



**Università
di Genova**

MASTER OF SCIENCE IN ROBOTICS ENGINEERING

**Motion Control of an Autonomous Underwater
Vehicle Subject to a Constant Unknown Ocean
Current**

Master's Thesis

Candidate

Mohammad Sabouri
Student ID: s5659227

Supervisor

Prof. Giovanni Indiveri
University of Genoa (Italy)



**Università
di Genova**

DIBRIS DIPARTIMENTO
DI INFORMATICA, BIOINGEGNERIA,
ROBOTICA E INGEGNERIA DEI SISTEMI

Academic Year 2025–2026

September 2026

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Prof. Giovanni Indiveri, for his guidance, careful discussions, and support throughout this work. I also thank the members of the GRAAL laboratory at the University of Genoa for providing a stimulating scientific environment and for their help with the research context in which this thesis was developed.

Abstract

This thesis develops a current-aware three-dimensional guidance law for an underactuated autonomous underwater vehicle moving through an unknown constant ocean current. At the kinematic level, the vehicle is described by its inertial position, unit surge direction, water-relative surge speed, and a roll-free angular-velocity input. The reference is a moving target constrained to a prescribed spatial path. The formal analysis is carried out with respect to the moving-reference error, while the geometric distance from the vehicle to the path is reported separately as an evaluation metric.

Starting from a published MED switching predecessor that alternates between target approach and path-tangent alignment, the proposed law replaces the discrete task switch with a smooth blend of the two mechanisms. The blend is defined on the set where both the desired direction and the required surge speed are mathematically well defined and physically feasible. A constant-current vector observer estimates the environmental disturbance, a current-feasible surge construction preserves physical realisability, and a signed along-path speed modulation removes a steady along-path error that may persist when direction alignment and nominal speed matching are used alone. For the reduced ideal kinematic model, the analysis proves exponential convergence of the current observer, convergence of the direction-alignment error from every initial direction except the antipodal configuration, and conditional convergence of the moving-reference position error under the stated regularity, feasibility, and admissible-domain assumptions. A full stability proof for the six-degree-of-freedom vehicle cascade is not claimed.

A frame-consistent interface maps the guidance commands to rudder, stern-plane, and propeller references of a six-degree-of-freedom REMUS 100 model driven by a PI cascade. On the ideal kinematic plant, the current-estimation error decays to a negligible numerical error floor relative to the true current magnitude. On the REMUS 100/PI implementation, the proposed law achieves RMS geometric path distances between 0.095 m and 0.431 m on the four curved or survey references and 0.970 m on the inclined line. Relative to the published MED switching predecessor, the RMS distance is reduced by approximately 59–97% across the five test paths. The same replacement removes the discrete task switch and reduces the surge-command roughness by factors from approximately 1.6×10^2 to 3.8×10^4 , while keeping steady-window actuator saturation at zero on the helix and lemniscate and below 8.2% on the remaining paths. A recent three-dimensional virtual-reference-point baseline attains lower RMS path distance on the helix and lemniscate; the proposed law gives lower RMS on the line, lawnmower, and terrain-following survey under the same transfer interface, together with explicit current estimation, feasibility monitoring, and implementation checks. An autonomous path-progression extension built on the same direction–speed structure is retained as future work.

