbol confirms the correct receipt of a command, and a dedicated symbol signals that an anomaly has been detected on the follower side. The specific definition of this alphabet and the state machine that interprets it at the leader are presented in Chapter 4 and Chapter 5, respectively.

2.6 Positioning of this thesis with respect to existing work

The preceding sections have surveyed implicit communication at four progressively narrower levels: the general principle of coordinating agents through physically grounded signals rather than explicit messages, the use of motion itself as a carrier of intent, and the reduction of that carrier to a single controlled axis, the detection-theoretic tools required to recover a weak, phase-uncertain periodic signal from a platform with its own significant dynamics, and the protocol-level mechanisms of acknowledgement and error handling developed in the networking literature. Positioning the work presented in this thesis against this background makes it possible to state precisely which design choices are inherited from established results and which respond to constraints specific to the bidirectional formation scenario considered here.

The return channel developed in this thesis modulates a single degree of freedom, the follower's altitude, with a sinusoidal waveform, and relies on an I/Q matched filter at the leader to recover the transmitted symbol despite its unknown phase, feeding the result into a finite-state machine that distinguishes a periodic liveness symbol, a positive acknowledgement, and an error indication. Each of these choices

follows a specific line of reasoning already introduced in this chapter: the single-axis encoding keeps the channel independent from the forward wake-sensing mechanism and from any additional hardware, the I/Q matched filter is adopted because the gesture's shape is known but its phase is not, and the three-symbol alphabet reproduces, in a minimal form, the acknowledgement and error-signaling logic of conventional reliable protocols.

Table 2.2 summarizes how this combination of choices compares with the reference points discussed above, along four dimensions that recur throughout the chapter: the physical carrier of the information, whether additional hardware beyond the platform's existing sensors is required, whether the scheme provides explicit acknowledgement and error handling, and its expected robustness to the noise floor introduced by the platform's own dynamics.

Table 2.2: Positioning of the proposed gesture-based return channel with respect to the reference approaches discussed in this chapter.

Approach	**Physical carrier**	**Extra hardware**	**ACK / error handling**	**Robustness to platform noise**
Motion-based intent communication [6, 9]	Full trajectory or flight path	Not required	Not addressed; no receiver-side protocol	Designed for human perception, not for detection under noise
Generic dedicated radio link	Radio-frequency signal	Required (transceiver)	Native to the protocol stack	High, but at added weight/power cost
Baseline one-way aerodynamic sensing channel	Leader-to-follower wake / attitude disturbance	Not required	Not present (one-way only)	Moderate, no return path to confirm receipt
Proposed gesture-based return channel	Single-axis (altitude) sinusoidal modulation	Not required	Explicit, via HEARTBEAT / ACK / ERROR alphabet	High, via I/Q matched filtering matched to platform dynamics

What emerges from this comparison is that no single reference approach simultaneously satisfies all four requirements imposed by the scenario considered in this thesis: avoiding additional hardware, providing explicit and timely feedback on command execution, and remaining robust to a noise floor that is shaped by the leader's own flight dynamics rather than by an external, stationary source. The gesture-based return channel is designed specifically to occupy this otherwise unaddressed point in the design space, extending the one-way baseline system with a return path that closes the communication loop while preserving its defining property of requiring no dedicated radio hardware. The remainder of this thesis develops this channel in detail: Chapter 3 formalizes the gesture and channel model anticipated in Sections 2.2 and 2.3, Chapter 4 details the symbol alphabet and generator design motivated in Section 2.5, and Chapter 5 presents the I/Q matched-filter decoder and finite-state machine whose theoretical grounding was laid out in Sections 2.4 and 2.5.

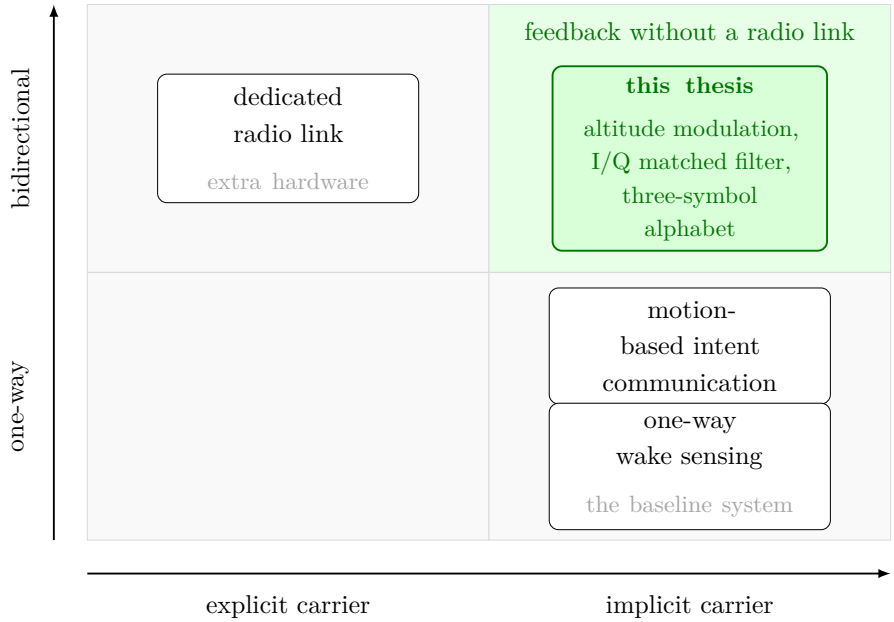


Figure 2.2: Positioning of the return channel with respect to the approaches surveyed in this chapter, along the two dimensions that separate them: whether information is carried by a dedicated signal or by the physics the vehicles already produce, and whether the resulting channel runs in one direction or two. The approaches that provide feedback do so over a dedicated link, and those that avoid a dedicated link provide no feedback. The upper-right quadrant is what this thesis occupies, and Table 2.2 states the same comparison in the four dimensions used throughout the chapter.

Chapter 3

Modeling the Follower-to-Leader Gesture Channel

3.1 Physical Principle of the Return Channel

The forward channel established in the preceding chapters relies on a geometric ordering that is, in this respect, asymmetric: in the nominal formation, the Leader flies above the Follower, and the rotor-induced flow trailing the Leader is directed downward and outward, arriving at the Follower as a measurable aerodynamic disturbance. Under this same ordering, however, no analogous exposure exists in the opposite direction. The wake generated by the Follower is carried away from the vehicle above it, so that the Leader remains, by construction, outside the region in which the Follower's own downwash could act as a detectable signal. A return channel built on the same physical carrier as the forward link is therefore not obtained simply by observing the existing formation geometry from the other side; it requires a deliberate, if temporary, departure from that geometry.

The mechanism adopted to create this exposure is a brief role-swap between the two vehicles. Starting from the nominal configuration, the Leader descends by a small, bounded amount while the Follower climbs to a position above it, inverting the vertical ordering of the pair for the duration of the manoeuvre. In this configuration, the wake generated by the Follower propagates downward directly onto the Leader, reproducing for the return direction the same physical exposure that the forward channel exploits in the nominal configuration. Once the manoeuvre concludes, the two vehicles resume their original vertical ordering, restoring the forward channel to its normal operating condition.

This choice is motivated directly by the design constraints established in Chapter 1: the return channel must be realized without additional hardware and without introducing a radio link between the vehicles. Reusing the aerodynamic wake as the carrier, and the onboard accelerometer as the sensing modality, satisfies both constraints at once, since neither element is specific to a single direction of communication. The same physical phenomenon and the same class of onboard measurement that allow the Follower to perceive the Leader are, under the inverted geometry produced by the role-swap, sufficient for the Leader to perceive the Follower. The return channel is in this sense not a separate communication system layered onto the formation, but a symmetric use of a carrier that was already available.

The quantity through which the Leader accesses this information is a vertical velocity estimate, denoted dz, derived from the onboard accelerometer during the role-swap window. The raw form of this signal is not, by itself, a clean representation of the Follower's gesture: its preprocessing, and the noise sources that shape its usable bandwidth, are treated in detail in Section 3.4, where the filtering stage that isolates the gesture-relevant component of dz is introduced. For the purposes of this section, it is sufficient to note that dz, rather than any externally sensed or visually derived quantity, is the sole channel through which the Leader becomes aware of the Follower's signalling.

It is worth emphasizing what the role-swap manoeuvre does not affect. The horizontal dynamics of the formation, and in particular the relative positioning and pursuit behaviour established by the forward sensing pipeline, are not altered by the vertical inversion: the manoeuvre is confined to a short, bounded window and does not modify the pair's lateral or longitudinal separation. Outside this window, the formation geometry and its associated sensing behaviour are unaffected by the existence of the return channel. This containment is itself a design requirement, since any broader disturbance to the nominal formation would compromise the operation of the forward channel that the return channel is built to complement, not to replace.

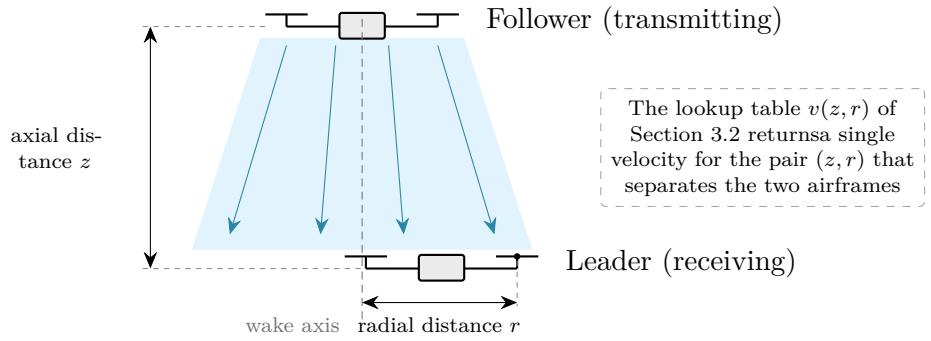


Figure 3.1: Geometry during the role-swap manoeuvre. The Follower climbs above the Leader, so that its wake propagates downward onto the vehicle below, reproducing for the return direction the same exposure the forward channel exploits in the nominal configuration. The two coordinates that parameterize the channel model are the axial distance z between the two rotor planes and the radial distance r of a given point from the wake axis.

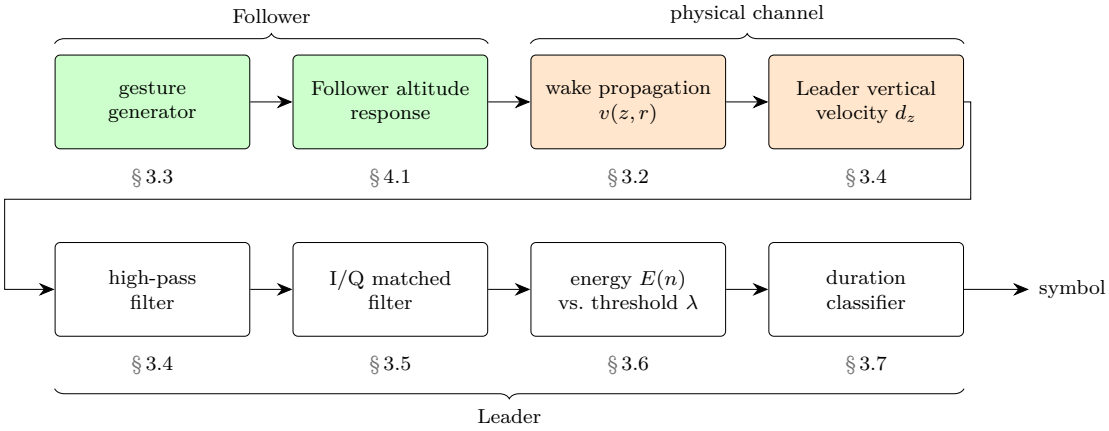


Figure 3.2: Overview of the complete follower-to-leader chain, with the section in which each stage is developed. The gesture is synthesized as an altitude reference offset on the Follower, tracked by its altitude loop, propagated to the Leader as an aerodynamic disturbance, and recovered from the Leader's own vertical velocity estimate through filtering, correlation and duration classification. Each stage is treated separately in the remainder of this chapter.

3.2 Gesture Signal Model and Design Variables

Closed-form models of rotor-induced downwash, including the turbulent-jet formulation adopted by Bauersfeld et al. [20], describe the flow field through a small

number of similarity parameters: a centreline decay law and a self-similar radial profile governed by a half-velocity radius. These parameters are, however, geometry- and scale-dependent, and the constants reported in the literature are calibrated for rotor sizes, thrust levels and measurement ranges that do not necessarily match the close-proximity, role-swap-specific geometry considered here. Rather than adopting such constants directly, this section characterizes the wake produced during the role-swap manoeuvre through a dedicated anemometric measurement campaign, using the same functional forms as a template but fitting their parameters to data collected for this specific configuration.

The measurements were taken with the vehicle held in stationary hover at a fixed altitude, so that the wake sampled at each station is the steady flow produced at hover thrust rather than a transient associated with a change in altitude. The anemometer was positioned at the station to be measured and the vehicle, not the instrument, defined the reference: axial distance is measured downward from the rotor plane, and radial offset is measured from the vertical axis through the centre of the airframe. Each station was sampled between three and five times, and the value entered into the fit is the mean of those repetitions, so that the scatter of the individual readings, which is appreciable in a turbulent jet, is reduced before the functional forms of this section are fitted rather than after. The repetitions also provide the measure of repeatability referred to below: it is the spread among readings at the same station, not a residual of the fit.

Measurements were taken with an Extech 407123 hot-wire thermo-anemometer, whose sensing element sits at the tip of a telescoping probe connected to the instrument body by a cable. The arrangement matters for a measurement of this kind: only the probe enters the jet, while the instrument and the operator remain outside it, so the flow being sampled is not the flow the measurement itself has disturbed. The instrument resolves velocities down to roughly $0.2\,\mathrm{m/s}$, two orders of magnitude below the centreline values recorded here, so the lower end of its range is never approached.

The campaign combines two complementary measurements. An axial profile of the wake centreline velocity was recorded at eight axial distances between $z = 10$ and $z = 50\,\mathrm{cm}$, with three repeated measurements at each station to characterize repeatability. A set of radial profiles was then recorded at five of these axial stations ($z = 20, 25, 30, 35, 40\,\mathrm{cm}$), each sampled at two azimuthal orientations (0° and 45°) to assess whether the wake could be treated as axisymmetric. Across the measured stations, the velocity recorded at matching radii along the two orientations differs only marginally, supporting the assumption of axisymmetry used in the model below and simplifying the lookup table to a function of axial distance and radius alone, without an azimuthal dependence.

The centreline velocity is fitted to the same decay law used in the reference

formulation,

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

giving $C = 444\,\mathrm{cm} \cdot \mathrm{m/s}$ and $z_0 = -34\,\mathrm{cm}$. Both parameters are required: constraining z_0 to zero, which would reduce the law to the $1/z$ decay characteristic of a fully developed turbulent jet, fails to describe the measurements at all. The reason is visible in the data themselves. Across the measured range the centreline velocity falls by a factor of 2.1 while the axial distance increases by a factor of 5, whereas a $1/z$ law would require the two to fall together. The measurements therefore do not lie in the far field of the jet but in the region where the flow is still developing, and the large negative z_0 is what allows the two-parameter form to reproduce this slower decay. It should accordingly be read as an empirical fit parameter rather than as the physical origin of an equivalent far-field jet, whose location it does not identify.

The radial profile is fitted to the self-similar shape

$$v(r, z) = U_C(z) \left(1 + k \left(\frac{r}{r_{1/2}}\right)^2\right)^{-2}, \tag{3.2}$$

with $k = \sqrt{2} - 1$ and a half-velocity radius $r_{1/2} = 9.6\,\mathrm{cm}$ held constant across all axial stations, the on-axis value at each station being taken from the axial fit above. Figure 3.4 shows the resulting profiles against the measured data. Holding $r_{1/2}$ constant rather than fitting it independently at each station is a deliberate simplification: station-by-station fits over the measured range do not exhibit the monotonic growth with z that a fully developed jet would display, which is consistent with the near-field character of the measurements noted above, and the spread between them is comparable to the scatter of the measurements themselves. A single value therefore describes the data no worse than five separate ones, while keeping the lookup table of this section a function of two variables with a single shape parameter.

Both functional forms are the same used by Bauersfeld et al. to characterize rotor-induced downwash, and the measured trends are broadly consistent with that reference: the centreline velocity decays hyperbolically with axial distance, and the radial profile retains the same self-similar, bell-shaped form at every axial station tested.

The fitted constants obtained here can be compared directly with those reported in the original work once the two are expressed in the same units. Bauersfeld et al. normalize both coordinates by the motor-to-motor distance l and the velocity by the induced velocity at hover U_H, and report $Bd = 10.11$, $S = 0.07668$ and $s_0 = -5.817$ for the unified fit. For the platform used in this campaign, a Parrot Mambo of mass $m = 63\,\mathrm{g}$ with propellers of diameter 65 mm and a motor-to-motor distance of

$l \approx 10\,\mathrm{cm}$, momentum theory gives $U_\mathrm{H} = 4.3\,\mathrm{m\,s^{-1}}$, and the published parameters translate into $C_\mathrm{ref} = U_\mathrm{H}\,B d\,l \approx 439\,\mathrm{cm\,m\,s^{-1}}$ and $z_{0,\mathrm{ref}} = -5.817\,l \approx -58\,\mathrm{cm}$.

Two points follow from this comparison. First, the negative development length recovered here is not an artefact of the present campaign: the reference fit is itself obtained with a large negative s_0, and the value found here, $z_0/l \approx -3.0$, is of the same sign and the same order as the published -5.82. The interpretation given above, that z_0 is an empirical fit parameter rather than the location of a physical virtual origin, is therefore consistent with the reference rather than a departure from it. Second, the constant half-velocity radius adopted here is close to what the reference model itself predicts over the measured range: evaluating $r_{1/2}(z) = S(z - z_{0,\mathrm{ref}})$ between $z = 20\,\mathrm{cm}$ and $z = 40\,\mathrm{cm}$ gives a growth from $6.0\,\mathrm{cm}$ to $7.5\,\mathrm{cm}$, a variation comparable to the scatter of the measurements themselves. Holding $r_{1/2}$ constant is thus not merely a convenient simplification but an approximation the reference model supports at these distances.

The remaining discrepancy is in the absolute scale of the radial spread, the value fitted here being wider than predicted, and this is consistent with the near-field character already established. Bauersfeld et al. report that the flows of the four individual propellers merge into a single turbulent jet only about 2.5 motor distances below the vehicle, which for this platform corresponds to $z \approx 25\,\mathrm{cm}$: a substantial part of the measured range therefore lies, by the reference's own criterion, upstream of the region in which the far-field model applies. A direct check confirms this. At $z = 20\,\mathrm{cm}$ the fitted centreline velocity is $8.2\,\mathrm{m\,s^{-1}}$, whereas $2U_\mathrm{H} = 8.7\,\mathrm{m\,s^{-1}}$ is the velocity of the fully contracted slipstream predicted by momentum theory. The measurements are therefore at the velocity scale of the developing slipstream rather than at that of the decayed far-field jet, which is precisely the regime in which the far-field decay law is not expected to hold and in which the fit parameters must be read as empirical descriptors of the measured data.

The model structure therefore transfers across configurations, while its parameters must be recalibrated for the specific geometry of the role-swap manoeuvre rather than taken from the literature: this is precisely the approach adopted at the beginning of this section, and the agreement found above confirms that it was the right one rather than a merely cautious one.

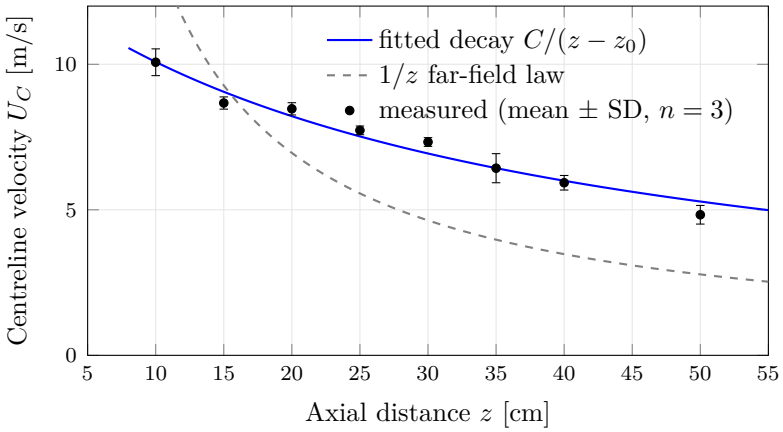


Figure 3.3: Axial decay of the Follower downwash. The fitted law $C/(z - z_0)$ is shown against the anemometric data (mean ± SD, $n = 3$), with the $1/z$ far-field law overlaid for comparison: the measured velocity falls by a factor of 2.1 while the axial distance increases by a factor of 5, which the $1/z$ law cannot reproduce.

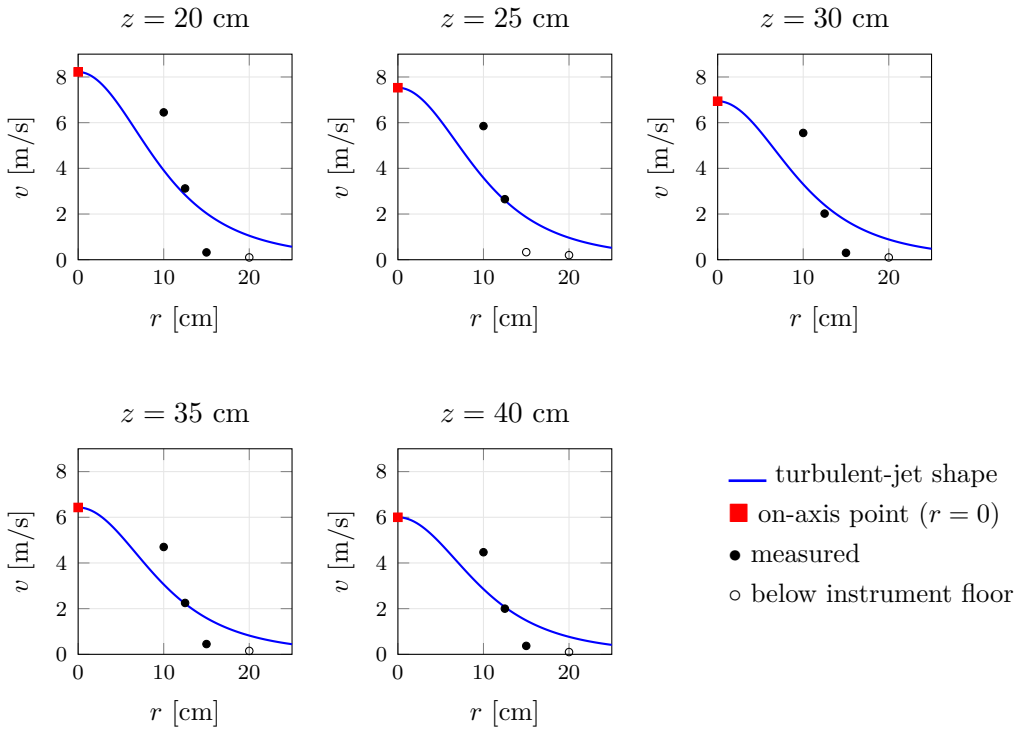


Figure 3.4: Radial profiles of the Follower downwash at the five axial stations. The fitted self-similar shape uses a half-velocity radius $r_{1/2} = 9.6\,\text{cm}$ held constant across all stations, the on-axis value at each station being taken from the axial fit of Figure 3.3. Open markers denote readings below the instrument's resolution floor.

3.3 Frequency, Amplitude and Duration Constraints

Having established in the previous section how a given altitude displacement propagates to the Leader, this section defines the gesture itself: the altitude reference offset that the Follower actually commands during the role-swap window. The gesture is formalized as a sinusoidal modulation of the Follower's altitude reference, shaped at its onset and offset by a Tukey window rather than switched on and off abruptly. This choice of waveform, rather than a step or ramp offset, keeps the commanded trajectory smooth and differentiable, which is a prerequisite for it to be tracked accurately by the altitude controller without exciting transients that would otherwise appear, at the Leader's side, indistinguishable from the gesture itself. Having established in the previous section how a given altitude displacement propagates to the Leader, this section defines the gesture itself: the

altitude reference offset that the Follower commands during the role-swap window. It is not a force or a torque, and it is fully specified by three quantities,

$$a(t) = A\, w(t)\, \sin(2\pi f_0 t), \qquad 0 \le t \le T_s, \tag{3.3}$$

where A is the commanded amplitude, f_0 the carrier frequency, T_s the symbol duration, and $w(t)$ a Tukey window that tapers the onset and offset of the interval and leaves the remainder at unity. Setting $w(t) \equiv 1$ recovers the rectangular case, whose spectral consequences are the subject of the discussion that follows. Two parameters govern the sinusoidal component: an amplitude $A = 0.12\,\mathrm{m}$ and a frequency $f_0 = 2.5\,\mathrm{Hz}$. Both values are constrained from opposite directions. The amplitude must be large enough to produce, once propagated through the channel characterized in Section 3.2, a disturbance at the Leader that remains well above the noise floor discussed in Section 3.4; at the same time, it must stay within the vertical excursion that can be regarded as safe during close-proximity flight, since the role-swap manoeuvre already reduces the nominal vertical separation between the two vehicles. Similarly, the frequency must be high enough to keep the gesture short and to place it outside the bandwidth of the vehicles' normal formation-keeping dynamics, so that it is not confused with ordinary altitude corrections; but it cannot exceed the closed-loop bandwidth of the Follower's altitude controller, beyond which the commanded oscillation would be attenuated before it is even executed.

Rather than being fixed independently, A and f_0 are calibrated jointly through the channel model established in Section 3.2. For a candidate amplitude, the lookup table $v(z, r)$ predicts the disturbance velocity expected at the Leader's position for the role-swap separation adopted in this work; this predicted value is then checked against the noise floor characterized in Section 3.4, providing an early estimate of detection margin well before the full detectability analysis of Section 3.5. The frequency is constrained in parallel by the same reasoning applied in the spectral domain: it must sit above the bandwidth of ordinary formation-keeping motion and below that of the altitude controller's tracking response, so that the gesture is neither mistaken for routine manoeuvring nor filtered out before execution. Both bounds on f_0 can be stated numerically. The lower one is the bandwidth of ordinary formation-keeping motion, which for this platform is of the order of half a hertz: the Leader's own closed-loop altitude response, identified in Section 7.1, sits at $0.5\,\mathrm{Hz}$, and the relative motion the pursuit controller produces is slower still. The upper one is the closed-loop bandwidth of the Follower's altitude loop, identified at $3.49\,\mathrm{Hz}$ with a damping ratio of 0.9. The chosen $f_0 = 2.5\,\mathrm{Hz}$ therefore sits five times above the motion it must not be confused with, and at 0.72 of the bandwidth beyond which it could not be executed at all. That it sits inside rather than well inside the second bound is a deliberate consequence of the first: pushing f_0 lower to gain tracking fidelity would move it toward the band occupied by ordinary

manoeuvring. The price is paid in amplitude rather than in detectability, and it is quantified in Section 7.2, where the second-order response at this frequency ratio attenuates the commanded oscillation to 0.73 of its amplitude and delays it by 69°.

The amplitude admits a comparable early estimate. Because the gesture displaces the Follower along the axis on which the wake decays, the separation seen by the Leader oscillates between $z - A$ and $z + A$, and the lookup table converts that oscillation directly into a modulation of the disturbance velocity. At a role-swap separation of 20 cm, the centreline law fitted in Section 3.2 returns 8.22 m/s at the nominal position and swings between 10.57 and 6.73 m/s over a cycle of $A = 0.12$ m: a modulation amplitude of 1.92 m/s, or 23 % of the mean. It is this modulation, and not the mean flow, that the return channel transmits, since the high-pass stage of Section 3.4 removes the standing component. The estimate is an upper bound in the sense that it assumes the commanded amplitude is executed; with the excursion actually achieved the modulation falls to 1.36 m/s, or 17 %.

The values adopted here, $A = 0.12$ m and $f_0 = 2.5$ Hz, are the result of this calibration process rather than an arbitrary choice, and are revisited quantitatively once the detection performance is analyzed later in this chapter.

The use of a Tukey window, rather than a rectangular one, follows directly from the frequency-domain consequence of switching the sinusoid on and off abruptly [21, 22]: a rectangular envelope introduces spectral sidelobes around f_0, widening the effective bandwidth that the gesture occupies and, as discussed in Section 3.5, reducing the selectivity of the matched filter used to detect it. Tapering the onset and offset of the sinusoid instead concentrates the gesture's energy closer to f_0, which is directly reflected in a sharper correlation peak at the detector and, correspondingly, a lower probability of confusing the gesture with broadband disturbances unrelated to it.

Amplitude, frequency and windowing together define the shape of a single gesture, independently of which symbol it is being used to represent; the distinction between symbols is introduced separately, through the duration for which the sinusoid is sustained, treated here only to the extent that duration is itself a design variable subject to the same kind of trade-off as amplitude and frequency: longer gestures are easier to detect reliably but slower to transmit and more likely to overlap with other formation events. The specific durations assigned to each symbol, and the alphabet they encode, are defined in Section 4.3, once the detection performance analyzed later in this chapter provides a basis for choosing them.