Contents

Acknowledgements	iii
Abstract	v
Nomenclature	xv
1 Introduction	1
1.1 Context and motivation	1
1.2 Research problem, aim, and objectives	2
1.3 Contributions and scope	2
1.3.1 Validation programme and scope	3
1.3.2 Relation to the published two-task formulation	3
1.4 Thesis organisation	4
1.5 Chapter summary	4
2 Background and State of the Art	5
2.1 Review scope and organising concepts	5
2.2 Foundations of path following	5
2.2.1 Geometric paths, progression, and error	5
2.2.2 Compact classification of formulations	6
2.2.3 Underactuation and spatial geometry	7
2.3 Guidance methods for underactuated marine vehicles	7
2.3.1 LOS and adaptive LOS	7
2.3.2 Virtual targets, vector fields, and offset points	8
2.4 Current compensation and dynamic realisation	9
2.4.1 Integral adaptation and explicit observation	9
2.4.2 Guidance–dynamics interface	9
2.5 Critical synthesis and thesis positioning	10
2.5.1 Local research gap	11
2.6 Chapter summary	11
3 Three-Dimensional Kinematic Guidance	13
3.1 Problem Formulation and Kinematic Preliminaries	13
3.1.1 Path-following formulation	13
3.1.2 Angular velocity of the geometric directions	14
3.1.3 State and measurement conventions	14
3.1.4 Standing assumptions	15
3.1.5 Path representation and error quantities	15
3.1.6 Control objective and scope of the proof	16

3.2	Initial Two-Task Guidance Formulation	16
3.2.1	Constant-current observer	16
3.2.2	Two tasks and hysteresis logic	18
3.2.3	Base surge speed and feasibility	18
3.2.4	Task 1: target approach	18
3.2.5	Task 2: path-tangent alignment	19
3.2.6	Physical angular-velocity command	19
3.2.7	Unified longitudinal overshoot policy	20
3.2.8	Properties of the initial formulation	20
3.3	Smooth Direction Design and Stability Analysis	20
3.3.1	Smooth blended target direction	21
3.3.2	Surge speed and direction command	22
3.3.3	Reduced mapping from desired direction rate to vehicle angular rate	23
3.3.4	Direction-alignment analysis	24
3.3.5	Position error under nominal surge	25
3.3.6	Modified surge policies	26
3.3.7	Continuous policy and zero-error convergence	28
3.3.8	Layered convergence statement for the ideal kinematic model	29
3.4	Chapter summary	30
4	Dynamic Model and Controller Integration	31
4.1	Six-Degree-of-Freedom REMUS 100 Model	31
4.1.1	Purpose and modelling scope	31
4.1.2	Frames and state	32
4.1.3	Ocean current and relative velocity	32
4.1.4	Kinematics	32
4.1.5	Equations of motion	33
4.1.6	Rigid-body and added-mass properties	33
4.1.7	Damping and restoring forces	34
4.1.8	Actuators and generalized input	34
4.1.9	Physical constraints	35
4.1.10	Model traceability and coefficient declaration	35
4.1.11	Reduced models for controller design	35
4.2	Frame-Consistent Ocean-Current Representation	36
4.2.1	Yaw-only mapping and attitude-dependent leakage	36
4.2.2	Full-rotation correction	37
4.2.3	Implementation boundary	37
4.3	Kinematic-to-Dynamic Command Interface	37
4.3.1	Purpose and scope	37
4.3.2	Closed-loop data flow	37
4.3.3	Conversion of the angular command	38
4.3.4	Water-relative surge and ground-frame measurement	38
4.3.5	Low-level actuator interface	39
4.3.6	DVL-aided current observer	39