This definition is implemented directly by the `gesture_generator` block, which produces the tapered sinusoidal altitude offset described above from the parameters A, f_0 and the requested symbol duration.

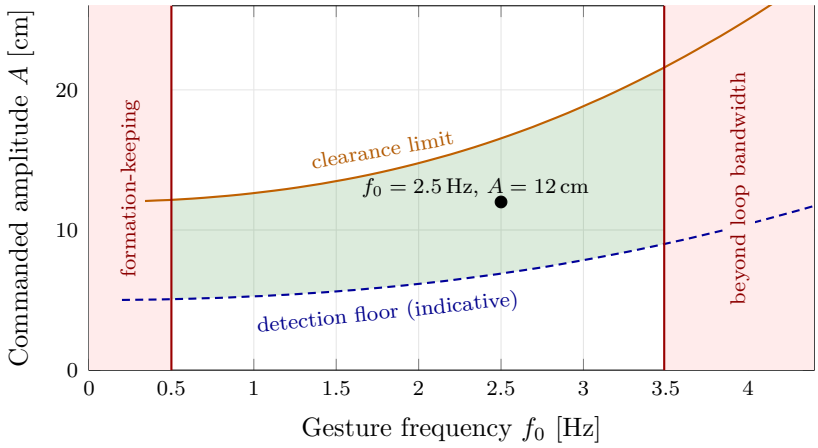


Figure 3.5: The admissible region for the two gesture parameters. The frequency is bounded below by the band occupied by ordinary formation-keeping motion and above by the closed-loop bandwidth of the Follower's altitude loop, both identified in Section 7.1. The amplitude is bounded above by the vertical clearance that must remain during a role-swap and below by the disturbance the Leader must be able to detect; both bounds rise with frequency, because the same loop that limits f_0 also attenuates what it executes, so a given excursion costs more commanded amplitude the closer f_0 sits to the bandwidth. The chosen pair lies inside the region with margin on every side. The lower boundary is indicative, since the noise floor on dz is an estimate rather than a quantity identified from a hover log.

3.4 Noise Sources in Leader-Side Gesture Observation

The physical channel characterized in Section 3.2 predicts the disturbance velocity that a gesture of given amplitude and frequency should produce at the Leader's position. What the Leader actually receives, however, is not this disturbance in isolation: the quantity available onboard is a vertical velocity estimate, dz, derived from the accelerometer during the role-swap window introduced in Section 3.1, and this estimate carries, superimposed on the gesture, everything else that contributes to the vehicle's vertical acceleration over the same interval.

Before any detection can take place, this raw estimate is passed through a first-order high-pass filter with a cutoff near 0.3 Hz. The stage is realized as a single recursion on the incoming estimate,

$$y[n] = \alpha \Big(y[n-1] + x[n] - x[n-1] \Big), \tag{3.4}$$

with $\alpha = 0.99$ at $F_s = 200$ Hz, which places the cutoff at the value quoted above

and gives the transient a time constant of $\tau = 1/(1-\alpha) \approx 99$ samples. Both figures matter later: the first sets what the stage removes, the second sets how long its output must be ignored after start-up. The role of this stage is to remove the slowly varying component of dz: sensor bias, thermal drift, and the residual low-frequency offset introduced by the role-swap manoeuvre itself as the Leader settles into its lowered altitude. None of these components carry gesture information, since the gesture is designed, as established in Section 3.3, to occupy a narrow band around $f_0 = 2.5\,\text{Hz}$, well above the filter's cutoff; the filter is accordingly transparent to the gesture itself while suppressing the part of the signal that would otherwise dominate it in amplitude.

What remains after this stage is not, however, a clean sinusoid. Three sources of residual disturbance are relevant to the detection problem addressed in the remainder of this chapter. The first is sensor noise intrinsic to the inertial measurement unit, present regardless of the manoeuvre being performed. The second is aerodynamic disturbance uncorrelated with the gesture itself: residual turbulence in the wake, or transient gusts, that perturb the Leader's vertical acceleration without being related to the Follower's commanded oscillation. The third is interference from the formation's own ordinary motion: velocity or heading adjustments that are part of normal formation-keeping and that, unlike the uncorrelated aerodynamic disturbances above, can be brief and structured enough to resemble part of a gesture if not accounted for.

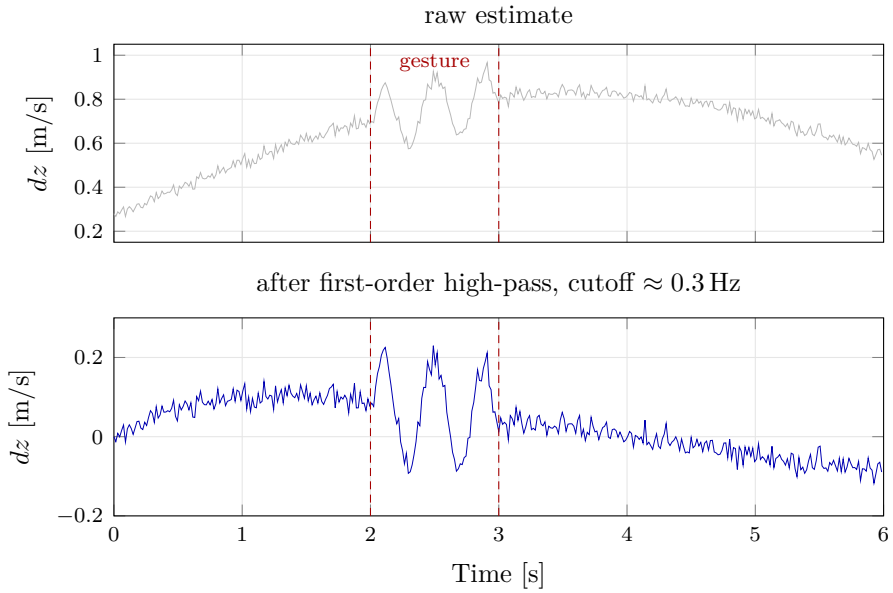


Figure 3.6: Illustrative effect of the high-pass stage on the Leader's vertical velocity estimate. Above, a simulated dz in which a one-second gesture at f_0 is superimposed on a sensor bias and a slow drift an order of magnitude larger than itself: the gesture is present but not visible, because the quantity is dominated by components that carry no information. Below, the same signal after the filter, which removes the bias and the drift while passing f_0 essentially unchanged, leaving the gesture standing clear of the residual noise. The traces are generated from a standalone simulation for illustration; filtered signals from the full model are reported in Chapter 7.

Taken together, these three sources define the noise floor against which the gesture must be detected. Their combined effect is not addressed by filtering alone: Section 3.5 introduces the matched filter used to extract the gesture from what remains after this preprocessing stage, and Section 3.6 discusses how the residual noise characterized here translates into false alarm and missed detection behaviour at the decision stage.

3.5 Detectability of Sinusoidal Altitude Gestures

Section 3.4 established that the signal available to the Leader, after high-pass filtering, is not a clean sinusoid at f_0 but a weak periodic component embedded in sensor noise, uncorrelated aerodynamic disturbance, and occasional interference from ordinary formation motion. Extracting a reliable indication of the gesture's presence from this signal requires more than searching for energy at f_0: the detector

must also be robust to a relative phase between the transmitted gesture and the observation window that is not controlled by either vehicle, since the Leader has no explicit synchronization signal telling it when a gesture begins.

This requirement motivates the use of an in-phase/quadrature (I/Q) matched filter rather than a single real-valued correlator. Correlating the filtered signal with a single reference sinusoid at f_0 produces an output that depends not only on the presence of the gesture but also on the phase offset between the gesture and the reference: for an unfavourable offset, the correlation can be small even when the gesture is present at full strength, making detection unreliable. The I/Q structure avoids this by correlating the signal separately against two reference components in quadrature, an in-phase term $I(n)$ and a quadrature term $Q(n)$, and combining them into an energy measure $E(n) = I(n)^2 + Q(n)^2$. Because $I(n)$ and $Q(n)$ respond to orthogonal phases of the same frequency, $E(n)$ is, to first order, independent of the gesture's phase relative to the correlation window: whatever phase the gesture arrives with, its energy is distributed between $I(n)$ and $Q(n)$ in a way that preserves the total. This is precisely the property that justifies the choice of a Tukey-tapered sinusoid over a rectangular one in Section 3.3: concentrating the gesture's spectral energy around f_0 sharpens the response of $E(n)$ to a genuine gesture and reduces its sensitivity to disturbances at nearby frequencies that a wider spectral footprint would have let through.

The energy $E(n)$ is compared against a fixed detection threshold to decide whether a gesture is present at time n. Values of $E(n)$ that remain below the threshold are interpreted as absence of a gesture, regardless of their cause; values that cross it are interpreted as the onset of a symbol, with the duration for which $E(n)$ remains above threshold used downstream, as discussed in Section 3.7, to classify which symbol was transmitted. The threshold itself is not fixed arbitrarily: its placement relative to the noise floor characterized in Section 3.4 and to the expected signal strength predicted by the channel model of Section 3.2 directly determines how often the detector reports a gesture that was not sent, or fails to report one that was, the two failure modes examined next, in Section 3.6.

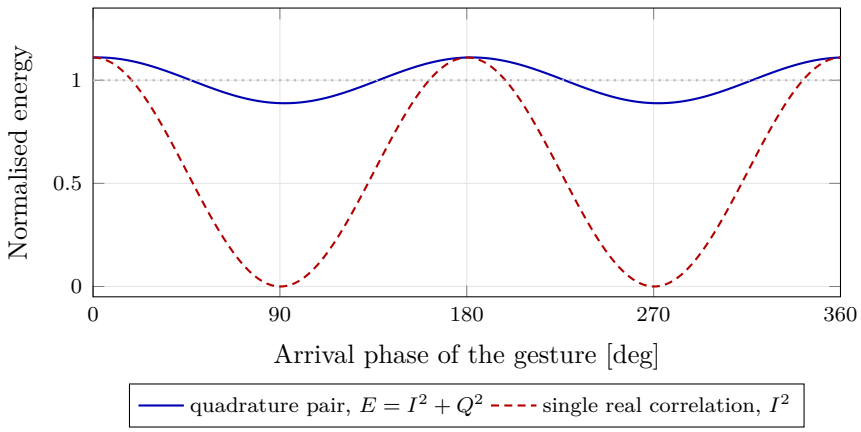


Figure 3.7: Why detection uses a quadrature pair. The same gesture is presented to the correlator at every arrival phase, a quantity the Leader cannot know in advance because it depends on when the Follower happened to begin transmitting. A single real correlation returns an energy that collapses entirely at the phases where the gesture is orthogonal to the reference, so a detector built on it would miss a transmission that was present and undistorted. The quadrature pair holds the energy within roughly $\pm 11\,\%$ of its mean. The residual ripple is not a property of the method but of the window length: the correlation spans two and a half cycles of f_0 rather than an integer number, and an integer-cycle window would remove it. That margin is small against the detection margin of Section 3.5, whereas the collapse of the real correlation is not.

3.6 False Alarm and Missed Detection Mechanisms

The detection rule introduced in Section 3.5 reduces the problem to a single comparison: the energy $E(n)$ against a fixed threshold. This simplicity, however, does not remove the two ways in which the comparison can go wrong, and both are consequences of the same noise sources characterized in Section 3.4 acting on either side of the threshold.

A false alarm occurs when $E(n)$ crosses the threshold in the absence of any transmitted gesture. Because $E(n)$ is computed continuously from a signal that always contains some combination of sensor noise, uncorrelated aerodynamic disturbance and, occasionally, structured interference from ordinary formation motion, its value fluctuates even when the Follower is not signalling; a sufficiently large fluctuation, however improbable, is indistinguishable at the level of $E(n)$ alone from the onset of a genuine symbol. The rate at which this happens is governed entirely by where the threshold is placed relative to the statistics of this residual energy: a

threshold set close to the typical noise-floor energy will be crossed frequently by chance, while one set further above it is crossed only by disturbances large enough to resemble a gesture in both frequency content and duration.

A missed detection occurs in the opposite circumstance: a gesture is transmitted, but $E(n)$ never crosses the threshold, or crosses it for too short an interval to be registered. This can happen for two distinct reasons. The first is that the gesture arrives at the Leader weaker than intended, because the role-swap separation realized in practice departs from the nominal one on which the amplitude A was sized in Section 3.3, so that the disturbance predicted by the channel model of Section 3.2 overstates what is actually received. The second is that the gesture, even at its intended strength, arrives superimposed on a transient disturbance large enough to distort the correlation without itself resembling a gesture, so that the resulting $E(n)$ trajectory crosses the threshold late, briefly, or not at all.

These two failure modes cannot be reduced independently: lowering the threshold to catch weaker gestures increases the rate of false alarms, and raising it to suppress false alarms increases the rate of missed detections. In principle, the threshold could be set from the statistics of the noise-floor energy alone, for instance by placing it at a high percentile of the energy observed in the absence of a gesture. This approach was not viable here: consecutive matched-filter outputs share all but one sample of their correlation window, so successive values of $E(n)$ are strongly correlated rather than independent, and a percentile computed over such a sequence does not carry the statistical meaning it would for independent samples. The threshold was therefore calibrated empirically instead, directly against the noise floor and the expected gesture energy rather than against an assumed noise distribution.

The threshold adopted, $\lambda = 0.030$, sits between two reference levels observed in simulation: a noise-floor energy of approximately 0.013, representing the largest fluctuation of $E(n)$ produced by the combined noise sources of Section 3.4 in the absence of any gesture, and a typical gesture-peak energy of approximately 0.25, representing the response of the matched filter to a gesture of the amplitude and frequency fixed in Section 3.3 and propagated through the channel model of Section 3.2. Expressed as ratios, the threshold sits at $\lambda / E_{\text{noise,max}} \approx 2.3$ (3.6 dB) above the noise floor, and at $E_{\text{gesture}} / \lambda \approx 8.3$ (9.2 dB) below the typical gesture peak. The first ratio is what keeps the false alarm mechanism described above rare, since the noise floor would have to exceed its observed maximum by more than a factor of two to cross λ; the second is what keeps the missed detection mechanism rare, since a gesture would have to lose most of its expected energy, whether through an unfavourable role-swap separation or through a large uncorrelated disturbance, before it stopped being distinguishable from noise. The resulting detector performance, evaluated through Monte Carlo simulation and confirmed in closed-loop validation, is presented in Chapter 7.

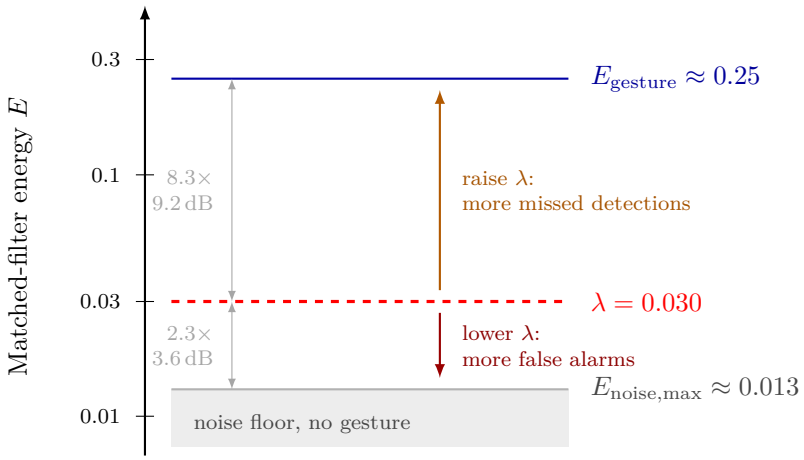


Figure 3.8: Placement of the detection threshold between the two failure mechanisms. The threshold cannot be derived from the statistics of the noise-floor energy alone, because consecutive matched-filter outputs share all but one sample of their correlation window and are therefore strongly correlated; it was calibrated instead against the two levels shown, both observed in simulation. The margins above and below are deliberately asymmetric: the smaller one guards against reporting a gesture that was never sent, the larger against failing to report one that was, and the two cannot be widened at the same time.

3.7 Timing Constraints for Reliable Symbolic Communication

Sections 3.5 and 3.6 established how a single gesture is detected and separated from noise. Turning that detection into a reliable symbolic message requires an additional set of timing constraints, governing not the shape of one gesture but the relationship between successive ones and the interval in which detection is allowed to occur at all.

Detection is disabled for the first 400 samples (2.0 s) after the matched filter starts running. Two distinct transients must expire before the detector's output carries meaning, and the interval is sized to the longer of the two. The first is the filling of the circular correlation buffer of Section 3.5, which requires $N = 200$ samples, since a correlation computed over a partially populated window bears no relationship to the energy the same window would produce once full. The second, and the one that actually determines the value, is the settling of the high-pass filter of Section 3.4. That filter has a time constant of approximately 99 samples, so its response to the initial sensor bias decays over several hundred samples; until it has, the filtered signal carries a slowly decaying component large enough to drive $E(n)$

above λ on its own. At 400 samples the filter has settled for four time constants, leaving under two per cent of that transient, whereas the 200 samples needed by the buffer alone would leave enough of it to produce a spurious detection immediately after the window closed.

Once this blanking window has elapsed, two further constraints apply to every detected peak. A minimum duration of 30 samples (0.15 s) is required for an energy excursion above λ to be accepted as a gesture at all, rejecting brief crossings caused by isolated noise bursts rather than by a sustained periodic component. A debounce interval of 100 samples (0.5 s) is then enforced between the end of one accepted detection and the start of the next, preventing a single physical gesture from being fragmented into multiple detections by a momentary dip of $E(n)$ below threshold near the gesture's edges.

Within these bounds, the classifier distinguishes between symbols using the duration for which $E(n)$ remains above λ, rather than the nominal duration commanded by the gesture generator. The two do not coincide: the matched filter correlates the incoming signal against a template of $N = 200$ samples, and as this template slides onto and off a gesture, the resulting energy excursion is measurably wider than the gesture's nominal duration. The extent of this widening is not derived analytically here but characterized empirically. At the nominal role-swap geometry the observed peak width is approximately 170 samples for HEARTBEAT, 280 for ACK and 475 for ERROR, against nominal durations of 100, 200 and 400. These widths are not constants: they narrow as the separation between the two vehicles increases and the received disturbance weakens. The values quoted above are those observed at the nominal geometry, while the ranges collected in Table 3.1 span the full operating region characterized in Section 7.5. The classification thresholds, 235 and 400 samples, are placed between the widths of adjacent symbols rather than at either symbol's nominal duration. Writing W for the width in samples of an excursion that has just closed, the assignment is

$$\text{symbol}(W) = \begin{cases} \textsc{heartbeat}, & W \leq 235, \\ \textsc{ack}, & 235 < W \leq 400, \\ \textsc{error}, & W \geq 400. \end{cases} \tag{3.5}$$

This placement is what resolves an apparent inconsistency in the nominal design values. The ACK symbol's nominal duration, 200 samples, sits numerically below the 235-sample boundary that separates HEARTBEAT from ACK; read naively, this would suggest that an ACK gesture risks being classified as a HEARTBEAT. In practice this does not occur, because classification acts on the observed peak width rather than the nominal duration, and that width exceeds the boundary by some 40 samples at the nominal geometry.

A second, related coincidence occurs at the opposite boundary. The ERROR

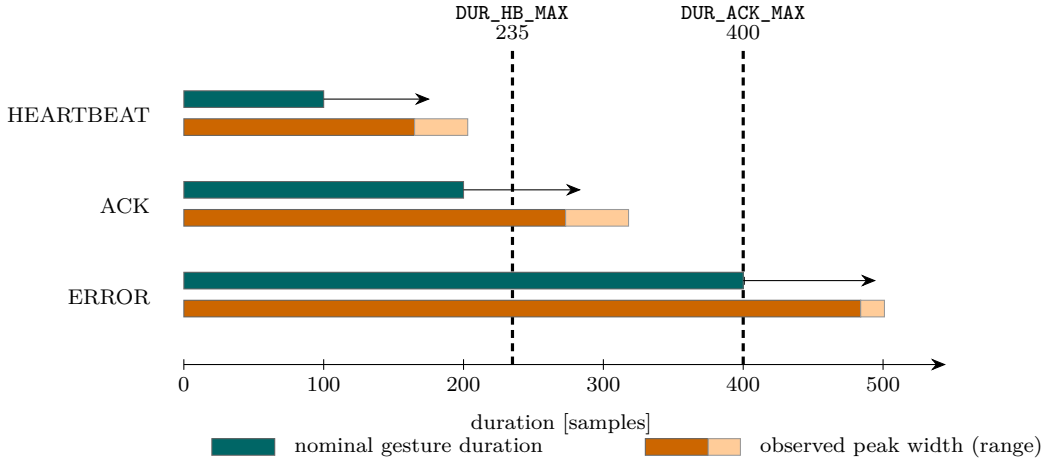


Figure 3.9: Nominal gesture duration compared with the width of the energy excursion the same gesture produces at the detector. The correlation window slides onto a gesture before it ends and off it only afterwards, so every observed width exceeds the duration that generated it. The classification boundaries are placed between the observed widths of adjacent symbols, which is what resolves the two apparent coincidences in the nominal values: the ACK's nominal 200 samples sit below the HEARTBEAT/ACK boundary, and the ERROR's nominal 400 coincide with the ACK/ERROR boundary exactly, yet neither symbol is misclassified, because classification acts on the observed width rather than on the nominal duration.

symbol has a nominal duration of exactly 400 samples, numerically identical to DUR_ACK_MAX, the same value used as the ACK/ERROR boundary. Here again the resolution lies in the observed rather than the nominal duration, the measured width exceeding the boundary by roughly 70 samples. The nominal equality between the ERROR duration and the classification threshold is therefore a property of the design values chosen in Section 3.3, not of the values the classifier actually observes.

Taken together, the blanking window, the minimum-duration filter, the debounce interval, and the classification thresholds define the complete timing behaviour of the return channel at the symbol level. Table 3.1 collects these values alongside the gesture and detection parameters introduced earlier in the chapter. Their combined effect on overall channel reliability is quantified in Chapter 7.

Table 3.1: Summary of Chapter 3 design parameters.

Parameter	**Symbol**	**Value**
Gesture (Follower) — Section 3.3		
Amplitude	A	0.12 m
Frequency	f_0	2.5 Hz
Tukey taper	—	0.1 s (20 samples)
Matched filter (Leader) — Sections 3.5, 3.6		
Template length	N	200 samples (1.0 s)
Detection threshold	λ	0.030
Noise floor (max, simulated)	$E_{\text{noise,max}}$	≈ 0.013
Typical gesture-peak energy	E_{gesture}	≈ 0.25
Threshold/noise margin	$\lambda / E_{\text{noise,max}}$	2.3 (3.6 dB)
Signal/threshold margin	$E_{\text{gesture}} / \lambda$	8.3 (9.2 dB)
Symbol timing (Leader) — Section 3.7		
Blanking period	—	400 samples (2.0 s)
Minimum duration	—	30 samples (0.15 s)
Debounce interval	—	100 samples (0.5 s)
HEARTBEAT/ACK boundary	DUR_HB_MAX	235 samples
ACK/ERROR boundary	DUR_ACK_MAX	400 samples
HEARTBEAT peak width (observed)	—	165–203 samples
ACK peak width (observed)	—	273–318 samples
ERROR peak width (observed)	—	470–501 samples

Chapter 4

Design of the Gesture-Based Return Channel

4.1 Altitude Gesture Generator

Chapter 3 established the shape of a single gesture: a sinusoid at $f_0 = 2.5\,\text{Hz}$ with amplitude $A = 0.12\,\text{m}$, tapered at onset and offset by a Tukey window, sustained for a duration that identifies which symbol is being transmitted. This chapter turns that specification into the component that actually produces it onboard the Follower: the gesture generator. Where Chapter 3 asked which parameters a detectable, unambiguous gesture must have, this section asks how those parameters are turned, at every control cycle, into the altitude reference offset that the Follower's controller is asked to track.

The generator is realized as a discrete-time block operating at the same sampling rate as the rest of the Follower's control loop, $F_s = 200\,\text{Hz}$ ($T_s = 0.005\,\text{s}$), consistent with the matched filter's sampling assumptions on the Leader's side (Section 3.5). Its external interface is narrow: a trigger pulse and a symbol code, the latter identifying which of the three alphabet entries defined in Section 4.3 is being requested. The amplitude and frequency are fixed constants, shared by every symbol and never passed in externally; only the symbol code varies between transmissions, and it is this code, not an externally supplied duration, that determines how long the sinusoid is sustained.

Internally, the generator's behaviour is organized around two operations that share a common piece of state rather than operating independently. On the sample at which a trigger is accepted, the symbol code is mapped to one of the three nominal durations established in Section 4.3, and a phase counter is reset to zero; this lookup happens once, at the onset of the symbol, not continuously. From that point on, at every subsequent sample, the generator advances the phase counter

and uses it, together with the duration selected at onset, to compute both the instantaneous value of the sinusoid and the value of the Tukey taper at that point in the symbol: the taper ramps up over the first samples of the gesture, remains at unit gain through the middle of the symbol, and ramps back down over the final samples, with the position of this final ramp determined directly by how close the phase counter is to the selected duration, exactly as motivated in Section 3.3. The tapered offset and the gating that determines when output is produced are therefore not independent operations on separate inputs, but two views of the same internal duration and phase state introduced at onset and advanced every sample thereafter.

The generator also exposes a gesture-busy flag, asserted for the entire duration of an active symbol and cleared once it ends. This flag is not consumed within the generator itself; it is the signal through which the arbitration logic developed in Section 4.5 knows whether the channel is currently occupied.

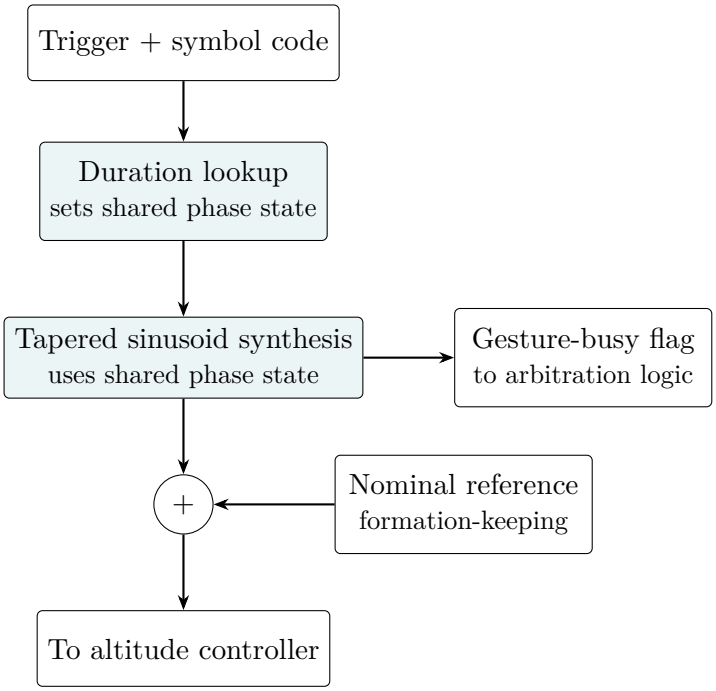


Figure 4.1: Block diagram of the altitude gesture generator: a trigger and symbol code select a nominal duration at onset, driving a shared duration and phase state that determines both the tapered sinusoid produced at each sample and the gesture-busy flag consumed by the arbitration logic of Section 4.5. The resulting offset is added to the Follower's nominal altitude reference before being sent to the altitude controller.

The output of the generator is not applied to the Follower's flight controller

directly as a new setpoint from outside the control loop, but added to the existing altitude reference already being tracked for formation-keeping purposes; this is what keeps the gesture consistent with the constraint, established in Section 1.5, that the return channel introduce no behaviour beyond what the existing control architecture already supports. The consequences of this addition for the Follower's tracking performance, and the specific operating frequency adopted at implementation level, are discussed next in Section 4.2.

4.2 Sinusoidal Modulation and Selected Operating Frequency

The values that define the gesture's sinusoidal component, $A = 0.12\,\text{m}$ and $f_0 = 2.5\,\text{Hz}$, were established in Section 3.3 together with the reasoning that constrains them; they are not re-derived here. What this section adds is the implementation-level question left open by that derivation: how the continuous sinusoid $A\sin(2\pi f_0 t)$ specified in Chapter 3 is actually realized by a generator operating in discrete time, at the sampling rate $F_s = 200\,\text{Hz}$ shared with the rest of the Follower's control loop.

At this sampling rate, one period of the gesture spans $F_s/f_0 = 80$ samples, corresponding to an angular resolution of $4.5°$ between consecutive samples of the sinusoid. This places the sampling rate roughly forty times above the Nyquist minimum required to represent a 2.5 Hz signal, and confirms that the operating frequency selected in Section 3.3 is bounded, in practice, only by the altitude controller's bandwidth and the formation-keeping dynamics discussed there, not by any additional constraint introduced by discretization: at this level of oversampling, the discrete sinusoid reproduces the continuous waveform closely enough that no distortion attributable to sampling reaches the Leader's detector.

Each gesture is generated with its phase reset to zero at the instant the symbol begins, so that $A\sin(2\pi f_0 t)$ is evaluated with t measured from the start of the current symbol rather than from a global clock. This choice has no bearing on detectability, since the I/Q matched filter introduced in Section 3.5 is constructed precisely to be insensitive to the phase at which a gesture arrives; its purpose here is implementational rather than functional, keeping the generator's internal state simple by avoiding the need to track absolute phase across successive, independently triggered symbols. Figure 4.2 shows the resulting offset for a single ACK symbol.

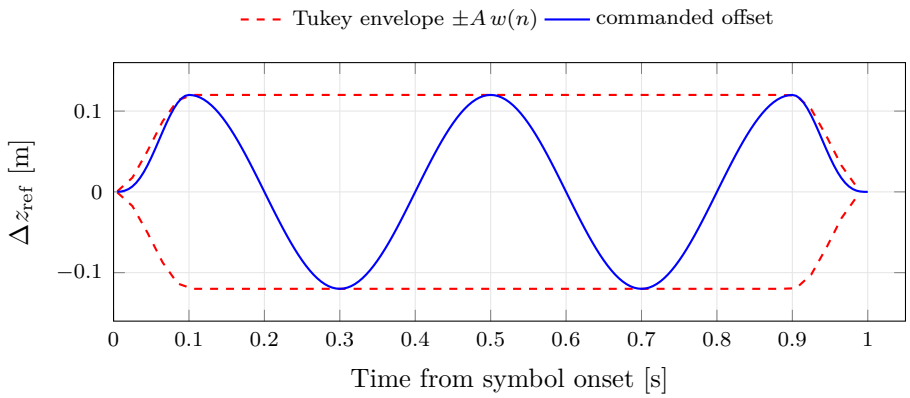


Figure 4.2: Commanded altitude offset for a single ACK gesture, generated at $F_s = 200\,\text{Hz}$. The Tukey envelope reaches full amplitude only after the first 20 samples and returns to zero over the last 20, so that the transition into and out of the gesture is continuous and the altitude controller is never asked to track a discontinuity.

The operating frequency is therefore fully specified by this point: fixed at $f_0 = 2.5\,\text{Hz}$ for every symbol, realized at $F_s = 200\,\text{Hz}$ with negligible discretization error, and generated independently for each transmitted gesture. The remaining degree of freedom in the alphabet, the duration for which this fixed waveform is sustained, is what distinguishes one symbol from another, and is defined next in Section 4.3.

4.3 Symbol Alphabet: HEARTBEAT, ACK and ERROR

Sections 4.1 and 4.2 established that every gesture shares the same waveform: a sinusoid at $f_0 = 2.5\,\text{Hz}$, amplitude $A = 0.12\,\text{m}$, tapered by a Tukey window. What distinguishes one transmitted symbol from another is therefore not the shape of the gesture but the duration for which it is sustained, as anticipated when this duration was first introduced as a design variable in Section 3.3. This section defines the resulting alphabet: the finite set of symbols the Follower can transmit, their semantic content, and the duration assigned to each.

Three symbols compose the alphabet. HEARTBEAT signals that the Follower is operating normally and remains within communication range, transmitted periodically regardless of any specific event, so that its absence over an extended interval is itself informative to the Leader. ACK confirms that a command issued through the forward channel has been received and acted upon, transmitted once in direct response to that event rather than on a fixed schedule. ERROR reports that

the Follower has detected an anomalous condition in its own state, transmitted as soon as the condition is identified, independently of the regular HEARTBEAT cycle. These three symbols correspond directly to the acknowledgement, liveness, and error-reporting requirements identified in Section 1.3 as the minimum functionality a return channel must provide, and reproduce, in a deliberately minimal form, the same class of signalling found in conventional acknowledgement-based protocols [17].

Table 4.1 lists the nominal duration assigned to each symbol. These values were fixed, together with the amplitude and frequency, as part of the design constraints discussed in Section 3.3: longer symbols are more easily distinguished from noise but slower to transmit and more likely to overlap with other formation events, and the three durations are spaced widely enough to remain distinguishable once the spreading introduced by the matched filter, characterized in Section 3.7 of Chapter 3, is taken into account. Figure 4.3 shows the three symbols on a common time axis. The classification boundaries used to separate these symbols at the receiver were derived from that same spreading behaviour and are not repeated here.

Table 4.1: Symbol alphabet and nominal gesture durations.

Symbol	**Meaning**	**Nominal duration**
HEARTBEAT	Periodic liveness indication	100 samples (0.5 s)
ACK	Forward command received and executed	200 samples (1.0 s)
ERROR	Anomalous condition detected onboard	400 samples (2.0 s)

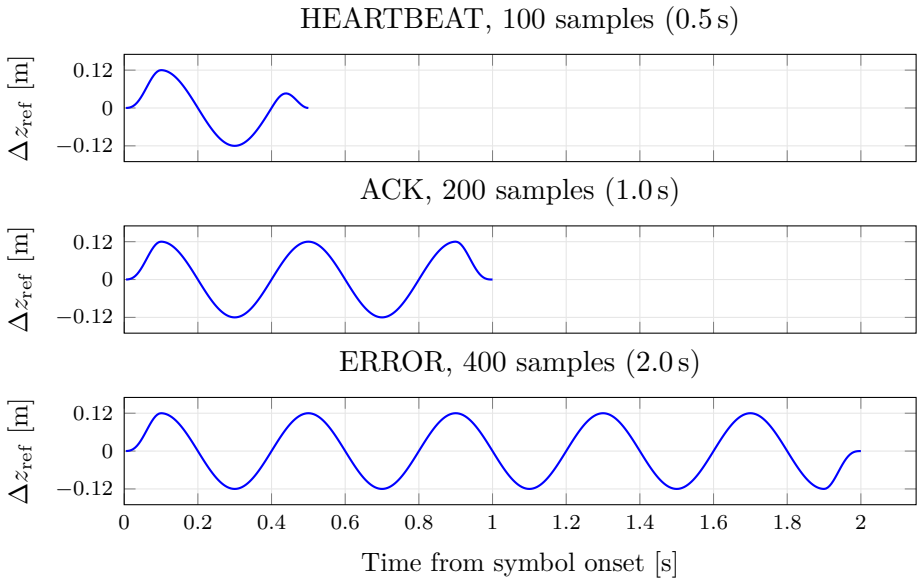


Figure 4.3: The three symbols of the alphabet on a common time axis. Amplitude, frequency and taper are identical in all three cases; only the duration for which the sinusoid is sustained distinguishes one symbol from another.

Encoding the alphabet in duration, rather than in frequency or amplitude, follows directly from the architecture already established for the gesture itself. Since every symbol shares the same f_0 and A, the Leader's receiver reduces to a single matched filter tuned to one frequency, rather than a bank of filters each tuned to a different candidate frequency or amplitude level; the alphabet is read entirely from how long the filter's output remains above threshold, a quantity the detector introduced in Section 3.5 already produces for the sole purpose of deciding whether a gesture is present at all. Distinguishing symbols by duration therefore adds no additional receiver complexity beyond what detection itself already requires, and ties the robustness of symbol classification directly to the same detection margin established in Section 3.6, rather than introducing a second, independent source of decision error.

4.4 Symbol Timing and Guard Intervals

Section 3.7 of Chapter 3 described how the Leader's receiver protects itself against ambiguous detections: a blanking period at start-up, a minimum-duration filter against noise bursts, and a debounce interval that prevents a single gesture from being registered twice. This section addresses a related but distinct problem, situated on the Follower's side of the channel: ensuring that when two genuinely

different symbols are transmitted close together in time, both are received as two separate events rather than as one merged detection or a lost one.

The risk arises directly from two properties already established elsewhere in this thesis. First, the matched filter's response to a gesture is wider than the gesture's nominal duration, spread by approximately the template length $N = 200$ samples as the correlation window slides onto and off the signal, as characterized in Section 3.7. Second, the same debounce interval that protects the receiver against fragmenting one gesture into multiple detections has the side effect of ignoring any new detection that begins before it expires, a property that is desirable when the second detection is a spurious echo of the first, but harmful if it is instead a genuinely new symbol. If the Follower were to transmit two symbols closer together than the combined extent of this spreading and the debounce window, the second symbol would risk being absorbed into the first's energy excursion, or discarded outright by the debounce logic, rather than being registered as a distinct event.

The generator therefore enforces a guard interval: a minimum spacing, measured from the end of one gesture's nominal duration to the start of the next, that any two consecutively transmitted symbols must respect. Sizing this interval is a matter of clearing both mechanisms with margin rather than exactly matching either: it must exceed the spreading tail produced by the first gesture, so that the second symbol's energy excursion opens only after the first has closed, and it must also exceed the debounce interval, so that the second detection is not suppressed as an echo of the first. The experiments of Section 7.4 locate this boundary empirically: two ACK symbols separated by 150 samples (0.75 s) are both recovered correctly, whereas at 100 samples the two excursions merge and the second is misclassified. A guard interval of 300 samples (1.5 s) is therefore adopted, twice the measured minimum, so that the spacing continues to clear both mechanisms as the excursions narrow at larger role-swap separations.



Figure 4.4: Why the guard interval is 300 samples. Each pair of blue blocks is two ACK symbols transmitted with the spacing shown; the bars beneath are the energy excursions they produce at the receiver, each wider than the symbol that generated it because the correlation window slides onto and off the signal. At 100 samples the two excursions overlap and the receiver reports a single event of the wrong duration class. At 150 they separate and both symbols are recovered, which locates the boundary empirically. The value adopted is twice that minimum, so that the spacing continues to clear both the spreading and the debounce interval as the excursions narrow at larger role-swap separations.

4.5 Encoding Strategy for Follower Status Feedback

The symbol alphabet defined in Section 4.3 gives the Follower three things it can say. This section describes how it decides what to say and when. Three conditions can arise independently of one another and of the generator's own timing. An onboard anomaly may be detected, a signal transmitted by the Leader through the forward channel may be received, and the periodic liveness indication is always eventually due. Since the generator introduced in Section 4.1 transmits one symbol at a time, these three conditions must be arbitrated into a single stream before reaching it.

The anomaly condition originates from a dedicated monitor that compares the Follower's pitch and roll, in radians, against a fixed bound of $\pm 0.5\,\mathrm{rad}$, and its altitude-tracking error, in metres, against a fixed bound of $0.25\,\mathrm{m}$. The monitor is disabled for the first 2000 samples ($10.0\,\mathrm{s}$) after start-up, so that the transient associated with take-off and initial stabilization is not itself read as an anomaly. Once an anomaly is flagged, the monitor stops evaluating these bounds entirely for

1000 samples (5.0 s). The condition itself may persist for longer than this interval, but the monitor reports it only once at the moment it is first detected, and does not report it again until the cooldown has expired and the bounds are checked once more. The result is one anomaly report per episode rather than a continuous signal for as long as the underlying condition holds.

The received-signal condition originates from the forward-channel decoder. That decoder does not reconstruct a message. It forms the norm of the Follower's pitch and roll response to the Leader's downwash and compares that norm against a threshold which adapts to the prevailing disturbance level rather than being fixed in advance, so that a crossing reflects a change in the wake the Follower is immersed in rather than an absolute magnitude. A crossing that survives the decoder's own debounce is reported as a single detection event. It is that event, and not any decoded content, which marks `ack_pending`: the Follower acknowledges having perceived a change in the Leader's behaviour, which is precisely what an acknowledgement means in this protocol, since the command it confirms was itself expressed as a change in the Leader's velocity rather than transmitted as a message.

Two properties of this output matter downstream. The first is that a detection carries a start and an end rather than a single instant, since the threshold must be crossed in both directions for the event to close; the duration between them is available to the Follower and is what allows the forward direction to carry an addressed announcement in Chapter 8. The second is that the delay between the Leader's change of velocity and the Follower's detection of it is small compared with the other contributions to the loop, and is treated as negligible in the budget assembled in Section 6.2. Whatever this signal is, once it arrives it is latched.

Both event-driven conditions are latched in the same way. The Follower sets a pending flag when either is raised, and that flag persists until the corresponding symbol has actually been transmitted, regardless of whether the generator happens to be busy at the moment the condition arises. Latching the anomaly condition is the more consequential of the two, and follows from an asymmetry examined in Section 5.7: a lost acknowledgement is absorbed by the timeout of Section 5.4 and produces no protocol-level consequence, whereas a lost error report leaves the Leader unaware of a fault the Follower has genuinely detected, with no mechanism elsewhere in the design to compensate for it. Since the anomaly monitor reports each episode once and then enters a cooldown, an ERROR pulse discarded because the channel was occupied would not be reissued until that cooldown expired. Because ERROR takes priority over ACK, a pending anomaly is served at the first sample on which the generator becomes free, so the latch delays the report by at most the residual duration of the symbol already in progress.

These two event-driven conditions, together with the passage of time, are arbitrated by a fixed priority order. ERROR is transmitted before ACK, and ACK before HEARTBEAT, whenever the generator is not currently transmitting.

This ordering reflects the relative urgency established when the alphabet was first motivated in Section 1.3. An anomaly is safety-relevant and should reach the Leader at the first opportunity. An acknowledgement can tolerate a short delay without consequence, since the signal it confirms is already in effect. The liveness indication is self-correcting, since a delayed HEARTBEAT costs nothing beyond a marginally later confirmation that the Follower is still present. When the channel becomes free and more than one condition is pending, only the highest-priority one is transmitted. Any lower-priority condition continues waiting and is re-evaluated at the next opportunity. Figure 4.5 summarises the resulting arbitration path.

The liveness indication is generated only in the complete absence of the other two conditions. A counter accumulates one step for every sample the channel is free and neither an anomaly nor a pending acknowledgement is being served. A HEARTBEAT is triggered once this counter reaches 1000 samples (5.0 s). The counter is reset to zero, not paused, on every sample during which the channel is occupied transmitting any symbol. As a result, two consecutive HEARTBEATs are never closer together than 5.0 s, but they can be further apart than that: every ACK or ERROR transmitted in between restarts the count from zero, delaying the next HEARTBEAT by at least the duration of that interruption.

The output of this arbitration is a trigger pulse together with the symbol code selected by priority. This is exactly the pair of inputs consumed by the gesture generator of Section 4.1. The busy flag that generator exposes is, in turn, the signal this section's arbitration logic uses to determine whether the channel is currently free.

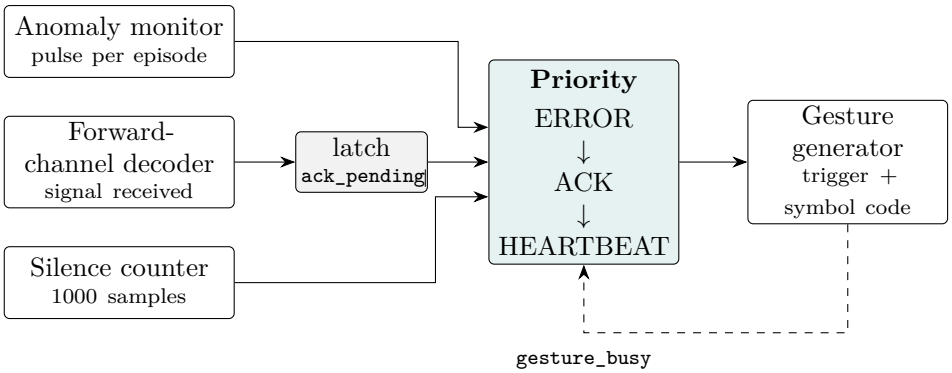


Figure 4.5: Arbitration of the three transmission conditions into a single stream of symbols. Only the forward-channel signal passes through a latch: an anomaly arriving while the generator is busy is not queued. The busy flag returned by the generator is what gates the whole selection.

4.6 Constraints Imposed by Flight Stability and Safety

Sections 3.3 and 3.1 already established two constraints that bound the gesture individually: the amplitude $A = 0.12\,\text{m}$ was sized to stay within a vertical excursion considered safe during close-proximity flight, and the role-swap manoeuvre itself already reduces the nominal vertical separation between the two vehicles before any gesture is added on top of it. Neither constraint, however, was examined in the context of the encoding strategy developed in Section 4.5, where symbols can be transmitted back to back whenever more than one condition is pending. This section considers what changes, from a safety standpoint, when a single gesture is no longer the unit of analysis.

The role-swap manoeuvre and the gesture amplitude interact directly. The descent performed by the Leader and the corresponding climb performed by the Follower already leave less vertical clearance between the two vehicles than in nominal formation flight, and the gesture's $\pm 0.12\,\text{m}$ oscillation is superimposed on top of whatever clearance remains once the manoeuvre is complete. The amplitude bound adopted in Section 3.3 accounts for the gesture in isolation, but does not by itself establish what minimum clearance the role-swap geometry must guarantee before a gesture is safely superimposed on it. The quantity that matters here is not the commanded amplitude but the excursion actually performed. As established in Section 7.2, the altitude loop transmits the reference at f_0 with attenuation rather than tracking it, so the peak excursion realized is 8.7 cm against the commanded $A = 0.12\,\text{m}$. It is this figure that sets the clearance during a role-swap, and Section 7.5 evaluates the resulting margin directly: at the lower end of the recommended separation range the instantaneous gap between the two vehicles closes to 11.3 cm at the bottom of every cycle, and since a Parrot Mambo measures 40 mm in height the two airframes together occupy 8.0 cm of it, leaving 3.3 cm of clearance. Contact would occur at a nominal separation of 16.7 cm, which lies below the separation at which classification has already begun to degrade, so the channel becomes unreliable before the manoeuvre becomes unexecutable. The clearance is therefore not the binding constraint, but neither does it carry a comfortable margin, and this is why the recommended range of Section 7.5 does not extend to the lower edge of the interval over which classification is still correct. The limiting case of two symbols transmitted in rapid succession does not worsen this. Each gesture is an oscillation about the same altitude reference, and the guard interval of Section 4.4 places 300 samples between the end of one and the start of the next, so the excursions do not overlap and the peak displacement is set by a single symbol however many follow it. What a succession of symbols changes is the length of time the formation spends in the reply configuration, not the depth to

which it descends within it, and that duration is bounded for a different reason in Section 6.4: a Follower that transmits continuously is a Follower that is not sensing. A related consideration follows from the arbitration logic itself. Because ERROR takes priority over ACK, and a pending ACK is served as soon as the channel frees up, a sequence of events on the Follower's side can, in the worst case, result in several symbols being transmitted with only the guard interval discussed in Section 4.4 between them, rather than the isolated, well-separated gestures assumed when the amplitude and frequency were first justified. Repeated oscillation over an extended interval is not obviously equivalent, from a flight-stability standpoint, to a single isolated gesture, particularly under degraded conditions such as reduced battery margin or ambient wind. No additional bound on the number of consecutive symbols is required on clearance grounds. The argument above establishes that the peak excursion is a property of one gesture rather than of a sequence, so the worst case for vertical separation is the case already analysed, and a constraint on symbol count would restrict the channel without buying any margin the single-gesture analysis has not already secured. The constraints that do bound how often symbols may be transmitted are of a different nature and are established elsewhere: the guard interval of Section 4.4, which exists so that two symbols are received as two events, and the sensing budget of Section 6.4, which exists so that the Follower spends the large majority of its flight in the configuration from which it can track the Leader at all.

These constraints are not treated here as newly derived safety bounds, but as open questions that the individual-gesture analysis of Chapter 3 does not answer by itself. Their resolution, together with the numerical values flagged as open in this section and in Section 4.5, is addressed in Chapter 7, where the role-swap manoeuvre and the encoding strategy are exercised together in simulation rather than considered separately.

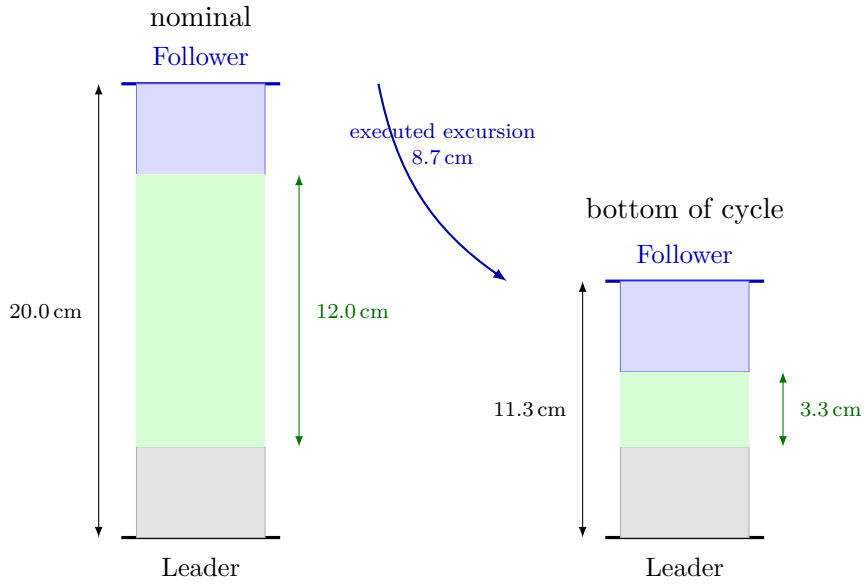


Figure 4.6: Vertical clearance during a role-swap. The gesture oscillates the Follower about the nominal separation, so the instantaneous gap closes by the excursion actually executed rather than by the amplitude commanded. At the lower end of the recommended range the gap reaches 11.3 cm at the bottom of every cycle, of which the two airframes occupy 8.0 cm. The remaining clearance is real but not generous, and it is exhausted at almost the same separation at which the classification margin is, which is why the recommended range stops short of the interval over which classification alone would still be correct.

4.7 Implementation of the Gesture Generator on the Follower UAV

The gesture generator described in Section 4.1, together with the arbitration logic of Section 4.5, is not a standalone algorithm but a component embedded directly in the Follower's control loop. This section describes how that embedding is realized, as a matter of software architecture rather than of the signal-processing content already covered.

Each of the functions involved, the anomaly monitor, the arbitration logic, and the generator itself, is implemented as a MATLAB Function block within the Simulink model of the Follower, targeting the Parrot Mambo minidrone platform. The implementation described here is not a simulation-only construct: the same

blocks are compiled to C and executed on the Mambo itself, and the results of Chapter 7 were obtained in simulation and then verified against the vehicle. Every such block carries the `#codegen` directive, marking it for C code generation rather than interpreted execution, and every piece of state that must persist from one control step to the next, the phase counter and selected duration inside the generator, the pending-acknowledgement and pending-anomaly flags and the silence counter inside the arbitration logic, the cooldown counter inside the anomaly monitor, is declared as a persistent variable local to its own block rather than shared through global state. This is not an incidental implementation choice: it is what allows each block to be compiled independently into embeddable C code while still retaining memory across calls, which a stateless function could not do.

Within the control loop, the generator occupies a fixed position relative to the rest of the Follower's altitude control chain. It is invoked once per control step, at the shared sampling rate $F_s = 200\,\text{Hz}$ established in Section 4.2, immediately after the arbitration logic of Section 4.5 has decided whether a new symbol should begin. Its output, the tapered altitude offset, is added to the nominal altitude reference before that reference reaches the altitude controller, exactly as illustrated in Figure 4.1; the controller itself is unmodified by the presence of the return channel and simply tracks whatever reference it is given; consistent with the requirement, established in Section 1.5, that the return channel introduce no behaviour beyond what the existing control architecture already supports.

The interface between the three blocks is narrow and matches exactly what was specified when each was introduced. The anomaly monitor exposes a single boolean output, consumed by the arbitration logic; the arbitration logic exposes a trigger pulse and a symbol code, consumed by the generator; and the generator exposes a busy flag, consumed back by the arbitration logic. No additional signal passes between these three blocks, and no other part of the Follower's software reads or writes their internal state, which keeps the return channel's implementation contained to a small, independently testable subset of the overall Simulink model rather than distributed across the Follower's broader control software.

Chapter 5

Leader-Side Decoding and Protocol Logic

5.1 I/Q Matched Filter for Gesture Detection

Section 3.5 in Chapter 3 established why detection is built around an in-phase/quadrature correlation rather than a single real-valued one: the resulting energy measure is insensitive to the unknown phase at which a gesture arrives. This section does not repeat that reasoning. It describes instead how the matched filter is actually realized as a running computation on the Leader, operating continuously on the incoming dz signal after the high-pass filtering stage of Section 3.4.

The detector operates on the output of the high-pass stage of Section 3.4, defined in Equation (3.4), whose transient time constant of $\tau \approx 99$ samples returns below in Section 5.3 as the quantity that sizes the blanking interval.

The two reference waveforms against which the incoming signal is correlated, an in-phase sine and a quadrature cosine, both at $f_0 = 2.5\,\text{Hz}$ and sampled at $F_s = 200\,\text{Hz}$, are not recomputed at every sample. They are generated once, the first time the detector runs, and normalized so that each reference carries unit energy over its own length; from that point on they are held fixed and reused for every subsequent correlation. This is possible because the reference waveforms depend only on the fixed operating frequency established in Section 4.2, never on the incoming signal itself, so there is nothing to update between one gesture and the next.

The incoming signal is held in a circular buffer of $N = 200$ samples, matching the length of the correlation template, so that at every sample the buffer contains exactly the most recent window of dz required for correlation. Rather than recomputing the correlation from scratch over the entire buffer at every sample, each new sample simply overwrites the oldest entry in the buffer, and the correlation is evaluated

against this sliding window once per control step, at the same $F_s = 200\,\mathrm{Hz}$ rate as the rest of the Follower and Leader control loops. The output of this stage, on every sample, is the pair of correlation values against the in-phase and quadrature references; how these two values are combined into a single decision quantity is treated next, in Section 5.2.

Writing $y[\cdot]$ for the high-pass output and s_I, s_Q for the two unit-energy references, the correlator evaluates at every sample

$$I[n] = \sum_{m=0}^{N-1} y[n-m]\, s_I[m], \qquad Q[n] = \sum_{m=0}^{N-1} y[n-m]\, s_Q[m], \tag{5.1}$$

and combines the two into a single scalar,

$$E[n] = I^2[n] + Q^2[n], \tag{5.2}$$

which is the quantity every subsequent stage of the decoder operates on. The sum of squares is what removes the dependence on the phase at which the gesture arrives, as established in Section 3.5.

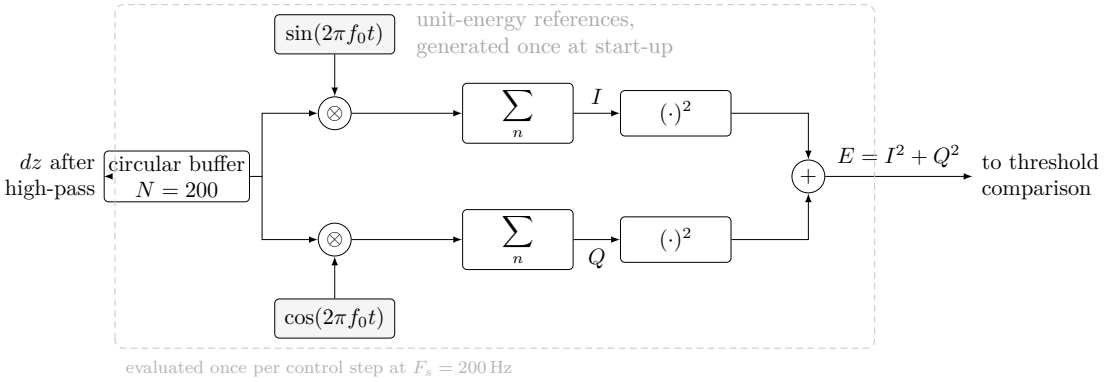


Figure 5.1: Structure of the in-phase/quadrature matched filter on the Leader. The high-pass-filtered vertical velocity estimate is held in a sliding window matching the correlation length, and correlated against two fixed unit-energy references at f_0. Squaring and summing the two correlation values yields an energy measure that does not depend on the phase at which the gesture arrives, which is the property established in Section 3.5 and the reason the quadrature pair is used rather than a single real correlation. The whole chain is evaluated once per control step; what becomes of E is the subject of Section 5.2.

5.2 Energy Detector and Decision Threshold

Section 5.1 left off with a pair of correlation values produced at every sample, one against the in-phase reference and one against the quadrature reference. This section

describes how that pair becomes a single detection decision, without repeating why the underlying threshold value was chosen, which was already established in Section 3.6.

At every sample, the two correlation values are combined into a single energy quantity by summing their squares, and this energy is compared against the fixed threshold $\lambda = 0.030$. The comparison itself is not a one-shot decision but drives a small state machine internal to the detector, distinct from the leader-side protocol FSM developed later in Section 5.4: while energy remains below threshold, the detector stays idle; once it crosses threshold, the detector begins counting how long the crossing persists; and once energy falls back below threshold, the accumulated count is what allows the detector to decide whether a genuine symbol occurred at all, and if so, which one.

A detection is declared while the energy remains above a fixed threshold,

$$d[n] = \begin{cases} 1, & E[n] > \lambda, \\ 0, & \text{otherwise}, \end{cases} \qquad \lambda = 0.030, \tag{5.3}$$

and the interval over which $d[n]$ remains at unity, rather than its amplitude, is what carries the symbol.