4.3.7	Realised-command error and analytical implications	40
4.4	Chapter summary	40
5	Numerical Results and Comparison	41
5.1	Scenarios and comparison protocol	41
5.2	Kinematic simulation results	44
5.2.1	Trajectory and position error	44
5.2.2	Observer convergence	46
5.2.3	Alignment law	47
5.2.4	Outer position consistency checks	48
5.2.5	Empirical check of the conditional bound	50
5.3	Kinematic–dynamic simulation results	50
5.3.1	Trajectory following on the helix	50
5.3.2	Position error under the dynamic plant	52
5.3.3	Command–response tracking on the inner loop	53
5.3.4	Observer convergence under the dynamic plant	55
5.3.5	Terrain-following outcome on the survey reference	57
5.4	Comparative evaluation	60
5.4.1	Main benchmark on the helix	60
5.4.2	Trajectory overlay and path-distance time history	61
5.4.3	Path-distance behaviour on the lawnmower	64
5.4.4	Path-distance behaviour on the terrain-following survey	64
5.4.5	Performance across reference geometries	65
5.4.6	Note on the straight-line dynamic behaviour	67
5.5	Discussion, limitations, and interpretation	67
6	Conclusions and Future Work	71
6.1	Conclusions	71
6.2	Future work	72
A	Supplementary Derivations	73
A.1	Derivative of a normalized vector	73
A.2	Reduced angular-command inversion and approximation errors	73
A.3	Blend simplification and singularity	74
A.4	Positive surge root	75
A.5	Alignment derivative	75
A.6	Proof of the target-projection inequality	76
A.7	Mode-wise proof for the hysteresis policy	77
A.8	Continuous-policy position decay	77
A.9	Uniform-decay condition for asymptotic position convergence	78
A.10	Dynamic interface error identity	78
A.11	Sufficient time-scale approximation condition	79
A.12	Bounded derivative of the position Lyapunov function	79
A.13	Modulation A and its effective speed ratio	79

A.14	Quantitative decay for the continuous policy	80
A.15	Layered kinematic convergence proof	80
A.16	Unit-direction integration	82
B	Observer and Current-Frame Derivations	83
B.1	Exponential convergence of the constant-current observer	83
B.2	Mismatch created by a yaw-only current transformation	85
B.3	Full-rotation correction and recovery of the nominal observer	86
C	Supplementary Numerical Results	87
C.1	Extended numerical consistency checks	87
C.2	Five-method comparison on all reference paths	87
C.3	Ideal-plant comparison	87
C.4	Extended dynamic response checks	94
D	Simulation Setup Details	99
D.1	Reference-path parameterisations	99
D.2	Guidance and inner-loop parameters	100
D.3	Initial conditions	102
D.4	Performance-metric definitions	103
D.5	Runtime feasibility and fallback protocol	103
D.6	REMUS 100 plant configuration	104
D.7	Ocean-current model	104
D.8	Numerical settings	105
D.9	Observer frame-consistency experiment	105
D.10	External comparison protocol	106
	D.10.1 Comparative evaluation protocol	107
E	Conditional Cascade-Stability Analysis	109
E.1	Interconnection model and standing hypotheses	109
E.2	Outer-loop input-to-state stability	111
E.3	Inner-loop input-to-state stability	112
E.4	Cascade input-to-state stability (Route 1)	114
E.5	Composite Lyapunov route (Route 2)	115
E.6	Conditional ultimate bound (Route 3)	117
E.7	Summary and unmodified controller	118

List of Figures

1.1	Guidance geometry for a current-disturbed AUV.	1
3.1	Internal block structure of the kinematic guidance layer.	25
4.1	The REMUS 100 autonomous underwater vehicle.	31
4.2	Complete closed-loop kinematic-dynamic cascade.	38
5.1	Vehicle trajectory on the helix under the four formally analysed self-configurations and the exploratory path-following variant (case K1).	45
5.2	Position-error norm on the helix (case K1).	46
5.3	Current-observer convergence on the helix (case K1).	47
5.4	Blend regularity, target projection, projection gap, and outer sign checks on the helix (case K1).	49
5.5	Vehicle following the helix under four analysed self-configurations and the exploratory path-following variant (case D1).	51
5.6	Position-error norm under the dynamic plant (case D1, helix).	52
5.7	Surge command–response tracking (case D1, helix).	54
5.8	Actuator commands against saturation limits (case D1, helix).	55
5.9	Current-observer convergence under the dynamic plant (case D1, helix).	56
5.10	Trajectory over the seabed for the four analysed internal configurations and the exploratory progression variant on the survey reference (case D4).	58
5.11	Altitude, seabed cross-section, and depth histories for the four analysed internal configurations and the exploratory progression variant on the terrain-following survey (case D4).	59
5.12	Five-method trajectory overlay on the helix (case D2).	62
5.13	Five-method geometric path distance on the helix, shown on linear and logarithmic scales (case D2).	63
5.14	Five-method geometric path distance on the lawnmower (case D3).	64
5.15	Five-method geometric path distance on the terrain-following survey (case D3).	65
C.1	Alignment storage for the four analysed internal configurations and the exploratory progression variant on the kinematic helix experiment.	88
C.2	Five-method geometric path distance on the inclined line.	89
C.3	Five-method geometric path distance on the lawnmower.	90
C.4	Five-method geometric path distance on the spatial Gerono lemniscate.	91
C.5	Five-method geometric path distance on the terrain-following survey.	92
C.6	Altitude, seabed cross-section, and vehicle-depth histories for the four analysed internal configurations and the exploratory progression variant on the oceanscan reference.	93
C.7	Attitude histories on the helix.	94