This internal state machine is where the protections discussed in Section 5.6 take effect at the implementation level: a crossing that ends before accumulating enough samples is discarded rather than reported, and a new crossing occurring too soon after the end of a previous one is ignored until the debounce interval has elapsed. Both checks are evaluated inside this same internal loop, before any output ever reaches the leader-side FSM. Figure 5.2 shows this loop.

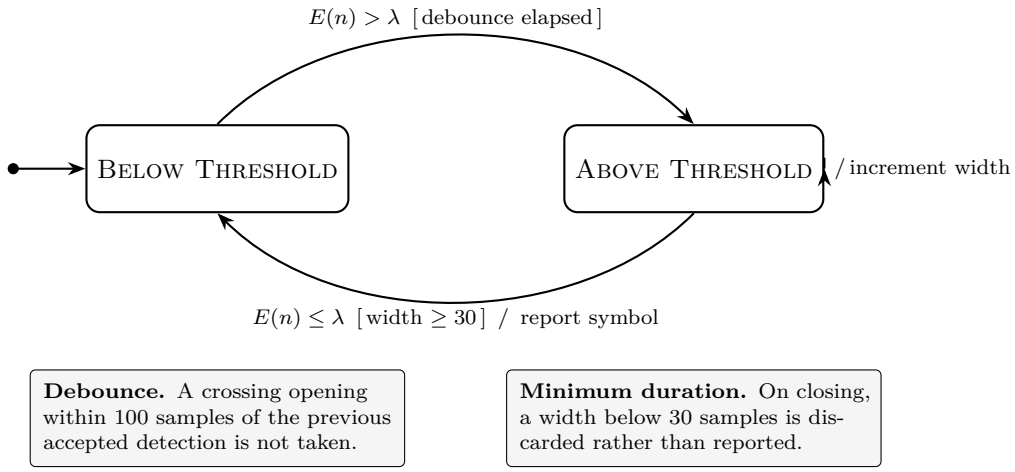


Figure 5.2: State machine internal to the energy detector, distinct from the protocol machine of Figure 5.4. Guards are shown in square brackets: nothing is reported while energy stays below threshold, the width accumulates for as long as a crossing persists, and the symbol is decided only once the crossing closes. Both protections act inside this loop, so a discarded crossing never reaches the protocol layer.

Only once a crossing has ended, and has been accepted as a valid symbol by the checks above, does the detector produce an output: a single pulse indicating that a symbol has been detected, together with a classification of which symbol it was, based on how long the energy remained above threshold. This classification reuses the boundaries established in Section 3.7 of Chapter 3 without re-deriving them here. It is this output, and only this output, that is visible to the leader-side protocol logic developed in the remainder of this chapter.

5.3 Windowing, Synchronization and Symbol Validation

The detector described in the previous two sections runs continuously from the moment the Leader's control loop starts, but its output is not meaningful from the very first sample. This section covers the conditions under which the detector's output is treated as valid at all, and the reasoning behind the absence of any additional filtering stage between the high-pass filter of Section 3.4 and the correlation itself.

Detection is suppressed for the first 400 samples (2.0 s) after start-up, for the reasons established in Section 3.7: the interval covers both the filling of the circular buffer introduced in Section 5.1 and, more restrictively, the settling of the high-pass

stage that precedes it. Suppressing the output over this window is what prevents the Leader from acting on a detection that reflects the receiver's own initial transient rather than any signal transmitted by the Follower. The interval is therefore sized to the receiver's own dynamics rather than to any property of the gesture, and no synchronization signal is exchanged between the two vehicles at any point: the detector never needs to know when a gesture begins, only whether its energy output crosses threshold, which is the property the quadrature structure of Section 3.5 was adopted to provide.

No band-pass filtering is applied between the high-pass stage and the correlation. The matched filter is already selective at f_0 by construction, since its reference waveforms contain no other frequency component, so a band-pass stage tuned to the same frequency would add no selectivity that the correlation does not already provide. It would, however, actively degrade detection: narrowing the input bandwidth reduces the number of statistically independent noise samples within the correlation window, and the processing gain of a matched filter depends directly on that number rather than on the window length alone. A sufficiently narrow band-pass stage upstream would therefore leave the correlation window nominally unchanged in length while reducing its effective content to a small fraction of independent samples, collapsing the gain the matched filter is there to provide.

The magnitude of this effect is easily estimated. Over a correlation window of $T = N/F_s = 1.0\,\mathrm{s}$, a signal sampled at $F_s = 200\,\mathrm{Hz}$ carries $N = 200$ statistically independent noise samples, and the matched filter's processing gain is correspondingly $10\log_{10} N \approx 23\,\mathrm{dB}$. A band-pass stage of bandwidth Δf placed upstream reduces the number of independent samples within the same window to $N_{\mathrm{eff}} \approx 2\,\Delta f\,T$, since the filtered noise is correlated over a time of order $1/\Delta f$. A band-pass of $\Delta f = 5\,\mathrm{Hz}$ around f_0, already generous relative to the gesture's own bandwidth, would leave $N_{\mathrm{eff}} \approx 10$ and a gain of $10\,\mathrm{dB}$, a loss of $13\,\mathrm{dB}$ with respect to the unfiltered case; a narrower stage of $1\,\mathrm{Hz}$ would leave $N_{\mathrm{eff}} \approx 2$ and collapse the gain to $3\,\mathrm{dB}$. The correlation window would remain nominally 200 samples long in both cases while delivering a small fraction of the gain it is there to provide.

Validation of an individual symbol therefore rests entirely on the detector's own output, without reference to any external timing or synchronization: a crossing is accepted as a symbol if it occurs after the blanking interval, if it persists long enough to clear the minimum-duration check, and if it is sufficiently separated from the previous accepted crossing. The specific thresholds involved were established in Chapter 3 and are examined, in terms of which false transitions they prevent at protocol level, in Section 5.6.

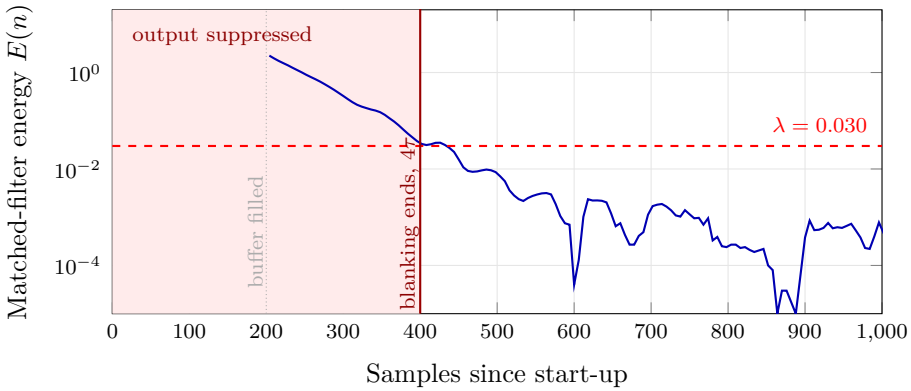


Figure 5.3: Why the blanking period is 400 samples rather than 200. At start-up the high-pass stage of Section 3.4 sees the Leader's standing vertical bias as a step, and its response decays with a time constant of $\tau \approx 99$ samples. The matched filter, correlating against that decaying transient, produces an energy that exceeds the detection threshold by orders of magnitude and remains above it well after the correlation buffer has filled. Sizing the blanking interval to the buffer alone would therefore leave the Leader acting on its own start-up transient; sizing it to four time constants of the filter that precedes the buffer brings the energy back below threshold before any output is accepted.

5.4 Finite-State Machine for Leader-Side Interpretation

The detector of Sections 5.1 and 5.2 reports isolated events: a symbol was detected, and it was of a given type. Nothing in that report indicates whether the symbol was expected, whether it answers a command the Leader issued earlier, or what the Leader should do about it. Interpreting those events in context is the role of the protocol state machine described here, which sits between the detector's output and the reaction logic of Section 5.5.

The machine has three states, and only one class of symbol is interpreted through them. In IDLE, the Leader has no outstanding command. The transition out of this state is driven not by the detector but by the Leader's own behaviour: when the Leader issues a command through the forward channel, the machine moves to WAITING ACK and starts a timer. A detected ACK then moves it to ACK RECEIVED, a transient state occupied for a single sample, which emits a pulse to the reaction logic and returns immediately to IDLE. An ACK detected in any other state is ignored, since an acknowledgement is meaningful only as the answer to a command that is actually outstanding.

If no acknowledgement arrives within 600 samples (3.0 s) of entering WAITING ACK, the timer expires and the machine returns to IDLE without emitting anything. This is a deliberate design choice rather than an omission: an unacknowledged command produces no protocol-level reaction at all. The Leader simply stops waiting and continues operating as before. Treating a missing acknowledgement as an error would conflate two distinct situations, since a command that goes unconfirmed does not by itself indicate that the Follower has detected a fault.

Two mechanisms operate outside this state graph, and both do so for the same reason: the conditions they report are not answers to anything the Leader has done. The first is error reporting. The anomaly monitor of Section 4.5 responds to the Follower's own attitude and altitude, independently of the forward channel, so a fault can be signalled at any moment and has no reason to fall within the interval during which the Leader happens to be awaiting an acknowledgement. A detected ERROR therefore emits its pulse regardless of the current state and returns the machine to IDLE. Confining error handling to WAITING ACK would have discarded the majority of reported faults, as shown in Section 7.6.

The second is liveness. A timer accumulates continuously and is reset to zero every time a HEARTBEAT is detected, in any state. If it reaches 2000 samples (10.0 s) without such a reset, the Follower is no longer considered active and a liveness flag is cleared. The flag recovers automatically: the next heartbeat to arrive resets the timer and restores it, without any state transition being required. Command acknowledgement is episodic and tied to a specific request, whereas liveness is a continuous property of the link that must remain observable even when no command has been issued for a long time.

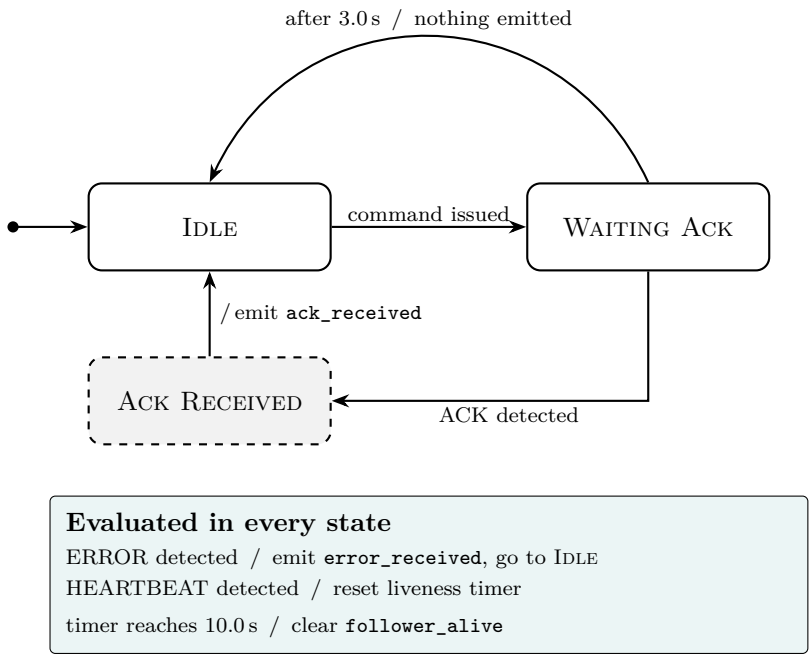


Figure 5.4: Leader-side protocol state machine, drawn as a statechart: transitions are labelled *event / action*, and the dashed state is transient, occupied for a single sample. Only acknowledgements are interpreted through the states, since an ACK is meaningful solely as the answer to an outstanding command. Error reporting and liveness are handled outside the state graph, because neither is a reply to anything the Leader has done.

The outputs of this machine are therefore three: a pulse when an acknowledgement is confirmed, a pulse when an error is reported, and a continuous flag indicating whether the Follower is still considered active. What the Leader actually does in response to each of these is the subject of the following section.

5.5 ACK and ERROR Handling Logic

The state machine of Section 5.4 reports what happened; it does not change how the Leader flies. That translation is the role of a separate reaction stage, which consumes the three outputs of the state machine and decides whether the Leader's nominal velocity command should be overridden. Keeping the two stages separate means the protocol logic can be reasoned about independently of the flight behaviour it triggers, and that the reaction stage never needs to know how a symbol was detected or classified.

The reaction stage holds three conditions of its own. Under normal operation

it does nothing: the Leader's velocity command passes through unmodified, and an acknowledgement pulse arriving in this condition is absorbed without effect. This is deliberate, since a confirmed acknowledgement means the Follower has received and acted on the command, which is precisely the case in which no change of behaviour is warranted. An error pulse, by contrast, moves the reaction stage into a recovery condition in which the Leader's velocity is overridden and held at zero, bringing the formation to a halt rather than allowing the Leader to continue advancing while the Follower is in a faulted state. The recovery condition persists for 2000 samples (10.0 s), after which the Leader resumes normal operation on the assumption that the Follower has had the opportunity to recover.

This recovery interval can also end early. If an acknowledgement arrives while the Leader is in recovery, the reaction stage returns immediately to normal operation without waiting for the remaining time to elapse. The reasoning is that an acknowledgement can only be produced by a Follower that is receiving forward-channel commands and responding to them, which is stronger evidence of recovery than the mere passage of time. The fixed interval therefore acts as an upper bound on how long the Leader waits in the absence of any positive indication, rather than as a fixed penalty applied to every reported error.

The liveness flag is handled differently from either pulse, because it is a continuous condition rather than an event. Whenever the flag indicates that the Follower is no longer considered active, the reaction stage enters a condition that takes precedence over both normal operation and error recovery: the Leader's velocity is overridden and held at zero for as long as the flag remains clear, with no timeout and no early exit. The Leader resumes only when the flag is restored, which, as established in Section 5.4, happens automatically when a HEARTBEAT is next detected. The absence of a timeout here is the substantive difference from error recovery: a reported error is a condition the Follower has itself signalled and can be expected to recover from, while a lost Follower has signalled nothing at all, and there is no basis on which the Leader could decide that continuing to advance has become safe again.

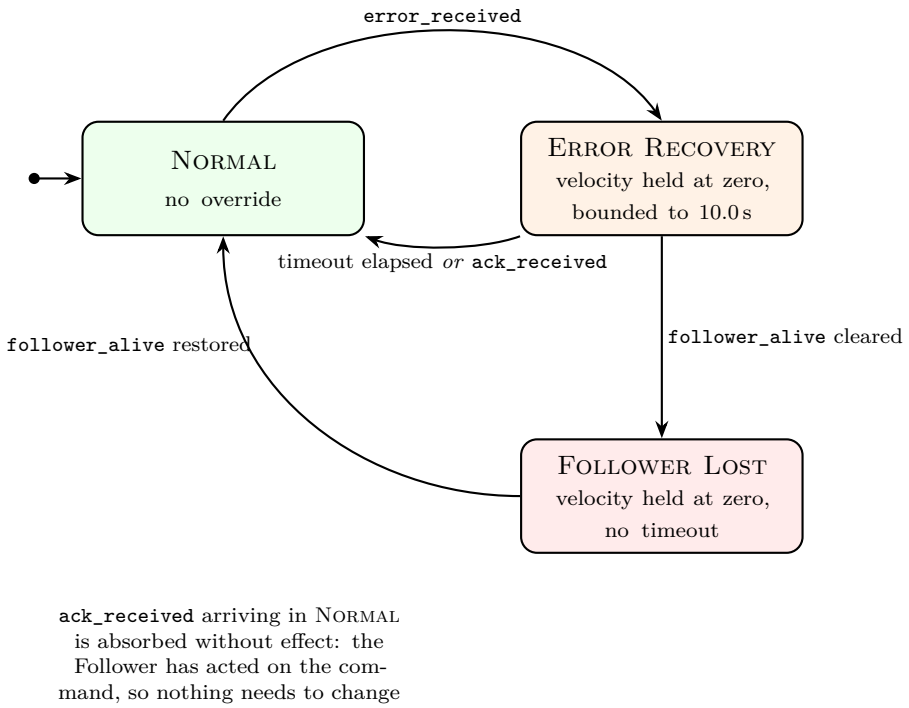


Figure 5.5: Reaction stage on the Leader. Loss of liveness takes precedence over both other conditions. The two halts differ in kind: error recovery is bounded and can be cut short by an acknowledgement, whereas a lost Follower halts the formation with no timeout, since silence carries no evidence that advancing has become safe.

Taken together, the two stages give the Leader a graded response. A confirmed acknowledgement changes nothing, a reported error stops the formation for a bounded interval that can be cut short by evidence of recovery, and a Follower that has stopped signalling altogether stops the formation indefinitely. In every case the response is a velocity override applied to the Leader alone, leaving the Follower's own control behaviour, and the forward channel through which it senses the Leader, untouched. Separating the two stages in this way is what keeps signal detection independent of flight behaviour: the state machine can be reasoned about, and its timing verified, without reference to what the Leader does in response.

5.6 Protection Against False Transitions

The protections built into the detector were introduced in Chapter 3 in terms of what they do to the energy signal: a crossing shorter than 30 samples is discarded, and a new crossing occurring within 100 samples of the previous one is ignored.

This section considers them from the opposite direction, asking which protocol-level consequences they exist to prevent, now that the state machine of Section 5.4 and the reaction stage of Section 5.5 have been defined.

The consequences are not symmetric across the three symbols, and this asymmetry is what makes the protections worth examining individually. A spurious HEARTBEAT is the mildest case: it resets the liveness timer, delaying by up to 10 s the moment at which a genuinely absent Follower would be recognized as lost. A spurious ACK is more consequential in one specific situation, since arriving during error recovery it terminates that recovery early, returning the Leader to normal operation on the strength of evidence that was never actually transmitted. A spurious ERROR is the most disruptive of the three, halting the formation for a full 10 s recovery interval in response to a fault that was never reported. This last case carries more weight than it would have under a state machine that examined incoming symbols only while an acknowledgement was outstanding, since error handling is placed outside the state graph in Section 5.4 and a spurious ERROR is therefore acted upon in any state.

The minimum-duration check addresses the mechanism most likely to produce any of these: an isolated noise excursion that briefly pushes energy above threshold. Such an excursion is short by nature, since it reflects a momentary fluctuation rather than a sustained periodic component, and requiring a crossing to persist for at least 30 samples removes it before it reaches the state machine at all. Notably, this check also shapes which spurious symbol a surviving excursion could be mistaken for: the classification boundaries of Section 3.7 assign the shortest accepted durations to HEARTBEAT, so a noise event that does clear the minimum-duration check is far more likely to be classified as the least disruptive of the three symbols than as an ERROR.

The debounce interval addresses a different mechanism, one that arises from a genuine gesture rather than from noise. Energy does not fall cleanly below threshold at the end of a gesture; it can dip below and recover briefly as the correlation window slides off the signal. Without protection, a single transmitted symbol could be reported twice, and the consequences at protocol level depend on which symbol was involved. A duplicated ERROR would restart the recovery interval just as the first was progressing. A duplicated ACK is discarded outright in IDLE, where no acknowledgement is outstanding, but could terminate a legitimate recovery early if it arrives while one is in progress. Suppressing any second report within 100 samples of the first removes this class of failure entirely, at the cost of an upper bound on how quickly two genuinely distinct symbols can be reported, which is the constraint discussed from the transmitting side in Section 4.4.

Neither protection acts on the state machine itself, which has no mechanism for rejecting an implausible symbol once it is reported. This is a consequence of the separation described in Section 5.5: the state machine interprets what it is told,

and the reaction stage acts on that interpretation, so both depend on the detector having already discarded what should never have been reported. The blanking interval of Section 5.3 completes this arrangement by suppressing the detector's output entirely during the period when its correlation carries no meaning, ensuring that the state machine begins operating only once the events it receives are ones it can meaningfully interpret.

5.7 Recovery from Missed or Corrupted Symbols

Section 5.6 considered what happens when the detector reports a symbol that was never transmitted. This section considers the opposite failure: a symbol that was transmitted but never reported, or reported as the wrong one. The return channel has no retransmission mechanism and no way for the Follower to learn that a symbol was lost, so recovery here does not mean repairing the loss. It means that the loss must not leave the system in a state from which it cannot proceed.

For HEARTBEAT the argument is direct. Each transmission is independent of the ones before it, and the liveness timer needs only one detection within its 10 s window to be reset. Since the Follower attempts a HEARTBEAT after every 5 s of channel silence, a single missed detection normally leaves a further opportunity within the window, and the Leader's assessment of liveness is unaffected. The margin is narrower when the channel has been occupied in the meantime, since the silence counter of Section 4.5 is reset rather than paused by any symbol transmitted, so consecutive heartbeats can be further apart than the nominal interval. Only a sustained inability to detect heartbeats causes the liveness flag to clear, which is the intended behaviour rather than a failure of it: from the Leader's perspective a Follower whose gestures cannot be detected is indistinguishable from one that has stopped transmitting, and in both cases halting the formation is the appropriate response.

A missed ACK is absorbed by the timeout established in Section 5.4. The state machine leaves Waiting Ack after 3.0 s regardless of whether the acknowledgement arrived, so no command can leave the Leader waiting indefinitely. The cost of the loss is that the Leader never learns the command was received, but as established in that section this produces no protocol-level reaction, and the Leader continues operating. A missed ERROR is the more serious case, since the Leader does not halt for a fault the Follower has genuinely detected. Two mechanisms bound this exposure. The first acts at the transmitter: as established in Section 4.5, the anomaly condition is latched, so a fault raised while the channel is occupied is transmitted at the first free sample rather than discarded. The second is the cooldown of the anomaly monitor itself: if the underlying condition persists, it is reported again once the cooldown expires. What remains uncovered is a transient

fault whose single latched report is transmitted but not detected, and for this the design offers no recovery. This residual case is accepted rather than solved, on the grounds that a fault severe enough to warrant halting the formation is unlikely to clear within the cooldown interval and therefore unlikely to go unreported altogether.

Misclassification is a distinct failure from loss, and its consequences follow the same asymmetry described in Section 5.6. A symbol classified as a shorter one than transmitted is the more likely direction of error, since it requires only that the energy excursion fall short of a boundary, and it is also the less damaging: an ERROR read as an ACK, or an ACK read as a HEARTBEAT, produces a weaker reaction than the one intended rather than a spurious strong one. In the second of these the reaction is not merely weaker but absent, since an ACK read as a HEARTBEAT leaves the Leader waiting for an acknowledgement that has already arrived, while resetting the liveness timer as though the Follower had reported nothing but its own presence. The opposite direction, in which a HEARTBEAT is read as an ERROR, would require the observed peak width to more than double, which the margins established in Section 3.7 make substantially less likely than the boundaries themselves would suggest.

What allows the system to tolerate all of these without a recovery protocol is that no symbol carries state forward. Each transmission is self-contained, no symbol acknowledges a previous symbol, and the state machine returns to IDLE after every event and every timeout. A lost or misread symbol therefore affects only the interpretation of that one event, and both vehicles remain able to communicate normally from the next transmission onward. The design is degraded by symbol loss rather than broken by it, which is the property that makes the absence of a retransmission mechanism acceptable in the first place.

5.8 Computational Implementation of the Decoder

The four stages described in this chapter, the correlation, the energy detector, the protocol state machine and the reaction logic, are implemented on the Leader in the same manner as their counterparts on the Follower described in Section 4.7. Each is a MATLAB Function block within the Simulink model, marked for C code generation, holding its own state in variables that persist across control steps rather than in any shared structure. The whole chain runs once per control step at $F_s = 200\,\text{Hz}$.

The decoder's position in the Leader's software is what distinguishes it from the generator on the Follower. It reads the vertical velocity estimate produced by the Leader's state estimator, which is already available as an output of the

existing control architecture, and it reads this signal in parallel with the control loop rather than in series with it. Nothing the decoder computes is fed back into altitude control, and the altitude controller is unaware that the signal it depends on is also being examined for gestures. This is what allows the return channel to be added to the Leader without modifying any part of its flight control behaviour, in keeping with the constraint established in Section 1.5.

The single point at which the decoder does affect the Leader's behaviour is the velocity override produced by the reaction stage of Section 5.5. This is applied to the Leader's velocity command, not to its altitude control, which keeps the return channel's influence confined to the horizontal motion of the formation and leaves the vertical axis, on which both channels depend for sensing, entirely undisturbed.

Two aspects of the implementation deserve note for what they avoid. The first is that the correlation is evaluated over the full template length at every control step, rather than incrementally. This is the more expensive of the available formulations, but it keeps the computation stateless beyond the circular buffer itself, avoiding the accumulation of numerical drift that a recursive update would introduce over the long run times involved.

The cost is bounded and small. At $F_s = 200\,\mathrm{Hz}$ each control step has a budget of $5\,\mathrm{ms}$, and a full-length correlation against two references requires $2N = 400$ multiply-accumulate operations per sample, or 8×10^4 per second. Even without exploiting the incremental formulation, this occupies a small fraction of the budget available on the target processor, and the detector is not the dominant cost in the Leader's control loop. The more expensive formulation is therefore adopted not because efficiency is irrelevant but because there is no need to trade numerical robustness for a saving that would not be noticed.

The second is that the detector determines a symbol's classification only after its energy excursion has ended, rather than attempting to identify the symbol while it is still in progress. This introduces a delay between transmission and detection of roughly the correlation window length, a delay that the state machine's timeouts of Section 5.4 were sized to accommodate rather than to eliminate.

The result is a decoder that occupies a small and clearly bounded part of the Leader's software. It consumes one signal that the platform already produces, it produces one override that acts on a single axis of motion, and it holds no state that any other part of the system can read or modify. The same properties that make the generator on the Follower independently testable, described in Section 4.7, apply here for the same reasons.

Chapter 6

Integration with the Forward Wake-Sensing Channel

6.1 Role of Leader-to-Follower Sensing in the Shared System

The preceding chapters treated the return channel largely in isolation, as a mechanism for carrying three symbols from the Follower to the Leader. That treatment was sufficient to design and analyse the channel itself, but it leaves unanswered a question that only becomes visible at system level: what the return channel contributes to the formation, given that the formation's objective is achieved through the forward channel alone. This chapter addresses the interaction between the two channels, beginning with the asymmetry in their roles.

The formation's objective is defined with respect to a fixed point in space, chosen before flight, which represents the rendezvous point the formation is intended to reach. This point is not a vehicle. It has no dynamics and no inputs, and its state is constant throughout the flight. Its role is to make the outcome of the formation a matter of design rather than of initial conditions: agents that agree among themselves converge to a value determined by where they happened to start, whereas agents anchored to a fixed reference converge to a value imposed in advance. The Leader knows this reference and is programmed to reach it. The Follower does not, and has no means of measuring its own absolute position at any point.

This is what makes the forward channel structurally indispensable. The Follower's only access to the reference is through the Leader, since the information reaches it exclusively along the path from the reference to the Leader and from the Leader to the Follower. The wake-sensing channel is the second link in that path.

If it fails, the Follower is not merely deprived of commands but disconnected from the reference entirely, with no alternative route by which the information could arrive. The forward channel therefore determines whether the formation objective is reachable at all.

The return channel occupies a different position, and one that is worth stating precisely because it is easy to overstate. It carries information from the Follower to the Leader, but the Leader's own behaviour with respect to the reference is not derived from that information: the reference is known to the Leader independently, and no symbol the Follower transmits alters it. The return channel thus never affects the value the formation converges to. What it affects is what happens when the assumptions behind that convergence stop holding. A Follower that has detected a fault, or one whose gestures have stopped arriving, is a Follower that may no longer be tracking the Leader, and without the return channel the Leader would continue advancing toward the reference regardless, increasing the separation between the two vehicles without bound. With the return channel active, as established in Section 5.5, the Leader halts instead.

The distinction can be put as follows: the forward channel decides whether the formation objective is reachable, and the return channel decides how that objective is preserved when conditions degrade. Neither substitutes for the other. A formation with sensing but no feedback still reaches its objective under nominal conditions but has no protection when the Follower fails; a formation with feedback but no sensing has protection against a failure it can no longer avoid, since the Follower is cut off from the reference regardless.

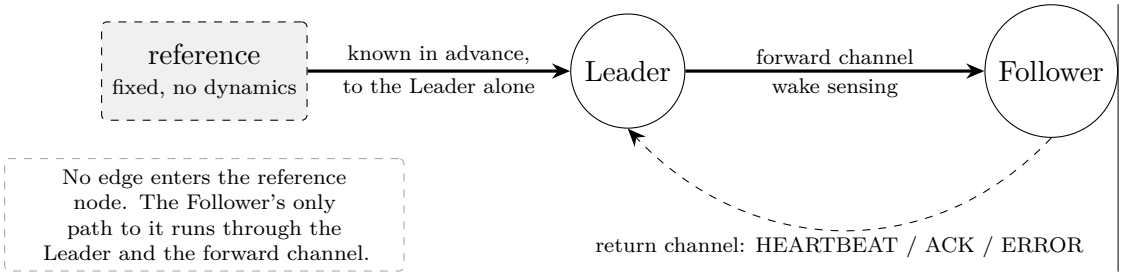


Figure 6.1: Augmented graph of the formation. The reference is a fixed point known to the Leader alone, with no dynamics and no incoming edges, so that the value the formation converges to is imposed in advance rather than determined by initial conditions. The Follower reaches the reference only through the forward channel, which makes that channel structurally indispensable. The return channel carries no reference information and never alters the value converged to; what it affects is the Leader's behaviour when the assumptions behind that convergence stop holding.

Table 6.1: Contribution of each channel to the formation objective.

Condition	Leader's behaviour	Outcome
Both channels available	Advances toward the reference; return channel active but without effect on convergence	Objective reached, and protected against Follower failure
Return channel lost	Advances toward the reference; no evidence of Follower state available	Objective reached under nominal conditions, but with no protection if the Follower fails
Forward sensing lost	Advances toward the reference, unaware that the Follower is no longer connected to it	Objective not reachable: the Follower has no alternative path to the reference

The graph formalisation itself, comprising the augmented topology, its Laplacian, the eigenvalue analysis and the convergence bound that follows from it, is developed in the companion work on the forward channel and is not reproduced here. What this section uses is its structure alone: the reference as a node that no edge enters, the sensing edge that carries from the Leader to the Follower, and the return edge that carries in the opposite direction. That structure is sufficient for the argument made here, which concerns which channel decides what rather than how quickly the formation converges, and the quantitative treatment is referred to rather than duplicated.

6.2 Bidirectional Timing Between Forward Commands and Return Feedback

The two channels are not merely coexistent: they form a closed loop whenever the Leader issues a command and the Follower acknowledges it. This section follows that loop from end to end, accounting for the delay contributed by each stage, and compares the total against the timeout that governs how long the Leader waits.

The loop begins on the Leader. A command is not transmitted as an explicit message but expressed as a change in the Leader's own velocity, which the forward channel conveys through the resulting change in the wake the Follower senses. The Leader marks the moment of transmission by detecting this change in its own commanded velocity, and it is this event, rather than any acknowledgement of

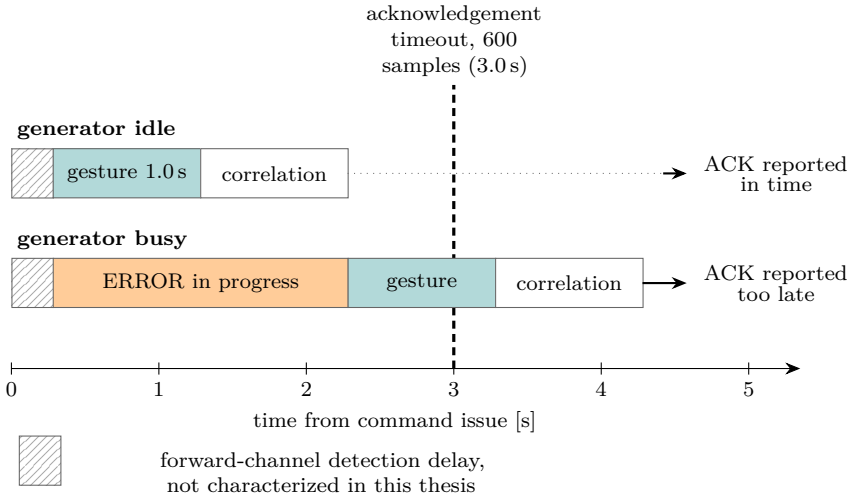


Figure 6.2: Latency of a complete command-acknowledgement exchange. In the favourable case the generator is idle when the acknowledgement becomes pending; in the unfavourable case a maximum-length ERROR is already being transmitted and must complete first. The correlation delay is structural rather than incidental, since the detector resolves a symbol only once its energy excursion has ended. The gesture and the correlation window each occupy 1.0 s. In the unfavourable case the acknowledgement arrives after the Leader has already stopped waiting for it.

transmission by the Follower, that moves the protocol state machine of Section 5.4 into WAITING ACK and starts its timer.

The second leg of the loop is the shortest of the four. As described in Section 4.5, the forward-channel decoder reports a detection as soon as the norm of the Follower's attitude response crosses its adaptive threshold, and the wake reaches the Follower at the speed of the flow rather than at the speed of a control loop. The interval between the Leader changing its velocity and the Follower marking `ack_pending` is therefore short compared with the gesture duration and the correlation delay that follow it, and is neglected in the budget below. The loop's latency is consequently dominated by the three contributions that remain, all of which are established in this thesis.

Once the Follower has marked an acknowledgement as pending, the delay before the corresponding gesture begins is determined entirely by the arbitration logic of Section 4.5. If the generator is idle, transmission starts on the following control step. If it is busy, the acknowledgement waits, and in the worst case waits for the longest symbol in the alphabet to complete, adding up to 2.0 s. The gesture itself then occupies a further 1.0 s, the nominal ACK duration.

Detection on the Leader's side adds a delay of its own, and one that is structural rather than incidental. As established in Section 5.8, the detector classifies a symbol

only after its energy excursion has ended, and that excursion is itself wider than the gesture that produced it. A gesture beginning at a given instant is therefore reported to the state machine roughly one correlation window later, on the order of 1.0 s after transmission began.

Summing these contributions gives the loop's total latency. In the favourable case, where the generator is idle when the acknowledgement becomes pending, the interval between the Leader issuing a command and the state machine leaving WAITING ACK is dominated by the gesture duration and the correlation delay, and amounts to roughly 2.0 s plus the forward-channel detection delay still to be determined. In the unfavourable case, where an ERROR is already being transmitted when the command arrives, a further 2.0 s is added. Against the 3.0 s timeout established in Section 5.4, the favourable case leaves margin, but the unfavourable case does not: an acknowledgement delayed behind a maximum-length symbol will arrive after the Leader has already stopped waiting for it. Extending the timeout is not the appropriate response, because the loss of the acknowledgement is not the consequence that matters in this case. A missed acknowledgement produces no protocol-level reaction, as established in Section 5.4, so an ACK that arrives after the timer has expired costs nothing by arriving late. What it does cost follows from the circumstance that delayed it. The acknowledgement was displaced by an ERROR, which the arbitration of Section 4.5 serves first, and by the time the acknowledgement reaches the Leader that ERROR has already been detected and the reaction stage of Section 5.5 has entered its recovery condition. An acknowledgement arriving in that condition terminates the recovery immediately, on the reasoning given in that section that a Follower able to acknowledge is a Follower that is receiving and responding. That reasoning does not hold for this acknowledgement. It was raised before the anomaly was reported and says nothing about the Follower's state after it, so it terminates a recovery interval on evidence that predates the fault the recovery exists to accommodate.

The correction belongs on the transmitting side and follows the asymmetry that Section 4.5 already establishes between the two event-driven conditions. Latching is justified for the anomaly condition, whose loss leaves the Leader unaware of a fault, and is not justified for the acknowledgement, whose loss is absorbed without consequence. The pending acknowledgement is therefore cleared rather than held when an ERROR preempts it, so that no acknowledgement raised before an anomaly can be delivered after it. The cost is one acknowledgement not transmitted, in exactly the circumstance in which the design already treats acknowledgements as expendable; the benefit is that the recovery interval retains the duration it was given.

6.3 Cross-Channel Interference Analysis

Two channels sharing the same physical medium invite the question of whether they interfere. This section argues that in nominal formation flight they do not, because they occupy different axes of motion, and then examines the one circumstance in which that separation does not hold.

The forward channel conveys information through the Leader's horizontal motion. A change in commanded velocity alters the wake the Follower senses, and it is that alteration the Follower detects. The return channel conveys information through the Follower's vertical motion, as an oscillation superimposed on its altitude reference. The two are carried by orthogonal components of the same physical phenomenon, and the quantities from which each is recovered are correspondingly distinct: the Follower reads the Leader's wake through its own attitude response, while the Leader reads the Follower's gesture through its vertical velocity estimate. A gesture transmitted by the Follower does not alter the Leader's horizontal velocity, and a command issued by the Leader does not produce a 2.5 Hz component in its own vertical velocity estimate. Neither channel writes into the quantity the other reads.

Frequency separation reinforces this. The gesture occupies a narrow band around $f_0 = 2.5\,\mathrm{Hz}$, chosen in Section 3.3 to sit above the bandwidth of ordinary formation-keeping dynamics. Forward-channel commands are expressed as changes in commanded velocity, which are, by comparison, slow. Even if some coupling between the axes existed, the matched filter of Section 5.1 would reject a disturbance at those frequencies, since it correlates against a reference containing no component outside its own narrow band.

The exception is the role-swap manoeuvre itself. In nominal formation the Follower flies below the Leader and the wake travels downward from Leader to Follower, which is the geometry the forward channel depends on. During a role-swap the vertical ordering is inverted, and for the duration of the manoeuvre the Follower is above, which is what makes the return channel possible in the first place. The wake always travels downward, so the channel that functions is determined by which vehicle is underneath, and the two are therefore never available simultaneously. This is not interference in the sense of one channel corrupting the other. It is mutual exclusion: while the Follower is transmitting from above, the forward channel is not merely degraded but absent, and the Follower is receiving nothing.

What makes this acceptable is that the interruption is bounded and self-terminating. A role-swap is a brief excursion of fixed duration, after which the formation returns to its nominal configuration and the forward channel resumes. The Follower is therefore cut off from the Leader's commands only for the length of a reply, and the sensing path that connects it to the formation objective, described in Section 6.1, is restored after every exchange.Whether a command issued during this window is lost depends on what a command is, and in this architecture it is

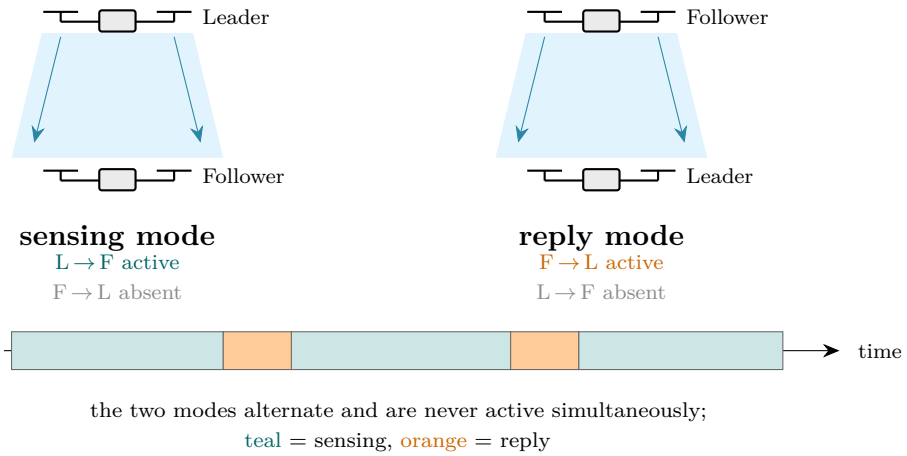


Figure 6.3: The two operating modes. The forward channel travels on the horizontal motion of the Leader and the return channel on the vertical motion of the Follower, so in nominal flight they do not interfere. The role-swap manoeuvre is the exception, and the relationship is not interference but mutual exclusion: while the Follower transmits from above, the forward channel is not degraded but absent. The interruption is bounded and self-terminating, and sensing resumes after every exchange.

not a discrete message. As established in Section 6.2, the Leader does not transmit an instruction but changes its own velocity, and the forward channel conveys that change through the wake the Follower senses. A velocity change is a persistent state rather than an event: it does not occur once and pass, it holds until the Leader changes it again. A Follower that was not sensing while the change was made therefore does not miss the command, it observes it later, discovering the new state at the moment sensing resumes rather than at the moment it was established. What is lost is not the instruction but the interval during which the Follower could have been acting on it, and the consequence of that interval is a growth in formation error rather than a failure to receive.

6.4 Coordination Between Sensing Updates and Gesture Transmission

Section 6.3 established that the two channels alternate rather than coexist, and that a role-swap suspends forward sensing for its duration. This section considers what follows from that alternation for the system's operation: how the formation moves between the two configurations, what the Follower does about sensing it cannot perform while transmitting, and what constrains how often a role-swap may

occur.

The formation operates in one of two configurations at any instant. In the sensing configuration, the Follower flies below the Leader at the nominal separation, aligned so as to remain within the region where the Leader's wake is measurable. In the reply configuration, produced by the role-swap of Section 3.1, the Follower is above the Leader and the separation is reduced. The formation's objective is defined with respect to the sensing configuration alone, since the relative geometry that the Follower controls, and the alignment it maintains with respect to the Leader, are meaningful only while it is positioned below. The reply configuration is not an alternative operating point but a bounded departure from the operating point, after which the formation returns.

Transitions between the two are announced by the Leader. The Follower in the sensing configuration is engaged in pursuit, regulating its geometry continuously against the vehicle above it, and a Follower that climbed without warning would be abandoning that pursuit while the Leader continued advancing toward its reference. The Leader therefore announces the opening of the reply configuration through its own wake, transmitting a sinusoidal modulation that the Follower detects through the forward channel; the Follower responds by interrupting its pursuit and holding position, after which the two vehicles exchange vertical ordering. The announcement is thus not a courtesy but the mechanism that arrests the pursuit loop, and it is required even in a formation of two vehicles. Nothing in this arrangement requires a return path, since the announcement travels from the Leader to the Follower along the forward channel that already exists.

While in the reply configuration, the Follower is not sensing. The forward channel provides no updates for the duration of the manoeuvre, and the Follower's estimate of the Leader's relative position is consequently frozen at whatever it was when the role-swap began. This does not degrade the estimate so much as suspend its refresh: the Follower does not receive incorrect information about the Leader, it receives none. The consequence is that any motion of the Leader during this window goes unobserved, and is discovered only once sensing resumes. The magnitude of the resulting deviation follows from the geometry of the situation rather than from any property of the channel. The Follower holds its last valid estimate while the Leader continues to advance, so the separation between the estimate and the truth grows at the Leader's own speed relative to the Follower and accumulates over the duration of the window: to first order the deviation is the product of that relative speed and the window length T_w, and it is bounded rather than divergent because the window is bounded. Two properties keep it small in practice. The first is that the formation is in the reply configuration for a small fraction of its flight, since a heartbeat is attempted only after 5.0 s of channel silence and the window itself is short by comparison. The second is that the deviation is recovered rather than retained: the pursuit loop resumes with a

correct observation of the Leader's position and closes the gap, so the reply window contributes a transient excursion to the formation error rather than a permanent offset.

The duration of the manoeuvre itself, as distinct from the transmission it makes possible, follows from the two altitude loops identified in Section 7.1. Entering the reply configuration and leaving it are each a commanded change of altitude, so the time either takes is the settling time of the loop that executes it.For a second-order response with damping ratio ζ and natural frequency ω_n, settling to within two per cent takes approximately $4/(\zeta\omega_n)$, which for the Follower's loop at 3.49 Hz amounts to 0.20 s and for the Leader's at 0.5 Hz to 1.41 s. The two vehicles move concurrently, so each transition is governed by the slower of them, and the manoeuvre window comprises two such transitions:

$$T_m \;\approx\; 2\,\frac{4}{\zeta\,\omega_{n,L}} \;\approx\; 2.8\,\mathrm{s}. \tag{6.1}$$

The asymmetry between the two figures is the substantive part of this result. The Follower reaches its commanded altitude seven times faster than the Leader does, so the manoeuvre is limited throughout by the vehicle that is merely making room, not by the vehicle that has something to transmit. The estimate is a lower bound, since it treats both loops as linear and therefore neglects any thrust or rate saturation encountered over the actual travel; what it establishes is the order of the quantity and, more usefully, which of the two loops determines it.

The value is consistent with the manoeuvre observed in the full model, so the estimate is used directly in the cycle budget of Section 8.4.

This suspension is what bounds how frequently symbols may be transmitted, independently of the guard interval discussed in Section 4.4. That constraint was concerned with whether two symbols would be correctly distinguished at the receiver. The constraint here is different in kind: even perfectly distinguished symbols cost sensing time, and a Follower that transmitted continuously would spend most of its flight above the Leader, unable to track it at all. The liveness mechanism of Section 4.5 reflects this directly, since the heartbeat interval of 5.0 s is long relative to both the gesture durations and the manoeuvre itself, keeping the formation in the sensing configuration for the large majority of the time and treating the reply configuration as the exception it is intended to be.

6.5 System-Level Behavior During Nominal and Error Conditions

The preceding sections examined the two channels separately and then their interaction. This section traces the behaviour of the complete system through four

scenarios: three in which the return channel determines what the Leader does, and a fourth in which it determines nothing at all. What the formation achieves in each is what distinguishes them.

Under nominal conditions both configurations alternate as described in Section 6.4. The Leader advances toward the reference it alone knows, the Follower tracks it through wake sensing and maintains the desired relative geometry, and heartbeats are transmitted whenever the channel has been silent long enough, interleaved with acknowledgements when commands are issued. The Leader's protocol state machine cycles between IDLE and WAITING ACK without ever entering an error condition, its liveness timer is reset before it can expire, and no velocity override is applied. The formation converges toward the reference, and the return channel, though continuously active, has no effect on that convergence.

The second scenario is a fault detected by the Follower. Its anomaly monitor flags the condition, the arbitration logic of Section 4.5 gives the resulting ERROR priority over anything else pending, and the symbol is transmitted at the next opportunity. On detection, the Leader halts and holds position for the recovery interval established in Section 5.5. The significance of this at system level is what it prevents rather than what it accomplishes: a Follower in a faulted state may no longer be tracking, and a Leader that continued advancing would increase the separation between the two vehicles without bound. Halting converts an unbounded divergence into a bounded one. The formation does not progress toward the reference during this interval, but it also does not come apart, and if the Follower recovers and resumes acknowledging commands the Leader returns to normal operation immediately rather than waiting out the full interval.

The third scenario is the loss of the return channel itself, whether because gestures are no longer transmitted or no longer detected. From the Leader's perspective these are indistinguishable, and both manifest as the absence of heartbeats. Once the liveness timer expires the Leader halts indefinitely, with no timeout and no mechanism for resuming other than the arrival of a heartbeat. This is the most conservative of the three responses, and deliberately so: unlike a reported error, silence carries no information about whether the Follower is recoverable, and there is no basis on which the Leader could conclude that advancing has become safe.

The fourth case stands apart from the preceding three, and is included here precisely because it falls outside what the return channel can address. If forward sensing fails, the Follower loses its only path to the reference, as established in Section 6.1. No behaviour of the return channel compensates for this. The Follower may continue transmitting heartbeats correctly, and the Leader may continue to regard it as active, while the formation objective has already become unreachable. This is the asymmetry between the two channels stated in its sharpest form: the return channel protects the formation against failures of the Follower, but not against failures of the path that makes the formation possible.

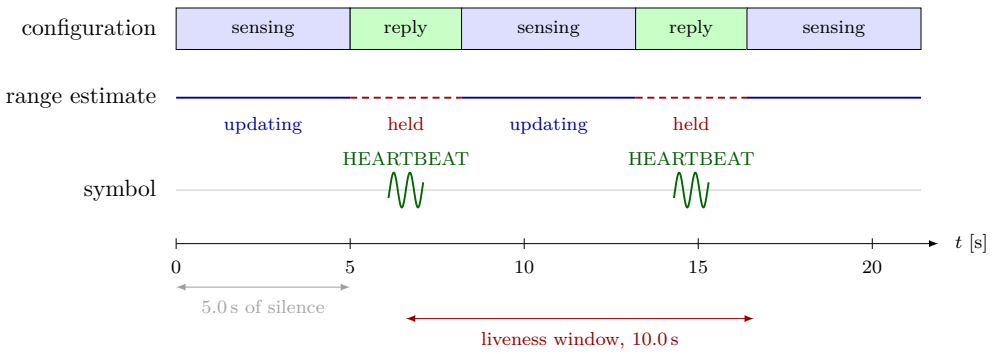


Figure 6.4: One period of the sensing/reply cycle. The formation spends most of its time in the sensing configuration, where the Follower tracks the Leader and updates its range estimate; a heartbeat is initiated after 5.0 s of channel silence, and for the length of the reply window the vertical ordering is inverted, so the Follower cannot sense the Leader and holds its last valid estimate rather than acquiring an incorrect one. Because the interval between consecutive heartbeats is half the window within which the Leader expects one, a single lost symbol does not clear the liveness flag.

Table 6.2: System-level behaviour across the four scenarios.

Scenario	Leader's response	Objective reachable	Error bounded
Nominal	No override; state machine cycles between IDLE and WAITING ACK	yes	yes
Follower fault	Halt for the bounded recovery interval, cut short by an acknowledgement	yes, after recovery	yes
Return channel lost	Halt indefinitely once the liveness timer expires; resumes only on a detected heartbeat	suspended	yes
Forward sensing lost	No response: heartbeats may continue and the Follower is still regarded as active	no	no

Chapter 7

Simulation-Based Validation of the Return Channel

7.1 Simulation Setup and Measurement Pipeline

The validation reported in this chapter is carried out in simulation, against a model of the channel built from the anemometric measurements of Section 3.2. The distinction between what is measured and what is modelled matters for interpreting the results that follow, and this section states it explicitly before presenting them.

The measured component is the channel itself. The lookup table derived from the measurement campaign supplies the disturbance velocity produced by the Follower's wake at any point in the operating region, and it is this table, rather than an analytical approximation, that the simulation evaluates at every control step. The conversion from that velocity to the force acting on the Leader follows the formulation used in the flight model. The table returns a single velocity for the current relative position, and the dynamic pressure this velocity carries acts over a fixed reference area with a fixed drag coefficient. The reference area is that of the airframe taken as a whole and enters the expression once, so the force reaching the Leader depends on the geometry of the formation only through the velocity the table returns, and not through any additional spatial weighting.

No moment accompanies that force. The offsets over which the flight model distributes it are symmetric about the airframe centre and cancel in pairs, so the wake acts on the Leader as a purely vertical disturbance and leaves the attitude axes unexcited. This is a consequence of the geometry rather than an approximation, and it is what allows the entire return channel to be treated on the vertical axis alone.

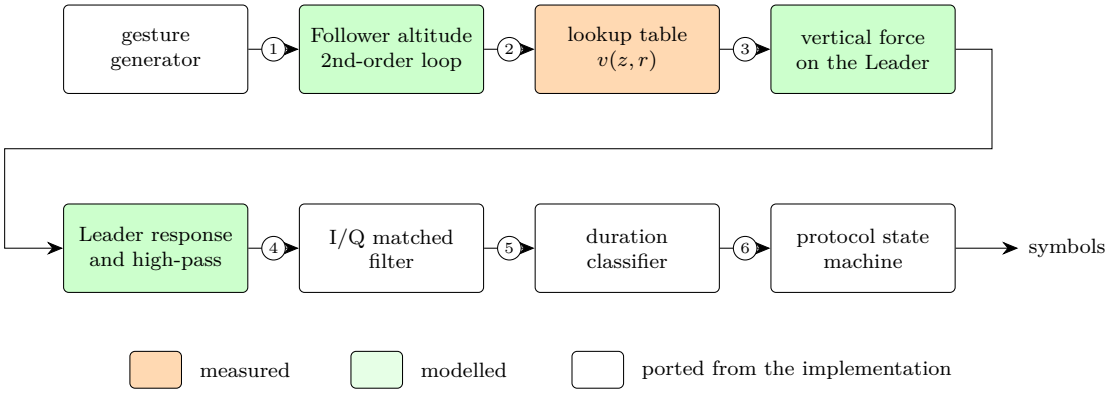
The modelled components are the two vehicles. The Follower's altitude response

and the Leader's response to the disturbance are each represented by a closed-loop second-order system, so that a commanded reference is tracked with finite bandwidth rather than instantaneously. Neither representation extends beyond the vertical axis, since the disturbance carries no attitude component for an attitude model to respond to. The Follower's loop is parameterised at 3.49 Hz with damping 0.9, the bandwidth being that which reproduces the executed excursion measured in Section 7.2; the Leader's is parameterised at 0.5 Hz with the same damping. Both are equivalent representations identified from the closed-loop behaviour rather than transcriptions of the firmware controllers, whose gains are internal to the platform. The signal chain downstream of these dynamics is not modelled but ported directly: the gesture generator, the high-pass stage, the matched filter, the classifier and the protocol state machine are transcriptions of the implementation described in Chapters 4 and 5, operating on the same sample rate and the same parameter values.

The measurement pipeline follows the signal through this chain, recording the commanded altitude offset, the altitude actually achieved, the vertical velocity induced at the Leader, the same quantity after high-pass filtering, the matched-filter energy, and the sequence of symbols the classifier reports. Comparing the first and last of these gives the end-to-end result; comparing intermediate pairs isolates where a discrepancy originates, which is how the attenuation reported in Section 7.2 and the widening reported in Section 7.3 are separated from one another.

The experiments combine two kinds of variation. The physical configuration is swept deterministically: the role-swap separation, the radial alignment between the two vehicles, the spacing between consecutive symbols, and the sequence of symbols transmitted. At each configuration the run is then repeated over independent realisations of the measurement noise on dz, whose standard deviation is the residual of the Leader's own vertical velocity estimate at hover. Repetition is therefore informative: the channel model is deterministic, but the signal the receiver observes is not, and the classification boundaries of Section 3.7 are crossed by a margin that the noise can erode. The standard deviation adopted is an estimate rather than a quantity identified from a hover log, and every rate reported in this chapter scales with it.

Results are accordingly reported as functions of the swept quantity, with detection and classification rates aggregated over the noise realisations at each point. This is what allows the operating window of Section 7.5 to be stated both as a range in centimetres and as a rate, and it is what makes the two edges of that window comparable: each is a point at which the observed peak width approaches a fixed boundary, and the rate at which it is crossed is a property of the noise rather than of the geometry alone. Figure 7.1 shows the chain together with the points at which each quantity is recorded.



Recording points: ① commanded offset, ② achieved altitude, ③ induced vertical velocity, ④ filtered signal, ⑤ matched-filter energy, ⑥ reported symbols.

Figure 7.1: Structure of the validation pipeline. The channel itself is not approximated analytically but evaluated from the lookup table derived from the anemometric campaign; the two vehicles are represented by closed-loop second-order responses; the signal chain downstream is transcribed directly from the implementation rather than re-modelled. Comparing the first and last recording points gives the end-to-end result, while comparing intermediate pairs isolates where a discrepancy originates.

7.2 Validation of Gesture Generation

The first property to verify is the one on which everything downstream depends: that the gesture actually transmitted by the Follower is the gesture specified in Chapter 4. Two distinct steps separate the specification from the transmission. The generator must produce the tapered sinusoid described in Section 4.1, and the Follower's altitude controller must then track it. This section examines both, in the absence of any disturbance from the Leader, so that the generator and the tracking loop can be assessed independently of the channel.

The commanded offset reproduces the specification exactly, which is expected given that it is computed in closed form at every control step rather than filtered or interpolated. The taper reaches full amplitude over the first 20 samples, holds unit gain through the body of the symbol, and returns to zero over the final 20, so that the commanded trajectory is continuous at both ends of every symbol and no discontinuity is ever presented to the controller. All three symbols share this envelope, differing only in the length of the constant-gain section between the two ramps.

The altitude the Follower achieves, however, is not the altitude it is asked to achieve. The commanded oscillation sits at $f_0 = 2.5\,\mathrm{Hz}$, close enough to the

bandwidth of the altitude loop that the reference is transmitted with attenuation and phase lag rather than tracked: the peak excursion realized is 8.7 cm against the commanded $A = 0.12\,\mathrm{m}$, a reduction of just under a third. The lag is 69° at f_0, equivalently 77 ms or 15 samples, and both figures follow from the same loop bandwidth quoted in Section 7.1. The altitude loop is second order, so its response at the gesture frequency is fixed by the natural frequency f_n and the damping ratio ζ alone; writing $r = f_0/f_n$,

$$|H(r)| = \frac{1}{\sqrt{(1-r^2)^2 + (2\zeta r)^2}}, \qquad \angle H(r) = -\arctan\!\left(\frac{2\zeta r}{1-r^2}\right). \tag{7.1}$$

With $f_n = 3.49\,\mathrm{Hz}$ and $\zeta = 0.9$ the ratio is $r = 0.72$, giving $|H| = 0.726$ and a lag of 69.3°. The measured values, a ratio of 0.725 between executed and commanded amplitude at a lag of 69°, agree to three significant digits. The agreement is a check in both directions: it confirms the loop parameters identified in Section 7.1, and it confirms that the attenuation observed here is the ordinary response of a second-order system rather than an artefact of the simulation. Figure 7.2 shows the two traces for a single ACK symbol. Neither effect compromises the channel, but the first has a direct consequence on the link budget: the disturbance the Leader receives is governed by the excursion actually performed, not by the amplitude commanded, and the gesture reaching the Leader is therefore weaker than the nominal value used when the amplitude was first sized in Section 3.3.

The phase lag has no equivalent consequence. As established in Section 3.5, the quadrature structure of the matched filter makes detection independent of the phase at which a gesture arrives, and a constant lag between commanded and executed altitude is indistinguishable, at the receiver, from a gesture that began slightly later. It shifts the detection instant marginally and leaves the classification, which depends on duration rather than phase, unaffected.