C.8	Body-frame velocity histories on the helix.	95
C.9	Body-frame angular-rate histories on the helix.	96
C.10	Cumulative angular effort and surge-command roughness on the helix. . . .	97

List of Tables

1.1	Relationship between the published formulation and the developments in this thesis.	4
2.1	Condensed development of marine path-following control.	6
2.2	Compact classification of path-following formulations.	6
2.3	Closest current-aware approaches and the position of this thesis.	10
3.1	State, estimate, command, and realisation conventions used in the thesis. . .	15
4.1	Approximate effective inertia of the REMUS 100 model.	33
4.2	Control-relevant REMUS 100 limits and nominal parameters.	35
5.1	Reference paths used in the validation programme.	42
5.2	Simulation case map.	42
5.3	Performance metrics used in Chapter 5.	43
5.4	Empirical UUB-radius sensitivity on the helix.	50
5.5	Own-method dynamic helix summary.	51
5.6	Terrain-following assessment.	57
5.7	Compared implementations.	60
5.8	Five-method helix benchmark.	61
5.9	RMS geometric path distance on the dynamic plant [m].	66
5.10	Settling time on the dynamic plant [s].	66
5.11	Cross-method interpretation.	67
C.1	Five-method RMS geometric path distance [m] on the ideal kinematic plant. Bold = best in column; <i>italic</i> = best of the thesis-lineage implementations when the overall best is a non-thesis-lineage method.	89
C.2	Five-method settling time T_s [s] on the ideal kinematic plant. Bold = best (smallest) finite value in column; <i>italic</i> = best of the thesis-lineage implementations when the overall best is a non-thesis-lineage method. . . .	90
C.3	M-B continuous-PP across plant levels.	91
D.1	Seabed model $z_{\text{bed}}(x, y) = z_c + \sum_i A_i \sin(k_{x,i}x + k_{y,i}y + \phi_i)$ used for the oceanscan reference.	100
D.2	Blended-alignment guidance law and observer gains.	101
D.3	Simulation parameters of the switching MED predecessor. The second column gives the values used in the present transfer study; these are not presented as a reproduction with the conference-paper tuning.	101
D.4	Inner-loop PI cascade gains and actuator limits.	101
D.5	Published outer-guidance gains retained for the two external baselines. . . .	102

D.6	Per-path initial-condition composition for the comparative runs.	102
D.7	REMUS 100 geometry, mass, and environment.	104
D.8	Hydrodynamic damping, added-mass, and lift/drag parameters.	104
D.9	Propeller and fin parameters.	105
D.10	Numerical integration settings.	106
D.11	Controlled conditions for the observer frame-consistency experiment (case D5).	106