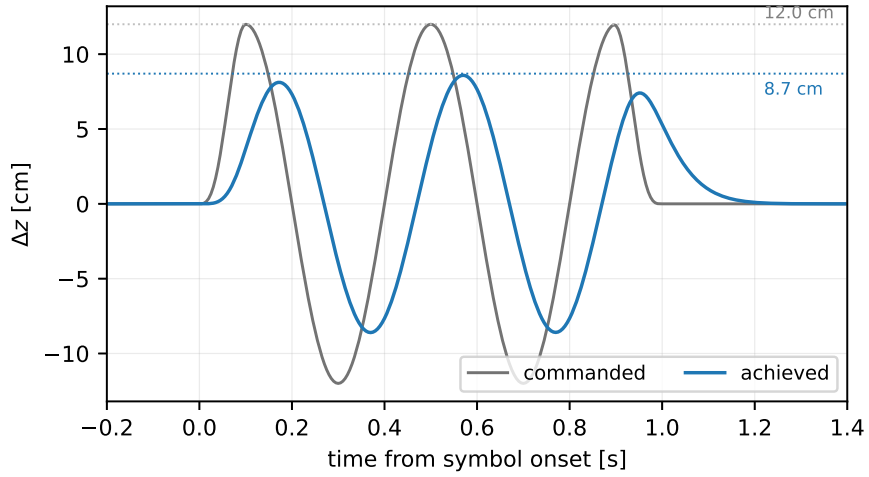


Figure 7.2: Commanded and executed altitude offset for a single ACK gesture. The altitude loop tracks the reference with attenuation and lag rather than reproducing it: the peak excursion falls from 12.0 cm to 8.7 cm, and the executed waveform trails the commanded one by 69°. The taper survives the loop, so the executed trajectory is still continuous at both ends, and the residual tail after the symbol ends is the loop settling rather than commanded motion.

What this section establishes is therefore narrower than it may appear. The generator behaves as designed, and the tracking loop introduces a predictable attenuation that carries forward into every result in this chapter. It is the achieved excursion, not the commanded one, that propagates through the channel model of Section 3.2 in the sections that follow.

7.3 Matched-Filter Response Under Controlled Conditions

With the transmitted gesture characterized, this section examines what the Leader's detector produces in response to it. Each symbol is transmitted in isolation, at the nominal role-swap separation of 25 cm, with no other traffic on the channel, so that the response can be attributed entirely to the symbol that produced it. The quantity of interest is the energy $E(n)$ defined in Section 5.2, and in particular two of its properties: how high it rises, and for how long it remains above λ.

The second property is the one the classifier acts on, and it behaves as Chapter 3 anticipated. The energy excursion is consistently wider than the gesture that produced it, because the correlation window slides onto the gesture before the gesture ends and off it only after. The widths observed here are 176 samples for

HEARTBEAT, 285 for ACK and 487 for ERROR, against nominal durations of 100, 200 and 400, with a spread over noise realisations of 6 samples or less. All three fall on the correct side of the classification boundaries of 235 and 400 samples, and all three symbols are classified correctly.

The first property is less intuitive and deserves a closer look, because the three symbols do not produce excursions of equal height. One might expect them to: amplitude, frequency and taper are identical across the alphabet, and only duration distinguishes one symbol from another. The energy nonetheless rises to markedly different levels, and the reason lies in the relationship between symbol duration and the length of the correlation template.

The template spans $N = 200$ samples. A symbol shorter than this never occupies the correlation window completely: part of the window is correlated against signal, and the remainder against whatever precedes or follows the gesture. HEARTBEAT, at 100 nominal samples and only 60 at full taper gain, occupies less than a third of the window and reaches roughly a sixth of the energy a continuous sinusoid would produce. ACK, at 200 nominal samples and 160 at full gain, comes close to filling it and reaches around nine tenths. ERROR, at 400 samples, fills the window entirely for a sustained interval and reaches the full value, which no longer symbol could exceed. Figure 7.3 shows the three responses on a common axis.

This behaviour is a consequence of the template length rather than a property of the symbols, and it has a practical implication worth stating. The detection margin is not uniform across the alphabet: HEARTBEAT, the symbol transmitted most frequently, is also the one whose energy sits closest to λ. Quantitatively, the peak energy clears λ by a factor of 2.96 for HEARTBEAT against 14.4 for ACK and 16.3 for ERROR. Any degradation that reduces the received signal therefore affects HEARTBEAT before it affects the other two, which is the mechanism examined in Section 7.5. The ordering is fortunate rather than designed: as argued in Section 5.6, HEARTBEAT is also the symbol whose loss carries the mildest consequence.

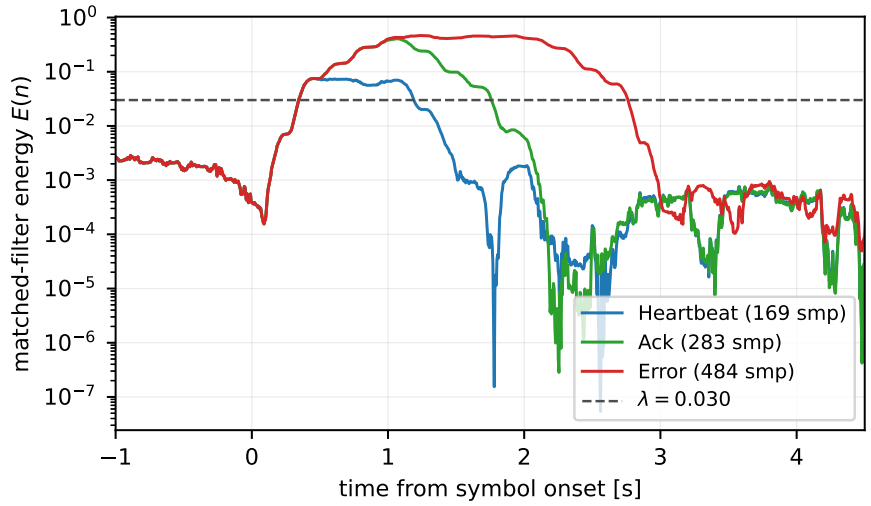


Figure 7.3: Matched-filter energy for the three symbols, aligned at symbol onset and drawn on a logarithmic scale. The peak reached differs by more than a factor of five across the alphabet although amplitude and frequency are identical: HEARTBEAT clears the threshold by 2.96 times, ACK by 14.4 and ERROR by 16.3. The width of each excursion, not its height, is what the classifier reads.

7.4 Decoder Accuracy in Simulated Flight Conditions

Section 7.3 examined symbols transmitted in isolation. This section moves to the condition the channel actually operates in, where symbols follow one another and the receiver must resolve each of them without reference to what came before. Sequences of five symbols were transmitted, mixing all three alphabet entries in randomised orders, over 40 independent noise realisations at each of 16 separations, and the sequence recovered by the classifier was compared against the sequence transmitted.

Every symbol in every sequence was recovered correctly within the operating window of Section 7.5. No symbol was lost, none was duplicated, and none was assigned to the wrong class. This holds regardless of the order in which symbols appear, which is the expected outcome given the structure of the receiver: each detection is resolved from a single energy excursion, and nothing in the classifier carries state from one symbol to the next. Two identical symbols in succession are treated exactly as two different ones, and a long symbol following a short one is unaffected by its predecessor.

Independence between consecutive symbols holds only while they remain far

enough apart, and the point at which it breaks down is worth locating precisely, because it bounds how fast the channel can carry information. Two ACK symbols were transmitted with progressively decreasing separation, measured from the end of the first gesture to the start of the second, and three regimes emerge as that separation closes.

Down to a separation of 200 samples (1.0 s) the two symbols behave as if transmitted alone. Each produces its own energy excursion, the two are divided by an interval in which the energy returns below λ, and both acknowledgements are recovered correctly with widths of 285 samples.

Between 200 and 130 samples the second symbol is not corrupted but lost outright. Its excursion opens within 100 samples of the end of the first, and the debounce interval of Section 5.6 rejects it before it reaches the protocol layer. The receiver reports a single, entirely correct ACK and gives no indication that a second symbol was transmitted at all. This is the cost of that protection stated quantitatively: the same rule that prevents one gesture from being reported twice also suppresses a genuine second symbol arriving too soon after the first.

Below roughly 130 samples the intervening dip disappears and the two excursions merge into one. The class of the resulting detection is erratic rather than monotone in the spacing, because the two carriers at f_0 overlap with a relative phase that repeats every 80 samples and adds or cancels accordingly: a spacing of 120 samples yields a single detection of 491 samples reported as ERROR, 100 samples yields 278 reported as ACK, and coincident gestures yield 213 reported as HEARTBEAT. Two acknowledgements can therefore be read as any of the three symbols in the alphabet, depending on a quantity neither vehicle controls.

The third outcome is the most damaging, and it explains why the failure cannot be left to be handled at protocol level. The first two degrade the channel in ways the protocol can absorb: a second acknowledgement suppressed by the debounce leaves the corresponding command unacknowledged, and the state machine of Section 5.4 eventually times out, which is a delay rather than a fault. A merged detection does not present itself as ambiguous at all. When the phase happens to be favourable it produces a well-formed ERROR that the state machine has no means of recognizing as spurious, triggering the recovery behaviour of Section 5.5 in response to a fault that was never reported. Preventing it is the responsibility of the transmitter, through the guard interval introduced in Section 4.4, and these results give that interval a lower bound: a spacing of 200 samples is required for both symbols to be reported, and the 300 adopted clears it by half as much again.

7.5 Robustness to Disturbances and Geometric Variation

The results of the preceding sections were obtained at a single role-swap geometry. This section varies that geometry and identifies the range over which the channel remains usable. At each separation the run is repeated over independent noise realisations, so that the outcome is a rate rather than a single verdict. The result is a bounded operating window, and although its two edges arise from unrelated mechanisms, both are governed by the same symbol.

The upper limit is set by the symbol with the narrowest margin. As established in Section 7.3, the three symbols do not produce excursions of equal height: HEARTBEAT is shorter than the correlation template and never fills it, so its energy peak sits closest to λ. Any degradation that weakens the received disturbance therefore reaches HEARTBEAT first. This is what the sweep shows. Up to 30 cm every symbol of every run is detected and classified correctly. At 32 cm the HEARTBEAT peak has narrowed from 207 samples at 20 cm to 122, and its spread has widened from 3 samples to 41: the excursion no longer clears the threshold cleanly but breaks into fragments, some of which fall below the minimum-duration check of Section 5.6 and are discarded. Detection falls to 99.2 % and classification to 98.0 %. By 34 cm classification has fallen to 77.0 %, and by 36 cm to 39.4 %.

The lower limit is reached by the same symbol moving in the opposite direction. As the separation closes, every observed width grows, and the HEARTBEAT width is the first to reach the 235-sample boundary above which a liveness indication is read as an acknowledgement. At 18 cm it averages 214 samples with a spread of 5 and every run is still classified correctly; at 17 cm it averages 218 with a spread of 7 and the noise carries a small fraction of the runs across; at 16 cm classification has fallen to 95.0 % and at 15 cm to 89.5 %, while detection remains complete throughout. The symbols are all seen; an increasing share of them are read as the wrong one.

A second constraint sits immediately below this one, and the two are close enough that neither can be regarded as comfortably dominant. During a role-swap the Follower climbs above the Leader and the gesture then oscillates about that position, so the instantaneous separation falls below the nominal one by the excursion actually executed. With the 8.7 cm excursion established in Section 7.2, a nominal separation of 20 cm closes to 11.3 cm at the bottom of every cycle. The Parrot Mambo measures 40 mm in height, so the two airframes together occupy 8.0 cm of that gap and the clearance between them vanishes at a nominal separation of 16.7 cm, below which the vehicles would contact at the lowest point of each cycle and the manoeuvre becomes unexecutable rather than merely unreliable. Since classification has already degraded by 17 cm, this contact limit is never the binding

constraint, but it lies only 1.3 cm below the point at which the channel stops being trustworthy, so the lower edge of the window admits no useful margin of either kind.

Table 7.1 collects the sweep, and Figure 7.4 shows the two rates against separation. Figure 7.5 shows the underlying quantity: the observed peak width of each symbol against the two classification boundaries.

Table 7.1: Detection and classification against role-swap separation, 40 noise realisations of five symbols at each point.

z [cm]	detection	correct class.	false alarms/run
14	100.0 %	76.8 %	0.00
16	100.0 %	95.0 %	0.00
18	100.0 %	100.0 %	0.00
20	100.0 %	100.0 %	0.00
24	100.0 %	100.0 %	0.00
28	100.0 %	100.0 %	0.00
30	100.0 %	100.0 %	0.00
32	99.2 %	98.0 %	0.00
34	93.2 %	77.0 %	0.00
36	75.0 %	39.4 %	0.00

No false alarm was recorded at any separation. This is a direct consequence of the two protections of Section 5.6 acting inside the detector: a noise excursion clearing the threshold is discarded unless it persists for 30 samples, and no second crossing is accepted within 100 samples of an accepted one. Both act before anything reaches the protocol layer, so the state machine is never presented with a symbol that was not transmitted.

The nature of the failure is the same at both edges, and it is the less convenient of the two possibilities. In neither case does the receiver report a weak or uncertain detection: it reports a different symbol, with no indication that anything has gone wrong, and the state machine of Section 5.4 has no means of recognizing the substitution. Above the window an acknowledgement read as a heartbeat leaves the Leader waiting for a reply that has already arrived; below it, a heartbeat read as an acknowledgement can terminate an error recovery that was still in progress. The channel fails by misclassification rather than by loss at both ends, which is the harder of the two failures to detect from outside.

Neither edge is a design value. At the lower edge the classification margin and the contact limit are exhausted within little more than a centimetre of each other, and at the upper edge the HEARTBEAT margin collapses over a comparable

distance, its width spread growing by an order of magnitude between 28 cm and 32 cm. A separation chosen near either boundary retains no tolerance for the variation in executed amplitude and alignment that ordinary flight introduces. The range recommended for operation is accordingly narrower, from 20 to 28 cm: the lower end leaves 3.3 cm of clearance above contact and 28 samples of margin below the classification boundary, and the upper end retains a HEARTBEAT width of 166 samples against that same boundary, with a spread of under 4 samples. Both edges of the wider window depend on quantities established elsewhere in this thesis, the lower on the executed gesture amplitude of Section 7.2 and the upper on the classification boundaries of Section 3.7 and the channel model of Section 3.2, whose measurement campaign spanned 20 to 40 cm and therefore covers the recommended range in full.

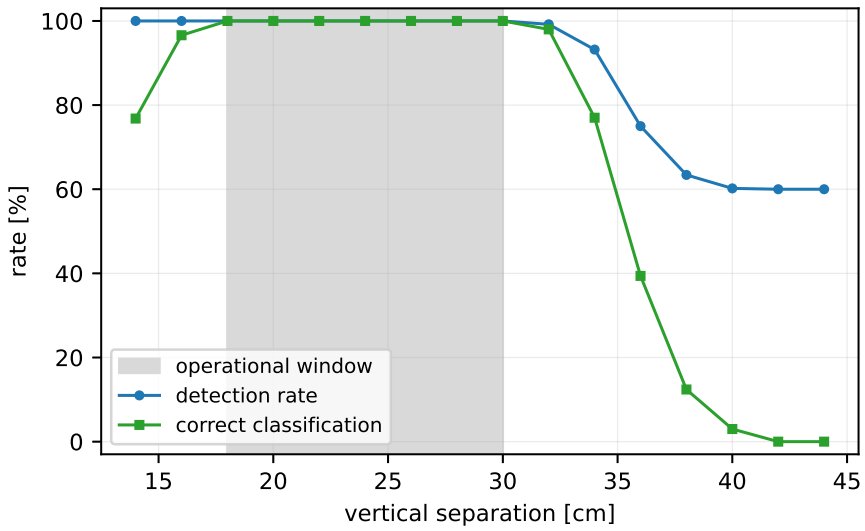


Figure 7.4: Detection and correct-classification rate against role-swap separation. Detection is complete across the shaded window and well beyond its lower edge; classification is the binding quantity at both ends, falling below 18 cm as the HEARTBEAT width grows past the 235-sample boundary and above 30 cm as the HEARTBEAT energy approaches λ and its excursion fragments. The same symbol sets both edges, through opposite mechanisms.

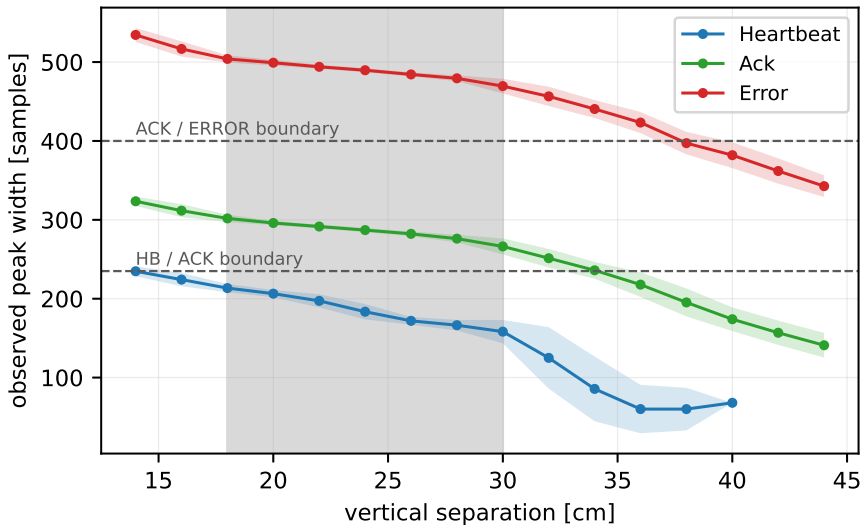


Figure 7.5: Observed peak width of the three symbols against role-swap separation, with the 235- and 400-sample classification boundaries. Shaded bands span one standard deviation over the noise realisations. Every width decreases monotonically with separation; the HEARTBEAT trace is both the lowest and the first to lose its margin at either end of the range, which is what sets both edges of the operating window.

7.6 FSM Behaviour During ACK, ERROR and Liveness Events

The preceding sections examined the receiver up to the point at which a symbol is classified. This section takes the classified symbols as input and exercises the protocol logic of Chapter 5, running the scenarios that the state machine is meant to distinguish. Figure 7.6 shows the four resulting traces on a common time axis.

In the nominal exchange the Leader issues a command, enters WAITING ACK, and the Follower replies. The acknowledgement is reported to the state machine 357 samples (1.78 s) after the command was issued, against a timeout of 600 samples (3.0 s). This total accumulates in three stages, and it is worth separating them. The interval before the Follower begins transmitting is negligible, at five samples. The gesture itself accounts for its nominal 200, and the detector for a further 152: the energy excursion closes once the correlation window has slid far enough off the gesture for the residual overlap to fall below λ, which happens before the window has left it entirely. The two dominant contributions are of comparable size, and the second is structural rather than incidental: it scales with the correlation length, and no choice of protocol timing can reduce it. The margin against the timeout is

accordingly narrower than the nominal symbol durations alone would suggest, a point already anticipated in Section 6.2.

The unfavourable case anticipated in Section 6.2 was exercised directly, with a command issued one sample after an ERROR gesture had begun. The acknowledgement is transmitted only once the generator frees, which happens 700 samples later: the 400 of the ERROR itself and the 300 of the guard interval that follows it. Adding the gesture and the correlation brings the total to 1,054 samples (5.27 s), so the acknowledgement reaches the state machine 2.27 s after the timeout has already returned it to IDLE. The symbol is classified correctly on arrival, with a width of 293 samples: the exchange fails on timing, not on detection. Consistent with Section 5.4 the loss produces no protocol-level reaction, but the margin assumed in Section 6.2 is not merely narrow in this case, it is absent, and restoring it would require a timeout of at least 1,100 samples.

When no reply arrives, the state machine leaves WAITING ACK at exactly the expected sample and returns to IDLE without emitting anything. No output distinguishes this from an exchange that never took place, which is the intended behaviour established in Section 5.4: an unacknowledged command produces no protocol-level reaction at all.

The error path required a correction to the state machine, and the scenario that exposed it is the ordinary one. The anomaly monitor of Section 4.5 responds to the Follower's own attitude and altitude, independently of any command the Leader may have issued, so an ERROR is transmitted whenever a fault is detected rather than in reply to anything. In the original formulation the state machine examined incoming symbols only while waiting for an acknowledgement, with the consequence that an ERROR arriving at any other moment was classified correctly by the detector and then discarded without effect. Since a fault has no reason to coincide with the three-second window following a command, this would have caused most reported errors to be lost.

The correction places error handling outside the state-dependent logic, alongside the liveness mechanism that was already structured this way. An ERROR is now acted upon in any state, and returns the machine to IDLE; an ACK continues to be accepted only while an acknowledgement is outstanding, so that a spontaneous ACK, transmitted without a preceding command, is correctly ignored. This also removes the need for a dedicated ERROR RECEIVED state, since the event can be emitted directly, reducing the machine to three states.

The liveness mechanism required no correction and behaves exactly as specified. Each detected HEARTBEAT resets the timer, and following the last heartbeat the Follower is declared lost after 2,000 samples (10.0 s), matching the specified timeout to within one sample. Because this mechanism sits outside the state graph, it continues to operate while a command is outstanding, so the two failure modes remain independent: a Follower can be declared lost while an acknowledgement

is still being awaited, and the resulting halt takes precedence over the pending exchange.

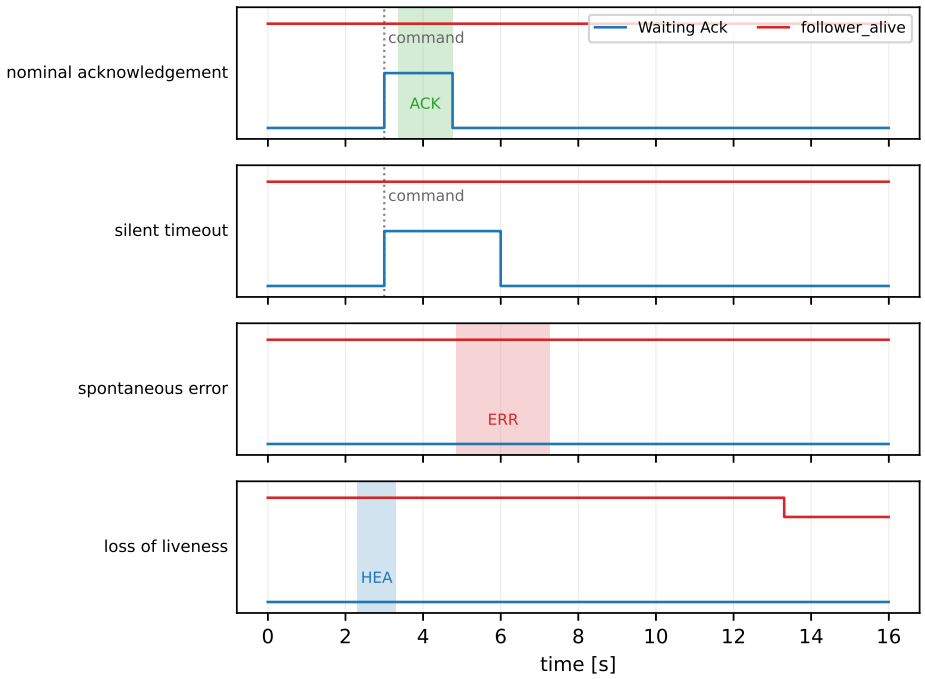


Figure 7.6: State and event traces for the four scenarios on a common time axis. Shaded bands mark detected symbols. In the nominal exchange the machine leaves WAITING ACK on the acknowledgement; in the silent case it leaves at the timeout with no output; the spontaneous error is acted on outside the state graph; and the liveness flag clears 2,000 samples after the last heartbeat, independently of whether a command is outstanding.

Chapter 8

Protocol Scalability in Multi-UAV Formations

8.1 Extending the Return Channel from One Follower to Many

Everything established in the preceding chapters concerns a formation of two vehicles, and most of it carries over to a larger formation without modification. The gesture waveform, the amplitude and frequency chosen in Section 3.3, the matched filter and its threshold, the duration-based classifier and the state machine that consumes its output are all properties of a single transmitter and a single receiver, and none of them acquires a dependence on the number of vehicles present. What does not carry over is the assumption underneath all of them: that whenever a symbol is transmitted, exactly one Follower is in a position to transmit it. This section establishes what becomes of that assumption when more than one Follower is present, and shows that the geometry of the channel, far from being an obstacle to the extension, determines the form the multi-Follower protocol must take.

The return channel does not operate from an arbitrary relative position. As established in Section 3.1, a symbol reaches the Leader only in the reply configuration, in which the Leader descends and the Follower climbs above it so that the Follower's downwash impinges directly on the Leader's airframe. The region from which this is possible is narrow in both directions. Axially, the sweep of Section 7.5 recommends operation between 20 and 28 cm of vertical separation; classification remains complete over the wider interval from 18 to 30 cm, but neither of its edges is a design value, since above it the HEARTBEAT excursion fragments before clearing the threshold and below it the same excursion widens past the 235-sample boundary, so that a liveness indication is read as an acknowledgement. Radially,

the self-similar profile of Section 3.2 has a half-velocity radius of $r_{1/2} = 9.6\,\mathrm{cm}$, so a Follower displaced laterally by more than roughly one body length from the Leader's vertical axis delivers a fraction of the signal the detector was dimensioned for.

Taken together these two bounds define a volume of the order of a Follower's own size, positioned directly above the Leader. Two Followers cannot occupy it at the same time, and the reason is not only that they would be flying within a few centimetres of one another. Even if they could be held there safely, the Leader's vertical velocity estimate would carry the superposition of both gestures rather than either of them. Because both Followers modulate at the same frequency f_0, the matched filter cannot separate the two contributions: it produces a single energy excursion whose extent depends on how the two transmissions happen to overlap in time, and the duration classifier of Section 3.7 assigns that excursion to one symbol class. Two simultaneous HEARTBEAT transmissions are therefore not received as two heartbeats, nor as a detectable collision, but as one symbol of a class determined by their relative timing, which the Leader has no means of recognising as anomalous.

Silent misclassification is a worse failure mode than loss, because the protocol acts on it. An excursion widened past 235 samples is delivered to the state machine of Section 5.4 as an acknowledgement, and one widened past 400 samples as an error, which halts the formation for the full recovery interval of Section 5.5 in response to a fault no vehicle reported. The reply configuration is thus a resource of unit capacity, shared by all Followers, and physically exclusive in a way that no protocol convention can relax.

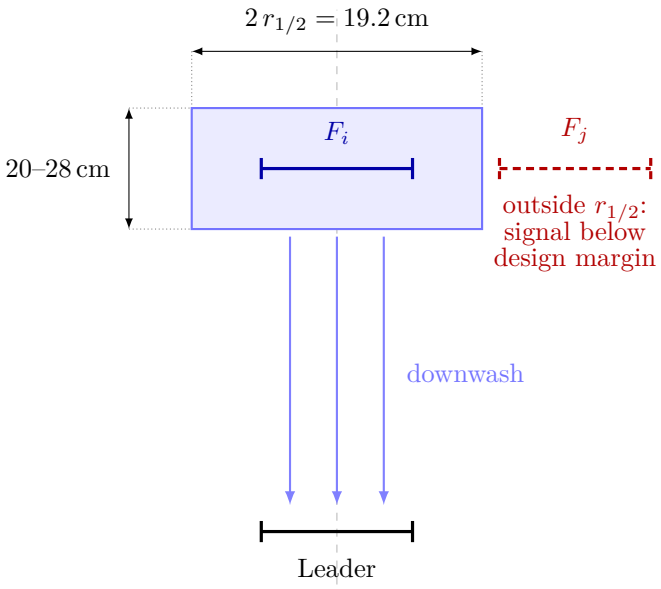


Figure 8.1: The reply configuration as a resource of unit capacity. The shaded region is the set of positions from which a Follower's gesture reaches the Leader with the margin established in Section 7.5: bounded axially by the recommended separation range and radially by the half-velocity radius of Section 3.2. Its extent is of the order of a single airframe, so a second Follower cannot be placed inside it, and one placed outside delivers a fraction of the signal the detector was dimensioned for.

This observation settles, before any comparison of alternatives, a question that would otherwise be open. A frequency-division scheme, in which each Follower is assigned a distinct gesture frequency and the Leader separates the concurrent transmissions spectrally, presupposes that the transmissions are concurrent, and therefore that several Followers occupy the reply volume at once. That presupposition is unavailable here. Allocating distinct gesture frequencies would consume design margin against the constraints of Section 3.3 without purchasing anything, since the resource being contended for is not bandwidth but position.

What the geometry admits instead is a time-division scheme in which the reply configuration is occupied by one Follower at a time, for a bounded interval, and released. The elements required to implement it are already present in the two-vehicle system. The duration of an occupancy follows from quantities already fixed: the longest symbol occupies 400 samples, the matched filter spreads the resulting energy excursion beyond the gesture that produced it, and the guard interval of Section 4.4 keeps the receiver from merging one transmission into the next. The mechanism that opens the reply configuration exists as well, and is described in Section 6.4: the Leader announces the opening through its own wake,

the Follower recognises the announcement and interrupts its pursuit, and the two vehicles exchange vertical positions. The Leader already maintains a liveness timer that expires after 2000 samples (10.0 s) without a detected HEARTBEAT and halts the formation when it does, as established in Sections 5.4 and 5.5. Chapter 6 already establishes the timing rules under which the two channels alternate. The multi-Follower protocol is therefore an extension of the existing architecture rather than a replacement of it, and the remainder of this chapter develops it as such.

A time-division scheme requires that occupancy be granted rather than assumed, and a grant is meaningful only if it is directed at one Follower and acted upon by that Follower alone. The consequence of a misdirected grant is not a lost message but two vehicles manoeuvring into the same small volume, so the mechanism must be built to make that outcome impossible by construction rather than merely improbable. The protocol developed in the following sections therefore places three independent conditions between a grant and a role swap, each necessary and none sufficient on its own.

The first condition is selection: the Leader's announcement must identify the Follower it is offered to, and a Follower interrupts its pursuit only on a match. Section 8.2 shows that this requires no new physical layer, since the announcement is already transmitted and already received by every Follower beneath the Leader; what it lacks is a destination, and the destination can be carried by the announcement's own duration. The second condition is geometric validation: an addressed Follower proceeds only if its own estimated separation from the Leader falls inside the range within which a role swap can be completed. That estimate is already produced onboard by the forward channel, so the check costs nothing and excludes a Follower that has been addressed in error but is not in a position to comply. The third condition is identity confirmation: the Leader establishes which Follower actually entered the reply configuration, rather than assuming that the vehicle it addressed is the vehicle that answered. Section 8.3 develops a means of establishing this from the rotor signature of the replying vehicle, and gives the alternative available on the current hardware.

The default behaviour when any of the three conditions fails is that no vehicle changes configuration. A failure of the addressing mechanism therefore leaves the formation in its sensing configuration, which is the state it is already designed to hold, rather than producing a manoeuvre. This property is what allows the protocol to be extended to a larger formation without weakening the safety argument made for two vehicles.

8.2 Addressing the Announcement

The mechanism by which the reply configuration is entered already exists in the two-vehicle system, and it is worth restating precisely before it is extended, because the reason it exists is not the one a multi-Follower reading would suggest. The Follower in nominal formation is engaged in pursuit: it tracks the Leader through wake sensing and regulates its relative geometry continuously. A Follower that simply climbed whenever it had something to transmit would be abandoning that pursuit without warning, and the Leader, unaware of the intention, would continue advancing toward its reference while the vehicle beneath it left the sensing region. The announcement transmitted by the Leader through its own wake is what prevents this. It tells the Follower that the formation is about to enter the reply configuration, and the Follower responds by interrupting its pursuit and holding position before the exchange of vertical ordering begins. The announcement is therefore not an addressing device that a two-vehicle formation happens not to need; it is the mechanism that arrests the pursuit loop, and it is required with a single Follower exactly as it is with several.

What the existing announcement does not carry is a destination. It is a broadcast statement that the reply configuration is opening, and with a single Follower this is sufficient, since the only vehicle that could interrupt its pursuit is also the only vehicle the Leader could have meant. The property fails as soon as a second Follower is present. The Leader's wake descends onto every vehicle beneath it, the announcement is one transmission rather than one transmission per recipient, and each Follower that detects it performs the same interruption and then converges on the same volume above the descending Leader. The result is the simultaneous occupancy that Section 8.1 established to be physically inadmissible, and it is reached not through a protocol failure but through every vehicle behaving exactly as designed. The extension required is consequently narrow and specific: the announcement must name the Follower it is offered to.

The information to be carried is small, one item drawn from the set of Followers, and the waveform that will carry it is already in the air. The Leader's announcement has, at present, a fixed duration, since duration conveys nothing when a transmission has only one meaning. Allowing that duration to vary turns the announcement into an alphabet of the same kind the return channel already uses, in which a symbol is identified by the interval over which the detector's output remains above threshold. Nothing else about the transmission changes: the frequency, the taper, the amplitude and the propagation path are those of the announcement already being sent, so the detection margin established for it at sensing separation is inherited rather than reopened. What the Follower must additionally do is measure the width of a detection it already produces, which is the operation the Leader performs on every incoming symbol using the classifier of Section 3.7. That the

width is available at all is established in Section 4.5: the forward-channel decoder reports a detection as an interval bounded by two threshold crossings rather than as a single instant. The modification is therefore the transfer of an existing classifier to the opposite vehicle, applied to a signal that vehicle already receives.

Two properties of the arrangement remove ambiguities that would otherwise have to be settled by convention. The first is that address symbols and reply symbols are never decoded by the same receiver: an address travels downward and is read only by Followers, a reply travels upward and is read only by the Leader. The two alphabets may reuse the same nominal durations without any possibility of a symbol being attributed to the wrong direction, so the classification structure validated in Section 7.4 transfers to the forward direction, though its boundaries must be recalibrated for the sensing separation rather than adopted from the reply configuration. The second is that an address is not a command. The Leader's instructions continue to reach the Follower through the Leader's own velocity, as established in Section 6.2; the addressed announcement carries one item of information only, namely which Follower is being offered the configuration that is about to open.

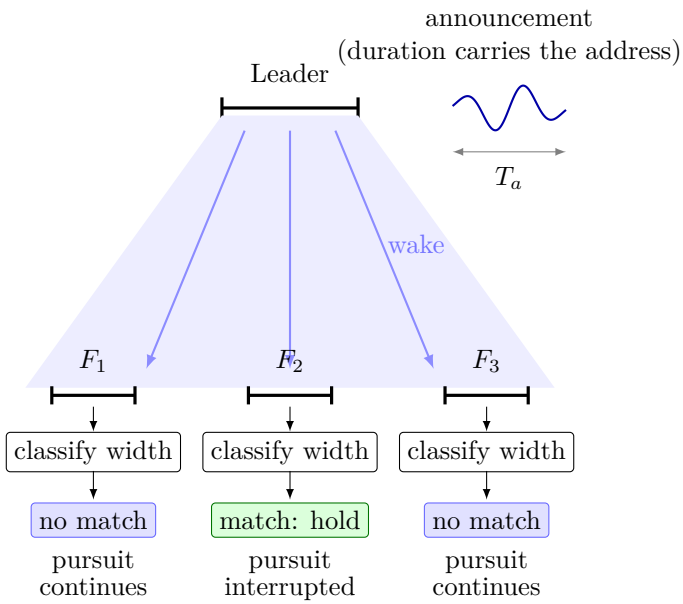


Figure 8.2: An addressed announcement in a three-vehicle formation. The transmission is broadcast by construction, since the wake descends onto every Follower beneath the Leader, and the address is carried by its duration. Each Follower classifies the width of a detection it already produces; only the vehicle whose address matches interrupts its pursuit, while the others continue tracking. Nothing about the transmission other than its duration differs from the announcement already sent in the two-vehicle system.

The number of Followers that can be named follows from the same argument that fixes the symbol alphabet of the return channel, and is a property of the detector rather than of the protocol. Because a detection is reported only once the excursion that produced it has ended, and that excursion is wider than the transmission itself, two announcements are separable only if their nominal durations differ by more than the spread of the excursions they generate. The measurements of Section 7.5 give the scale of that spread in the reply direction, where a HEARTBEAT of nominal duration 100 samples is observed as an excursion of 207 samples at 20 cm. The forward direction is not governed by the same mechanism: the decoder described in Section 4.5 compares an attitude norm against an adaptive threshold rather than correlating against a template of fixed length, so its excursions are not widened by a correlation window and the achievable spacing is set by its debounce interval instead. The addressable set is bounded by the ratio between the longest announcement the grant cycle can accommodate and this minimum spacing, and grows only as the former is extended. Each additional address is therefore paid for in time, in the same currency as the cycle of Section 8.4, and not in bandwidth.

Lengthening the announcement lengthens an interference the two-vehicle system already contains. While the Leader modulates its wake, the centreline velocity from which each Follower reconstructs its separation is being modulated too, so the estimate produced during a transmission varies with the announcement rather than with the relative motion of the two vehicles. Section 6.4 already establishes that the Follower's estimate is suspended rather than corrupted whenever the forward channel is unavailable, and the same treatment applies here: the Follower holds its last valid estimate for the duration of the announcement and resumes updates once it has ended. What changes is only the length of the hold, bounded by the longest address symbol, and it remains short compared with the timescale over which the formation geometry evolves.

The failure modes inherit the asymmetry that motivated the three-condition structure of Section 8.1, and one of the three responses is already implemented. An announcement that fails to cross the detection threshold produces no match on any Follower, hence no interruption of pursuit and no manoeuvre; the formation continues in the configuration it already occupies, and the Leader, receiving nothing, returns to sensing. An announcement misclassified into a neighbouring duration class is the more consequential error, since it names a Follower the Leader did not intend, and it is against this case that the two remaining conditions were introduced. The named vehicle proceeds only if its own estimated separation places it inside the range within which the manoeuvre can be completed, and the Leader establishes which vehicle answered before treating the reply as belonging to the one it addressed, by the means developed in Section 8.3. A misclassified announcement therefore terminates in the same state as a lost one, with the formation returning to sensing and no symbol having been attributed to the wrong vehicle.

8.3 Passive Identification Through Rotor Signatures

The third of the three conditions of Section 8.1 asks the Leader to establish which Follower replied, rather than to assume that the vehicle it addressed is the vehicle that answered. The distinction matters because the two preceding conditions are both evaluated onboard the Followers: the address is classified by each Follower against its own identifier, and the geometric check is performed by each Follower against its own estimate. Neither produces any evidence the Leader can observe. If both were to fail together on the same vehicle, through a misclassified address that happened to name a Follower which was also within range, the Leader would receive a symbol it had no means of attributing to anyone but the vehicle it intended to address, and would act on it accordingly. The purpose of the third condition is to place at least one check on the Leader's own side of the channel.

The physical basis for such a check is that a rotor is not a silent source. A rotor with N_b blades turning at n revolutions per second pulses the flow once per blade passage, producing a periodic component at the blade-passing frequency

$$f_{\mathrm{BPF}} = N_b \, n, \tag{8.1}$$

together with its harmonics. Where this component falls is set by the blade count and the shaft speed alone; the strength of the component, and of the broadband content around it, depends on geometry and material, but its position does not. Two otherwise identical vehicles fitted with propellers of different blade count therefore produce components at different frequencies even when their rotors turn at the same speed, which is what makes the blade count usable as an identifier.

A bench campaign was conducted to establish whether this holds in practice for propellers of the scale involved here. The bench carried a single motor of the 2312E class, rated at 960 KV, on a rigid mount, driven by an electronic speed controller of the E305 family taken from a quadrotor airframe with the original flight controller bypassed, so that the throttle command reached the controller unmodified. A microcontroller board acted as signal generator, producing a pulse width between 1000 and 2000 μs at 50 Hz, which allowed an operating point to be set repeatably and returned to across propellers. Shaft speed was taken either from the electrical frequency of the motor phase or from the position of the shaft component in the acquired spectrum, the two agreeing wherever both were available.

The propellers were produced for the campaign rather than sourced, and this is what makes the comparison meaningful: commercial propellers of different blade count differ in profile, chord and material as well, so the effect of blade count alone cannot be separated from them. Each was generated from a single parametric model with a NACA 4412 section and a pitch-to-diameter ratio of 0.5, and printed by fused deposition on the same machine and in the same material. That last point is what allows the comparison to isolate the blade count: any difference between two spectra is a difference in geometry, and not in how the geometry was made. Two families were produced: one varying the blade count between two, three and four at a fixed diameter of six inches, and one varying the diameter between six, eight and nine inches at two blades. All shared a self-tapping hub with a bore of 5.8 mm, so that any propeller could be fitted to the same shaft without adapters.

Each operating point was recorded acoustically for approximately fifteen seconds at 48 kHz, windowed and transformed. Identification of the shaft component was anchored to the blade count rather than performed by unconstrained peak search, since the blade-passing component must fall at the integer multiple fixed by the number of blades, and requiring that relation resolves the ambiguity that arises when a harmonic is stronger than the fundamental. Recordings were made with a consumer-grade microphone held at a fixed position across propellers. The microphone-to-rotor distance does not enter any of the claims made here: the

position of a blade-passing component is a frequency, and a frequency does not depend on where it is observed from, while the spectra of Figures 8.3 and 8.4 are normalised to their own maximum, so the relative levels within each are unaffected as well. The only quantity the arrangement would not support is an absolute sound-pressure level, which this work does not use.

Figure 8.3 shows the result for the blade-count family at a common throttle command. The shaft component remains where it is across the three propellers while the blade-passing component moves with the blade count, appearing at 139, 208 and 272 Hz for two, three and four blades against a shaft frequency near 69 Hz: in every case the integer multiple predicted by Equation (8.1), to within one per cent. Three blade counts rather than two matter here, because a factor of two between two and four blades admits explanations that a factor of three-halves between two and three does not.

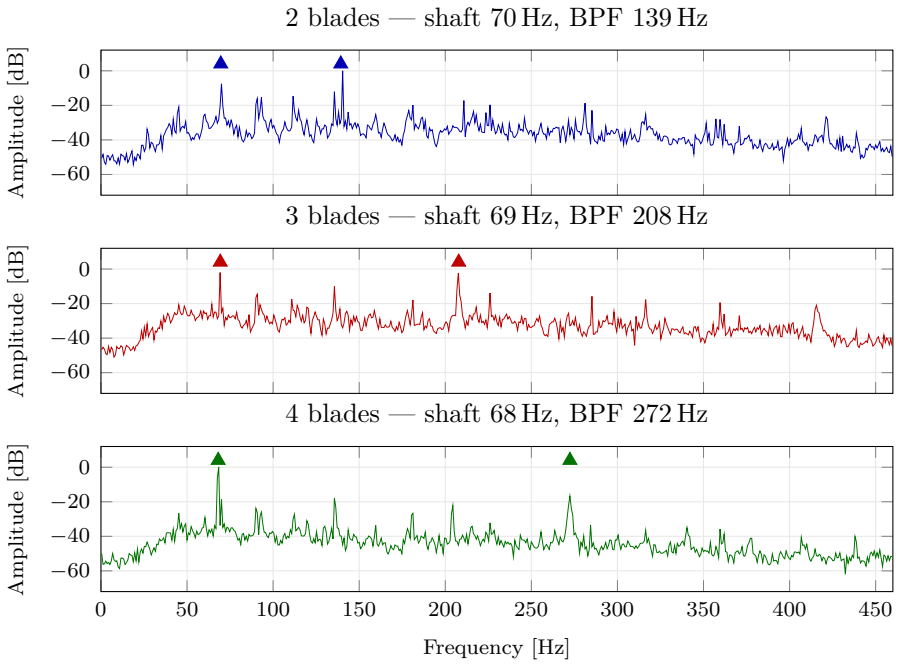


Figure 8.3: Acoustic spectra of the three six-inch propellers at the same throttle command, and therefore at very nearly the same shaft speed. The shaft component stays where it is while the blade-passing component moves with the blade count, appearing at 139, 208 and 272 Hz for two, three and four blades against a shaft frequency of approximately 69 Hz: in each case the predicted integer multiple, to within one per cent. Because the three propellers come from the same parametric model and differ only in blade count, the displacement cannot be attributed to profile, chord or material.

Diameter acts differently, and the difference is worth stating because it bounds what the identifier can be built on. A larger propeller loads the motor more heavily, so at the same throttle command it turns more slowly, and the blade-passing component moves accordingly; it also changes the strength of the harmonics and the broadband content between them. What it does not do is change where the component falls relative to the shaft speed, which remains fixed by the blade count. Diameter therefore alters the operating point and the texture of the spectrum, while blade count alters the position of the line, and only the latter is a usable address.

For two vehicles to be distinguished, two signatures must be separable when both are present. Figure 8.4 shows the spectra of two such superpositions. Both blade-passing components survive and remain individually resolvable, including in the closer of the two cases, where they lie 65 Hz apart. The superpositions were formed numerically from recordings taken separately, since the bench carries one motor and two propellers could not be run on it simultaneously; this is admissible because the transform is linear, but it means the result is a statement about spectral separability and not about how two physically adjacent jets would interact. The nonlinear merging of two nearby wakes is a separate question, and nothing here addresses it.

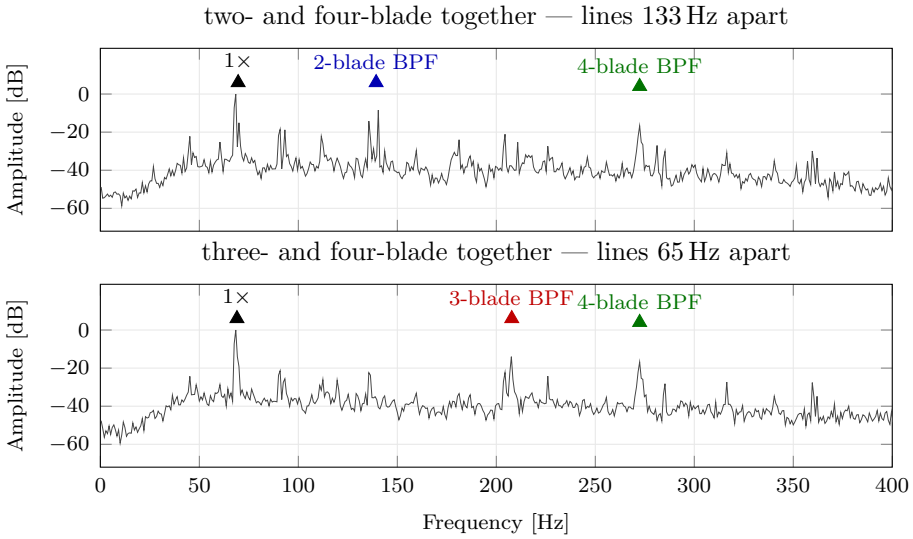


Figure 8.4: Two signatures present at once. Each panel is the spectrum of the superposition of two recordings taken separately at the same throttle command; since the bench carries a single motor, two propellers could not be run simultaneously, and the superposition is formed numerically, which is admissible because the transform is linear. Both blade-passing components survive and remain individually resolvable, including in the closer case where they lie 65 Hz apart. The result is a statement about spectral separability, and not about how two physically adjacent jets would interact.

What this offers the protocol is an identifier carried by the vehicle rather than by the message. A Follower fitted with a propeller of distinct blade count announces its identity continuously, for as long as its rotors are turning, without transmitting anything and without occupying any part of the reply configuration. Two consequences follow. The first is the one the third condition requires: the Leader, receiving a symbol in the reply configuration, can attribute it to the vehicle whose signature accompanies it rather than to the vehicle it addressed, and can discard a reply whose identity does not match. The second is that the silent misclassification described in Section 8.1 becomes detectable. Two Followers transmitting simultaneously produce a single merged energy excursion at f_0, which the classifier cannot recognise as anomalous, but they also produce two distinct signatures at their respective blade-passing frequencies, and the presence of two where one was expected identifies the merged excursion as a collision rather than as a symbol. A collision that can be detected can be discarded and retried; one that cannot is acted upon.

The identification is available only while the replying vehicle is above the Leader. The wake travels downward, so a Follower in the sensing configuration directs

neither its downwash nor the periodicity carried within it toward the vehicle above it, exactly as established in Section 6.3 for the gesture itself. Rotor signatures therefore identify the vehicle that occupies the reply configuration; they do not provide the Leader with continuous evidence about vehicles that are not occupying it, and in particular they do not offer a way to establish liveness without granting a slot. This limitation is structural rather than instrumental, and it is what keeps the timing budget of Section 8.4 proportional to the number of Followers.

A second limitation is instrumental, and it concerns the current hardware rather than the principle. The components measured on the bench lie between roughly 140 and 510 Hz, so recovering them requires a sampling rate above 1 kHz. The decoder of Chapter 5 runs at $F_s = 200$ Hz, whose Nyquist limit of 100 Hz falls below the entire measured band, and logs recorded on the Parrot Mambo confirm what this implies: with a Leader present the amplitude of the Follower's roll residual rises substantially, but the rise is broadband rather than tonal, and no component at any blade-passing frequency appears. There is a further gap which the sampling rate alone does not close: the bench measures acoustic pressure, whereas the quantity available to the Leader is an aerodynamic moment, so the campaign establishes that the source produces a tonal component at the predicted frequency, not that the component is transferred to the moment with usable amplitude. The mechanism described in this section is therefore not observable on the platform used for the rest of this thesis, and its adoption presupposes an inertial sensor of the class routinely fitted to first-person-view airframes, which sample in the kilohertz range and would place the whole measured band within the readable spectrum.

Where that sensor is unavailable, the third condition can still be satisfied, at a cost. The replying Follower may transmit an identity codeword ahead of its symbol, drawn from the same duration-coded alphabet the return channel already uses and classified by the same receiver. This requires no hardware beyond what the two-vehicle system already carries, and it places the check on the Leader's side as required, but it lengthens every occupancy of the reply configuration by the duration of the codeword together with the guard interval of Section 4.4, and it does not make collisions detectable, since two codewords transmitted simultaneously merge exactly as two symbols do. The two mechanisms are therefore not equivalent: the passive signature identifies without cost and detects collisions, while the transmitted codeword identifies at the cost of cycle time and leaves collisions silent.

8.4 The Grant Cycle and Its Timing Budget

The preceding sections establish that the reply configuration is granted to one Follower at a time and that the grant can be addressed and confirmed. This section assembles those elements into a cycle and asks how many Followers it can serve,

which is where the scalability of the protocol is actually decided.

The cycle consists of an addressed announcement, an interruption of pursuit by the named Follower, an exchange of vertical ordering, the transmission of one symbol, a return to the sensing configuration and a resumption of pursuit. Writing T_a for the announcement, T_m for the manoeuvre, T_s for the symbol and T_g for the guard interval that separates one occupancy from the next, and k for the number of grants each Follower receives per liveness window, the admissible size of the formation follows from a constraint the two-vehicle system already imposes. The Leader considers a Follower active only while its liveness timer has been reset within 2000 samples (10.0 s) by a detected HEARTBEAT, and clears the flag otherwise, halting the formation until a heartbeat next arrives, as established in Sections 5.4 and 5.5. In a formation of N Followers the Leader holds N such timers, each of which must be reset within its own window, and the reply configuration serves one Follower at a time:

$$N \;\le\; \left\lfloor \frac{T_{\mathrm{live}}}{k\,T_c} \right\rfloor, \qquad T_c \;=\; T_a + T_m + T_s + T_g. \tag{8.2}$$

All four terms are fixed by the design already established. The symbol that resets a liveness timer is a HEARTBEAT, so $T_s = 100$ samples (0.5 s); grants carrying an acknowledgement or an error report are longer, but they are occasional and do not set the budget. The guard interval is $T_g = 300$ samples (1.5 s) from Section 4.4. The announcement is of the order of a return-channel symbol and is taken as 0.5 s, pending the calibration of Section 8.2. The manoeuvre is the quantity estimated in Section 6.3 from the two altitude loops, and at $T_m \approx 2.8$ s it dominates everything else:

$$T_c \;=\; 0.5 + 2.8 + 0.5 + 1.5 \;=\; 5.3\,\mathrm{s}, \tag{8.3}$$

of which the role-swap accounts for fifty-three per cent, the guard interval for twenty-eight, and the symbol the Leader is actually waiting for for nine.

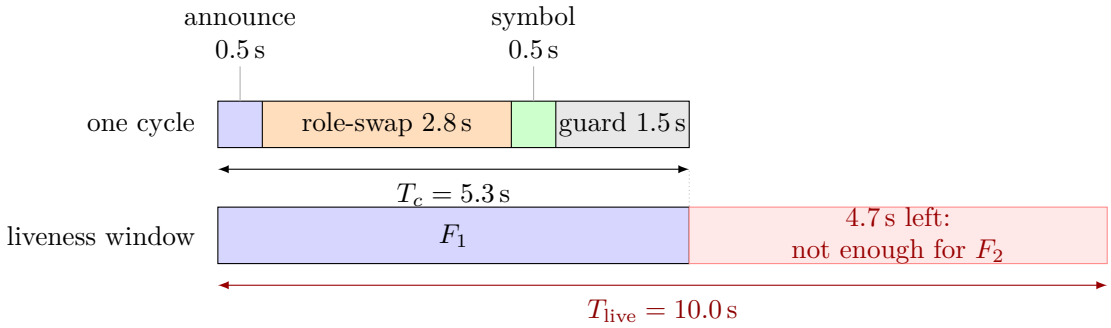


Figure 8.5: The grant cycle against the liveness window. With the manoeuvre duration estimated in Section 6.3, a cycle carrying one heartbeat occupies 5.3 s, of which the role-swap alone accounts for more than half and the symbol the Leader is waiting for accounts for less than a tenth. Two such cycles do not fit inside the window within which every Follower must be heard from, so the liveness budget inherited from the two-vehicle system admits one Follower and no more. What the protocol costs is not channel time but manoeuvre time.

Substituting into Equation (8.2) with $k = 1$ gives $N \leq 1$. The multi-Follower protocol, dimensioned as it stands, admits no second Follower, and it is worth being precise about why, because the reason is not the return channel. The liveness window of 10.0 s was sized in Section 5.4 for a formation of two vehicles, where a Follower enters the reply configuration whenever the Leader opens it and the window need only accommodate one heartbeat with margin. Under a granted scheme that same window must accommodate one full cycle per Follower, and a cycle is more than ten times the symbol it carries.

The bound is therefore a statement about the liveness window rather than about the channel, and it is relieved by enlarging the window rather than by improving detection:

$$T_{\text{live}} \;\geq\; k\,N\,T_c \;=\; 5.3\,kN \quad [\text{s}], \tag{8.4}$$

so that two Followers require 10.7 s and three require 16.0 s at one grant apiece. What is paid for the enlargement is detection latency: a Follower that has failed is recognised as failed only after the window has elapsed, and since the Leader's response to that recognition is to halt, the enlargement extends the interval over which the formation continues to advance past a vehicle that has stopped signalling. Formation size and fault-detection latency are bought from the same budget, in direct proportion.

A second reading of Equation (8.3) points somewhere the protocol cannot reach on its own. More than half the cycle is spent executing a manoeuvre whose duration is set by the bandwidth of the *Leader's* altitude loop, seven times slower than the Follower's, as Section 6.3 establishes. Raising that bandwidth to 1 Hz would

shorten the cycle to 3.9 s and admit two Followers within the window as it stands; at 2 Hz the cycle falls to 3.2 s and three become feasible without enlarging the window at all. The most effective single change available to a designer wishing to scale this protocol is therefore not to the alphabet, the detector or the scheduling discipline, but to a control loop that has nothing to do with communication.

A further property of the cycle deserves to be recorded, because it is a consequence of granting rather than of scheduling. In the two-vehicle system the Follower transmits when it has something to say, and the arbitration of Section 4.5 resolves the case in which more than one condition is pending by a fixed priority order, ERROR before ACK before HEARTBEAT. Under a granted scheme the Follower can transmit only when it holds the slot, so a fault detected between two grants waits for the next one. The latch described in that same section guarantees that the report is not discarded, so the consequence is a delay rather than a loss, but the delay is now bounded by the cycle rather than by the residual duration of a symbol in progress. Error reporting is therefore the element of the protocol most sensitive to the cycle duration, and this is the strongest argument for keeping the cycle short that does not come from the liveness budget.

8.5 Feedback Topology of a Multi-Follower Formation

Section 6.1 described the two channels as edges of a graph containing a reference node, the Leader and the Follower, and used that description to distinguish what the forward channel decides from what the return channel decides. This section extends the same description to N Followers, and asks which of its properties survive the extension. The quantitative treatment of that graph, its Laplacian and the convergence bound that follows from its spectrum, belongs to the companion work on the forward channel and is not reproduced here; what is used below is its structure.

The structure generalises without modification. The reference is known to the Leader and to no other vehicle; each Follower senses the Leader and no Follower senses the reference; and no edge enters the reference node, so nothing any vehicle does alters the value the formation is converging toward. In the sensing configuration the path from the reference to each Follower runs through the Leader and through that Follower's own sensing edge, so every Follower is reachable from the reference as long as it holds that edge, exactly as in the two-vehicle case. The return channel contributes one edge per Follower, directed from that Follower to the Leader, and none of these edges points at the reference. The conclusion drawn in Section 6.1 therefore holds for any N: the return channel does not participate in determining what the formation converges to.

What the return channel changes is not the graph in nominal operation but the graph under fault, and this is the one respect in which the feedback edges act on convergence rather than merely observing it. The reaction stage of Section 5.5 responds to a cleared liveness flag by overriding the Leader's velocity and holding it at zero. In the language of this section, that override removes the edge from the reference to the Leader: the Leader is still the only vehicle that knows where the formation is going, but it has stopped regulating toward it. For the duration of the override the formation is therefore disconnected from the reference, and with no edge entering the graph from outside it the state ceases to evolve toward anything. It holds.

Three properties follow, and together they are the clearest statement of what the return channel contributes. The first is that convergence is suspended rather than degraded: because the reference edge is removed rather than reweighted, the formation does not converge to a different value during an override, it converges to nothing at all, and resumes from the state it held when the edge is restored. The second is that the total settling time of the formation is therefore the nominal settling time plus the accumulated duration of every override, and nothing else; a Follower that reports promptly and recovers promptly costs the formation only the time it was actually faulted. The third is that the two override conditions differ precisely in how the edge is restored. An ERROR removes it for a bounded interval, after which it returns whether or not anything has changed, because the Follower has signalled a fault it can be expected to recover from. A cleared liveness flag removes it with no timeout at all, because a Follower that has stopped signalling has given the Leader no basis on which to decide that advancing has become safe again. In graph terms, the Leader will not reconnect itself to the reference until every vehicle it is responsible for has been heard from.

The second of those properties can be made quantitative, because both of its terms are known. The nominal figure deserves one line of justification, because it is not the time constant of a single loop. The Leader is driven directly by the reference and its error decays as $e_0 e^{-kt}$, but the Follower is the second stage of a cascade forced at its own natural rate, so its error carries a secular factor, $e_F(t) = e_0\,(1 + kt)\,e^{-kt}$. The formation is settled when its slowest member is, so it is the second expression that must be solved against the tolerance band, and at $k = 0.08\,\mathrm{s}^{-1}$ and $e_0 = 0.30\,\mathrm{m}$ this gives $t^\star \approx 55\,\mathrm{s}$ to a band of 2 cm. Writing $T_{\mathrm{ov},i}$ for the duration of the i-th override, the time to settle is then

$$T_{\mathrm{settle}} = t^\star + \sum_i T_{\mathrm{ov},i}. \tag{8.5}$$

Nothing else enters, because the overrides do not alter the value being converged to and the intervals between them are ordinary convergence. Each reported anomaly contributes the fixed recovery interval of 10.0 s from Section 5.5, so a formation

that reports one fault during its approach settles in roughly 65 s rather than 55. A cleared liveness flag is different in kind: it contributes an interval that is not bounded at all, since the reference edge is not restored until a heartbeat arrives, and T_{settle} is then undefined rather than merely longer. That is the intended behaviour and not a deficiency of the expression, since a formation that has lost a Follower is not a formation that should settle.

This is the sense in which the return channel does not make the formation converge faster. It makes convergence interruptible on evidence, which a formation without a return path cannot do: lacking any report from the vehicles behind it, such a Leader remains connected to the reference under all conditions and continues to advance toward it, and the separation between a failed Follower and the formation it belongs to grows without bound.

What does not generalise unchanged is the rule by which the Leader combines the N liveness flags it now holds. Halting whenever any flag is clear preserves the guarantee established for two vehicles, that no Leader advances while a vehicle it is responsible for has stopped signalling, and extends it to every member of the formation; but it also makes the formation's ability to advance conjunctive in N, so that the rate at which the formation is halted grows with its size even if the reliability of each individual link is unchanged. The sweep of Section 7.5 bounds that rate favourably. Across the recommended range no symbol was misdetected or misclassified in any realisation, so the per-symbol failure probability is not merely small but unobserved, and the appropriate statement is an upper bound rather than an estimate: with no failure in n transmissions it lies below $3/n$ at 95 % confidence, which for the sweep conducted here places it in the low single digits of a per cent. Writing p for that bound, the probability that a window elapses without a detected HEARTBEAT is at most p^k for a given Follower and at most $1 - (1 - p^k)^N$ for the formation.

Two consequences follow, and they pull against each other. A single grant per window already yields a spurious-halt rate of a few per cent per window, which grows only linearly with N and is therefore not what limits the formation. Granting twice per window squares the bound and reduces the rate to a fraction of a per cent, but does so by halving the admissible N in Equation (8.2). Reliability and formation size are purchased from the same budget, and the choice between them is not resolved by the channel measurements, which bound p but say nothing about how many spurious halts an application can tolerate.

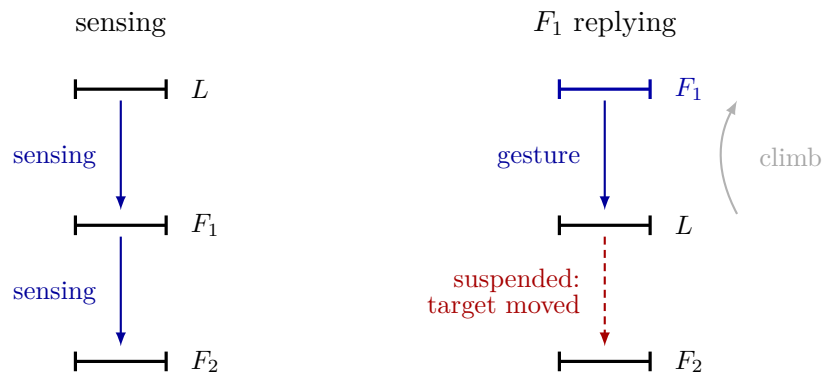
The asymmetry between the two directions is what makes the conjunctive rule acceptable despite this cost. A Follower that has genuinely failed and one whose gestures merely went undetected are indistinguishable from the Leader's position, and were already treated as equivalent in Section 5.7 on the grounds that halting is the appropriate response in both cases. Enlarging the formation does not change that argument; it changes only how often it is invoked. The alternative, in which the

Leader advances while some subset of its Followers is unaccounted for, reintroduces precisely the failure the return channel was built to prevent, and does so for the vehicles least able to report it.

8.6 Chains and Hierarchies

The formations considered so far are stars: one Leader, several Followers, each sensing the same vehicle. The alternative arrangement, in which a Follower serves as the Leader of the vehicle behind it, distributes the sensing load and is the natural response to the bound of Section 8.4, since it caps the number of Followers any one vehicle must serve. It also introduces a coupling that the star does not have, and this section examines it.

In a chain, an intermediate vehicle occupies both roles: it senses the wake of the vehicle ahead and is sensed by the vehicle behind. When such a vehicle is granted the reply configuration, it climbs above the vehicle it reports to, and in doing so it leaves the position from which the vehicle behind it was sensing. A reply is therefore not a local event. It suspends the forward channel of the tier above, as established for two vehicles in Section 6.3, and simultaneously suspends the forward channel of the tier below, whose sensing target has moved. The suspension is of the same kind in both cases, since Section 6.4 establishes that a Follower deprived of sensing holds its last valid estimate rather than acquiring an incorrect one, but it affects two links rather than one, and the vehicle at the bottom of a chain of k tiers is exposed to a suspension whenever any tier above it replies.



A reply by an intermediate vehicle suspends the sensing edge above it and the sensing edge below it at the same time.

Figure 8.6: Why a reply is not a local event in a chain. An intermediate vehicle senses the tier ahead and is sensed by the tier behind, so when it is granted the reply configuration and climbs, it simultaneously leaves the position from which the vehicle below was sensing it. Both suspensions are of the kind established in Section 6.4, in which a deprived Follower holds its last valid estimate rather than acquiring an incorrect one, but two links are affected rather than one. Propagating grants downward, so that a vehicle granted a slot does not itself issue one until it has returned to sensing, keeps at most one tier in the reply configuration at any instant.

The remedy follows from the same exclusivity that produced the grant in the first place. Because the reply configuration is granted rather than seized, and because grants are issued by the vehicle that will receive the reply, a chain can enforce that at most one tier is in the reply configuration at any instant simply by propagating the grant downward: a vehicle that has been granted a slot by the tier above does not itself grant one to the tier below until it has returned to sensing. No new mechanism is required, only an ordering of the grants that already exist. The cost is that the liveness budget of Section 8.4 now applies to the chain as a whole rather than to each tier independently, since the cycles cannot overlap, and the admissible chain length is bounded by the same ratio that bounded the number of Followers in a star.

The bound this places on the return channel is not the only one a chain is subject to, and the two are of different origin. The forward channel accumulates estimation error along the chain, since each vehicle regulates its geometry with respect to a target that is itself regulating its own, and the conditions under which that accumulation remains bounded are a property of the sensing and pursuit loops rather than of the feedback protocol. Those conditions are established in

the companion work on the forward channel and are not reproduced here. What the return channel contributes is orthogonal to them: it does not reduce the accumulated error, and it does not enter the loop that produces it. It bounds instead the consequence of a tier ceasing to track at all, by ensuring that the vehicle ahead of a silent Follower stops rather than continuing to advance, so that a failure at one tier does not propagate as an unbounded separation at every tier behind it.

8.7 Reliability of the Granted Protocol

The two-vehicle protocol tolerates the loss of a symbol without a retransmission mechanism, and Section 5.7 established why: no symbol carries state forward, the state machine returns to its idle condition after every event and every timeout, and a lost or misread symbol therefore affects the interpretation of one event only. This section asks whether that argument survives the introduction of a grant, since a granted protocol does carry state between the two vehicles, namely the fact that a slot is open and to whom.

Three losses are possible where the two-vehicle system admitted one. The announcement may be lost, in which case no Follower interrupts its pursuit and the formation remains in the sensing configuration; the Leader, receiving nothing within the interval it allotted, resumes and may address the same Follower again at the next cycle. The reply may be lost, which is the case already covered by Section 5.7, with the additional consequence that a cycle has been spent without resetting the liveness timer it was intended to reset. The announcement may be misclassified, which names the wrong Follower and is the case the geometric check and the identity confirmation of Section 8.1 exist to contain. In none of the three does the formation enter a configuration it cannot leave, and in all three the recovery is the same: the cycle ends, the vehicles are in the sensing configuration, and the next cycle proceeds independently of the outcome of the previous one.

This is the property that the two-vehicle argument turned on, and it survives because the grant is not an agreement but an offer with a deadline. The Leader does not wait indefinitely for a Follower to occupy a slot it has opened, and a Follower does not occupy a slot it was not addressed in. Neither vehicle holds a commitment that the other must discharge before the protocol can proceed, so no loss can leave the pair in disagreement about whose turn it is. What a granted protocol adds is not a new class of failure but a new cost for the existing ones, since a cycle consumed by a failed exchange is a cycle unavailable to the other Followers, and the liveness budget of Section 8.4 is spent whether or not the exchange succeeded.

The one case that has no counterpart in the two-vehicle system is the undetected collision, which arises only if two Followers occupy the reply configuration simultaneously. Section 8.1 established that this produces a merged excursion the classifier

reads as a single symbol of the wrong class, and Section 8.3 established that rotor signatures make it detectable where the sampling rate permits and leave it silent where it does not. On hardware where it remains silent, the three conditions of Section 8.1 are what prevent the situation from arising, and their independence is what makes that prevention credible: a collision requires that two Followers both classify the same announcement as their own address and both find themselves within the range the geometric check admits, which are unrelated events.

8.8 Operating Envelope of the Multi-Follower Protocol

The constraints developed in this chapter are of three kinds, and this section collects them and asks which of them binds first, since that determines what would have to change for the formation to grow.

The geometric constraint is the hardest and the least negotiable. The reply configuration admits one vehicle, within the recommended axial range of 20 to 28 cm and a radial extent set by a half-velocity radius of 9.6 cm, and no protocol arrangement relaxes it. It does not, however, bound the size of the formation: it bounds the number of vehicles that may reply at once, which the grant mechanism has already reduced to one by construction. The geometric constraint is therefore binding on the design of the protocol rather than on its scale, and having been accommodated it does not reappear.

The addressing constraint bounds how many Followers can be named, and is set by the number of separable duration classes in the forward direction. It is not quantified here, since the spread of a detection at sensing separation has not been characterised, but its scaling is known: additional addresses require longer announcements, and longer announcements lengthen the cycle. Addressing therefore does not impose an independent limit so much as it converts one resource into another, spending cycle time to buy names.

The timing constraint is the one that binds, and it does so with a number rather than a scaling. A cycle serving one Follower occupies 5.3 s by Equation (8.3), against a liveness window of 10.0 s, so the protocol as dimensioned admits a single Follower. Every mechanism developed in this chapter that improves scalability acts on one of the two quantities in that ratio: passive identification removes the codeword from the numerator, hierarchical decomposition divides the population among several windows, and enlarging the window buys formation size at a latency cost proportional to it. None of them acts on the term that actually dominates the cycle, which is the manoeuvre, and which is governed by the Leader's altitude loop rather than by anything in this chapter.

Two conclusions follow, and they point in opposite directions. The protocol

extends to a multi-Follower formation without any change to the physical layer, the detector, the alphabet or the state machine: the announcement that already arrests the Follower's pursuit carries the address, the estimate the Follower already computes performs the geometric check, and the mechanisms that already handle a lost symbol handle a failed grant. In this sense the extension is structural rather than inventive, which is the strongest evidence that the two-vehicle design was the right one. But the formation it extends to is small, and it is small for a reason that has nothing to do with the return channel's own design: the reply configuration is a physical position, and a vehicle can only be in one position at a time. The protocol spends nine per cent of its cycle communicating and fifty-three per cent manoeuvring, and no refinement of the communication can recover the difference. Scaling it substantially would require either that liveness be observable without occupying that position, which the sensing geometry of Section 8.3 forbids while the wake travels downward, or that the formation be decomposed until no vehicle serves more Followers than its liveness window admits, which is what Section 8.6 describes.

Chapter 9

Conclusions

9.1 Summary of the Individual Contribution

This thesis set out to give a leader-follower formation something the baseline system did not have: a path for information to travel from the Follower back to the Leader. The problem was not that no such path could be built, but that the two constraints inherited from the baseline left almost nothing to build it from. A radio link was excluded by assumption, since the entire premise of implicit communication is that coordination should not depend on one, and adding a sensor to either vehicle was excluded by the platform, which carries what a nano-quadrotor carries and nothing more. What remains, once both are ruled out, is the physics the vehicles already produce and the instruments they already have: a rotor that generates a wake whether or not anyone intends it to carry information, and an inertial estimate that registers a disturbance whether or not anyone intends to read one.

The channel developed here is built entirely from those two. The transmitter is the Follower's own altitude loop, commanded to oscillate about its nominal position so that its downwash is modulated at a chosen frequency; the receiver is the Leader's own vertical velocity estimate, from which that modulation is recovered. Nothing is added to either vehicle beyond the code that generates the oscillation and the code that detects it, and both run inside the control loop that was already executing. In this sense the contribution is not a new device but a demonstration that the device was already present, unused.

Around that carrier, the thesis develops a complete communication system rather than a single detection experiment. Chapter 3 models the channel from measurement, fitting the wake decay and the radial profile from a dedicated campaign rather than adopting published constants, and from that model derives the constraints that bound the gesture frequency, its amplitude and its duration. Chapter 4 defines what the channel carries: an alphabet of three symbols distinguished by duration

alone, sized so that the classifier can separate them, spaced by a guard interval that keeps consecutive transmissions from merging, and arbitrated by a fixed priority that decides which condition is reported when more than one is pending. Chapter 5 builds the receiver: a high-pass stage that removes what carries no information, a quadrature matched filter whose energy does not depend on the phase at which a gesture arrives, a threshold placed between two measured levels rather than derived from a noise model that does not apply, a duration classifier, and a state machine that converts symbols into decisions about whether the formation may continue to advance.

Two further contributions concern the system rather than the channel. Chapter 6 establishes that the two channels cannot operate simultaneously, since the return channel requires the vertical ordering of the vehicles to be inverted and the forward channel requires it not to be, and develops the coordination that lets them alternate without either corrupting the other. Chapter 8 extends the protocol to a formation of several Followers and shows that the extension requires no change to the physical layer, the detector, the alphabet or the state machine, because the mechanism that grants the reply configuration already exists in the two-vehicle system and lacks only a destination.

The thesis also states, throughout, what the channel does not do. The return path does not make the formation converge faster and does not alter what it converges to; the graph argument of Section 8.5 shows that no feedback edge reaches the reference. What it changes is the behaviour of the formation when the assumptions underneath it stop holding. A Leader without a return path remains connected to its reference under every condition and continues to advance toward it, so a Follower that has failed is left behind at a separation that grows without bound. A Leader with one disconnects itself from the reference and halts. The contribution of this work, stated as precisely as it can be, is that it makes the convergence of a wake-coupled formation interruptible on evidence.

9.2 Main Experimental Findings

The measurements reported in Chapters 3 and 7 support the design at the level of individual parameters rather than only at the level of the system, and several of them produced results that shaped the design rather than confirming it.

The channel model was fitted rather than assumed. The centreline decay and the self-similar radial profile of Section 3.2 reproduce the anemometric campaign with a half-velocity radius of $r_{1/2} = 9.6\,\mathrm{cm}$, and the negative development length recovered by the fit, which would be difficult to interpret as a physical location, is consistent in sign and order with the published reference. The model matters because everything downstream is dimensioned from it: the disturbance the Leader

receives during a gesture is the modulation this law predicts, and the amplitude of the gesture was sized against that prediction.

The altitude loop does not execute what it is commanded. At $f_0 = 2.5\,\mathrm{Hz}$ the loop transmits the reference with attenuation and lag rather than tracking it, and the executed excursion is 8.7 cm against 12 cm commanded, at a lag of 69°. This is not a defect but a consequence of placing f_0 close to the bandwidth deliberately, and the second-order response of Equation (7.1) reproduces both figures to three significant digits from the identified loop parameters alone. The consequence for the link budget is direct: the signal reaching the Leader is set by what the Follower performs, not by what it asks for.

The operating window is narrow and was located by sweep rather than by argument. Classification is complete between 18 and 30 cm of vertical separation and degrades quickly outside, and the recommended range of 20 to 28 cm stops short of both edges because neither carries margin. The two boundaries fail differently, which is the useful part of the result: above the upper edge the energy excursion fragments before clearing the threshold, while below the lower edge it widens past the first classification boundary, so that a liveness indication is delivered to the state machine as an acknowledgement. The second failure is the more consequential, because the protocol acts on it.

Vertical clearance was found to be exhausted at almost the same separation as the classification margin. With the executed excursion, the instantaneous gap at the bottom of every cycle closes to 11.3 cm at a nominal separation of 20 cm, of which the two airframes occupy 8 cm; contact would occur at a nominal separation of 16.7 cm. That the two limits coincide so closely was not designed and is a coincidence of the platform, but it means the recommended range is bounded from below twice over.

The detection threshold could not be derived from the statistics of the noise floor, because consecutive matched-filter outputs share all but one sample of their correlation window and are therefore strongly correlated. It was placed instead between two measured levels, a noise-floor maximum of 0.013 and a typical gesture energy of 0.25, at $\lambda = 0.030$: margins of 3.6 dB below and 9.2 dB above. The asymmetry is deliberate, since the two failure mechanisms it guards against are not equally costly, and the two margins cannot be widened at once.

Two timing parameters were fixed empirically and one of them by a result that arrived unexpectedly. The guard interval was set to 300 samples after two symbols separated by 150 were found to be the closest pair still recovered as two events, the value adopted being twice the measured minimum. The blanking interval was set to 400 samples, four time constants of the high-pass stage rather than the 200 samples needed to fill the correlation buffer, and the start-up transient was subsequently found to cross the detection threshold at almost exactly that point: sizing the blanking to the buffer alone would have left the Leader acting on its own

initialisation.

At the formation level, the grant cycle is dominated by the manoeuvre and not by the channel. With the role-swap window estimated from the two altitude loops, a cycle carrying one heartbeat occupies 5.3 s, of which the manoeuvre accounts for fifty-three per cent, the guard interval twenty-eight, and the symbol the Leader is waiting for nine. Against the liveness window of 10 s inherited from the two-vehicle system, this admits one Follower, and the relation $T_{\mathrm{live}} \geq 5.3\,N$ states what must be relaxed to admit more. The most effective way to shorten the cycle acts on none of the communication parameters: the manoeuvre is limited throughout by the Leader's altitude loop, seven times slower than the Follower's, so raising that bandwidth shortens the cycle more than any change to the protocol could.

Finally, the bench campaign of Section 8.3 established that a rotor carries an identity in its own acoustic signature. Two-, three- and four-bladed propellers produced from the same parametric model place their blade-passing component at 139, 208 and 272 Hz at a common shaft speed near 69 Hz, each within one per cent of the integer multiple predicted, and two such components remain individually resolvable in superposition even when they lie 65 Hz apart. The result is not usable on the present platform, for reasons given below, but it establishes that the identifier exists and is well behaved.

9.3 Limitations of the Gesture-Based Feedback Channel

The limitations fall into three groups, and they are of quite different weight.

The first concerns validation. The results of Chapter 7 were obtained from the full model and verified against the vehicle, which establishes that the decoder behaves as designed and that the parameters are consistent, but it is not the same as a flight campaign with statistics collected over many trials in varying conditions. The noise environment in particular is modelled rather than measured: the standard deviation attributed to the Leader's vertical velocity estimate is an estimate rather than a quantity identified from a hover log, and while no conclusion in this thesis depends on its exact value, the lower boundary of the design space in Section 3.3 is indicative for that reason. All testing was conducted without external wind, which for a channel whose carrier is a low-frequency aerodynamic disturbance is a significant idealisation.

The second concerns the channel itself, and these limits are structural rather than provisional. The return channel occupies a physical position, not a frequency band, and the volume from which a gesture reaches the Leader is of the order of one airframe. One vehicle can be in it at a time, so the capacity of the channel is not increased by any protocol arrangement, and every reply costs a manoeuvre whose

duration is set by flight dynamics rather than by signal processing. The alphabet is correspondingly small: three symbols distinguished by duration, sufficient for the three things the one-way channel could not carry but not extensible without lengthening the cycle that already dominates the formation budget. The channel is also strictly half-duplex with the forward channel, since the two require opposite vertical orderings, so every transmission is paid for with an interruption of sensing.

The third concerns the mechanisms this thesis proposes but could not demonstrate on the platform. The rotor signature that would let the Leader identify which Follower replied lies between roughly 140 and 510 Hz, entirely above the Nyquist limit of the 200 Hz decoder, and logs recorded on the vehicle confirm that with a Leader present the Follower's attitude residual rises broadband rather than tonally. A second gap separates the bench measurement from the intended application even where sampling permits: the campaign measured acoustic pressure, whereas the quantity available to the Leader is an aerodynamic moment, so what has been established is that the source produces a tonal component, not that the component survives into the channel the protocol would read it from. For the same reason, the multi-Follower protocol of Chapter 8 is developed and dimensioned but not demonstrated, and the number of addresses it can distinguish remains unquantified because the spread of a detection in the forward direction has not been characterised.

9.4 Future Work on Robust Bidirectional Implicit Communication

The most direct continuation follows from the limitation that closes the previous section. An inertial sensor of the class routinely fitted to first-person-view airframes samples in the kilohertz range and would place the entire blade-passing band within the readable spectrum, which is the necessary condition for passive identification. It is not the sufficient one: what must be established before that mechanism can be relied upon is whether the tonal component rises above the broadband turbulent floor in the aerodynamic moment, which the bench campaign cannot answer because it measured a different quantity. A dedicated measurement on a kilohertz-class gyroscope, with a receiver that keeps the time dynamics rather than averaging them away, would settle it either way, and either answer would be worth having.

The second direction is the one the scalability analysis pointed to unexpectedly. Because more than half the grant cycle is spent on a manoeuvre limited by the Leader's altitude loop, and because that loop is seven times slower than the Follower's, the single most effective change available to a designer wishing to scale this protocol is to that loop rather than to anything in the communication chain.

Whether the Leader's bandwidth can be raised without compromising the stability of the pursuit it anchors is a control question rather than a communication one, and answering it would convert a protocol that admits one Follower into one that admits three.

The third is experimental and is what the thesis most obviously lacks. A multi-vehicle campaign would test the extension of Chapter 8 where it is most exposed, namely in the assumption that the three conditions placed between a grant and a role swap are independent in practice as well as in principle, and would provide the flight statistics that the present validation approximates. The same campaign would characterise the spread of a detection in the forward direction and thereby quantify how many Followers can be addressed, which is the last parameter of the protocol that remains open.

Two smaller items would strengthen what is already here. Identifying the noise floor of the Leader's vertical velocity estimate from a hover log would anchor the lower boundary of the design space and convert a qualitative statement about detection margin into a quantitative one. And testing the channel in the presence of external wind would establish whether the high-pass stage, which was designed to remove sensor bias and slow drift, also removes a disturbance that is neither slow nor stationary, or whether the carrier frequency would have to move to avoid it.

Beyond these, the result that seems most worth pursuing is the one this thesis could only demonstrate as a spectral property. That two rotors of different blade count remain individually resolvable when their signatures are superposed says that the aerodynamic channel can carry more than one sublevel, and it says so for a reason as robust as the linearity of the transform. Whether two physically adjacent wakes merge in a way that preserves that separability is a question about fluid dynamics rather than about signal processing, and nothing in this work addresses it. If they do, then the identifier developed here as a way of naming a Follower is also a way of letting several of them speak at once, and the constraint that has bounded this protocol throughout, that the reply configuration admits one vehicle at a time, would be the first thing worth revisiting.

Bibliography

[1] F. Pettinato. «Downwash-Based Leader Sensing and Relative-State Estimation for Implicit UAV Formation Control». MA thesis. Turin, Italy: Politecnico di Torino, 2026 (cit. on p. 1).

[2] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. «Swarm Robotics: A Review from the Swarm Engineering Perspective». In: *Swarm Intelligence* 7.1 (2013), pp. 1–41 (cit. on p. 11).

[3] E. Şahin. «Swarm Robotics: From Sources of Inspiration to Domains of Application». In: *Swarm Robotics*. Berlin, Germany: Springer, 2005, pp. 10–20 (cit. on p. 11).

[4] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. «Novel Type of Phase Transition in a System of Self-Driven Particles». In: *Physical Review Letters* 75.6 (1995), pp. 1226–1229 (cit. on p. 11).

[5] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. «Collective Memory and Spatial Sorting in Animal Groups». In: *Journal of Theoretical Biology* 218.1 (2002), pp. 1–11 (cit. on p. 11).

[6] A. D. Dragan and S. S. Srinivasa. «Generating Legible Motion». In: *Proc. Robotics: Science and Systems*. Berlin, Germany, June 2013 (cit. on pp. 12, 18).

[7] A. D. Dragan, K. C. T. Lee, and S. S. Srinivasa. «Legibility and Predictability of Robot Motion». In: *Proc. 8th ACM/IEEE Int. Conf. Human-Robot Interaction (HRI)*. Tokyo, Japan, Mar. 2013, pp. 301–308 (cit. on p. 12).

[8] B. Capelli, C. Secchi, and L. Sabattini. «Communication Through Motion: Legibility of Multi-Robot Systems». In: *Proc. Int. Symp. Multi-Robot and Multi-Agent Systems (MRS)*. 2019, pp. 126–132 (cit. on p. 12).

[9] D. Szafir, B. Mutlu, and T. Fong. «Communication of Intent in Assistive Free Flyers». In: *Proc. ACM/IEEE Int. Conf. Human-Robot Interaction (HRI)*. Bielefeld, Germany, Mar. 2014, pp. 358–365 (cit. on pp. 12, 18).

[10] D. Szafir, B. Mutlu, and T. Fong. «Communicating Directionality in Flying Robots». In: *Proc. 10th ACM/IEEE Int. Conf. Human-Robot Interaction (HRI)*. Portland, OR, 2015, pp. 19–26 (cit. on p. 12).

[11] B. A. Duncan, E. Beachly, A. Bevins, S. Elbaum, and C. Detweiler. «Investigation of Communicative Flight Paths for Small Unmanned Aerial Systems». In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. Brisbane, Australia, 2018, pp. 602–609 (cit. on p. 12).

[12] S. M. Kay. *Fundamentals of Statistical Signal Processing, Volume II: Detection Theory*. Upper Saddle River, NJ: Prentice-Hall, 1998 (cit. on pp. 13, 14).

[13] H. Urkowitz. «Energy Detection of Unknown Deterministic Signals». In: *Proceedings of the IEEE* 55.4 (Apr. 1967), pp. 523–531 (cit. on pp. 13, 14).

[14] G. L. Turin. «An Introduction to Matched Filters». In: *IRE Transactions on Information Theory* 6.3 (June 1960), pp. 311–329 (cit. on pp. 13, 14).

[15] D. O. North. «An Analysis of the Factors Which Determine Signal/Noise Discrimination in Pulsed-Carrier Systems». In: *Proceedings of the IEEE* 51.7 (July 1963), pp. 1016–1027 (cit. on p. 14).

[16] H. L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I*. New York: Wiley, 2001 (cit. on p. 14).

[17] A. S. Tanenbaum and D. J. Wetherall. *Computer Networks*. Fifth. Boston, MA: Pearson, 2011 (cit. on pp. 16, 44).

[18] D. Bertsekas and R. Gallager. *Data Networks*. Second. Englewood Cliffs, NJ: Prentice-Hall, 1992 (cit. on pp. 16, 17).

[19] N. A. Lynch. *Distributed Algorithms*. San Francisco, CA: Morgan Kaufmann, 1996 (cit. on p. 16).

[20] L. Bauersfeld, K. Muller, D. Ziegler, F. Coletti, and D. Scaramuzza. «Robotics Meets Fluid Dynamics: A Characterization of the Induced Airflow Below a Quadrotor as a Turbulent Jet». In: *IEEE Robotics and Automation Letters* 10 (2025), pp. 1241–1248 (cit. on p. 22).

[21] C. Bingham, M. Godfrey, and J. Tukey. «Modern Techniques of Power Spectrum Estimation». In: *IEEE Transactions on Audio and Electroacoustics* 15.2 (June 1967), pp. 56–66 (cit. on p. 29).

[22] F. J. Harris. «On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform». In: *Proceedings of the IEEE* 66.1 (Jan. 1978), pp. 51–83 (cit. on p. 29).