Nomenclature

Symbol	Description	Unit
Reference frames and path geometry		
$\mathcal{F}_n$	North-East-Down inertial reference frame	–
$\mathcal{F}_b$	Body-fixed reference frame	–
$\mathbf{q}$	Vehicle position vector	m
$\mathbf{r}$	Moving-reference position vector	m
$\mathbf{r}_\mathcal{P}$	Geometric path parametrisation	m
$\mathbf{r}_H$	Helical reference used in the simulation benchmark	m
$\mathbf{r}_L$	Three-dimensional lemniscate reference	m
$\mathcal{I}$	Domain of the geometric path parameter	–
σ	Geometric path parameter	–
σ_0	Initial value of the path parameter	–
$\sigma_T, \sigma_{T,0}$	Terrain-path coordinate and its initial value	m
$\dot{\sigma}$	Virtual-target path-progression rate	s ^{−1}
$\dot{\mathbf{r}}$	Moving-reference velocity vector	m s ^{−1}
$\ddot{\mathbf{r}}$	Moving-reference acceleration vector	m s ^{−2}
$\mathbf{b}$	Unit tangent vector of the reference path	–
$\mathbf{e}$	Vehicle-to-reference position-error vector	m
$\bar{\mathbf{e}}$	Unit position-error direction	–
$\mathbf{e}_\perp$	Local cross-track error vector	m
$e_\parallel$	Generic along-track component used in the literature review	m
s_e	Tangential target-relative error in the problem formulation	m
s	Signed tangential error used in the smooth-controller analysis	m
ℓ	Tangential-equilibrium coordinate	m
e_ℓ	Signed tangential error used in the two-task formulation	m
R	Helix radius	m
ω_p	Reference-path angular rate	rad s ^{−1}
ω	Lemniscate parameter rate used in the simulation definition	rad s ^{−1}
v_z	Reference-path vertical progression speed	m s ^{−1}
v_∞	Limiting speed of the asymptotically straight reference	m s ^{−1}
A_x, A_y, A_z	Transient displacement amplitudes of the asymptotically straight reference	m
$\mathbf{r}_T$	Terrain-following reference path used in the frame test	m
k_y, k_z	Terrain-following spatial frequencies	m ^{−1}
v_σ	Terrain-following path-progression speed	m s ^{−1}
τ	Transition time constant of the asymptotically straight reference	s
a	Lemniscate scale parameter	m
z_a	Lemniscate vertical-amplitude parameter	m

Symbol	Description	Unit
z_0	Lemniscate vertical offset	m
χ_d	Desired planar course angle	rad
χ_p	Local planar path-tangent angle	rad
e_y	Signed planar cross-track error	m
Δ_{LOS}	Line-of-sight look-ahead distance	m
Kinematic model, observer, and guidance law		
$\mathbf{h}$	Unit vehicle surge-direction vector	—
u	Water-relative surge speed	m s^{-1}
ω	Inertial angular-velocity command	rad s^{-1}
$\mathbf{v}_c$	Unknown constant ocean-current vector	m s^{-1}
$\mathbf{v}_c^n$	Ocean-current vector expressed in the NED frame	m s^{-1}
$\hat{\mathbf{v}}_c$	Estimated ocean-current vector	m s^{-1}
$\tilde{\mathbf{v}}_c$	Ocean-current estimation-error vector	m s^{-1}
$\tilde{\mathbf{v}}_{c,\text{eff}}$	Effective-current estimation-error vector	m s^{-1}
$\hat{\mathbf{e}}$	Estimated position-error vector	m
$\tilde{\mathbf{e}}$	Position-observer error vector	m
$\mathbf{v}_1$	Estimated inertial vehicle-velocity vector	m s^{-1}
$\bar{\mathbf{v}}_1$	Unit estimated inertial-velocity direction	—
$\mathbf{a}$	Unnormalised blended guidance vector	—
$\bar{\mathbf{a}}$	Normalised blended guidance direction	—
p	Blend weighting coefficient	—
λ	Guidance look-ahead length	m
κ_a	Blend normalisation factor	—
Π_e	Position-error projection on the blended direction	m
ω_e	Angular velocity of the position-error direction	rad s^{-1}
ω_{v_1}	Angular velocity of the estimated inertial-velocity direction	rad s^{-1}
$\omega_{v_1}^{\text{des}}$	Desired angular velocity in the two-task formulation	rad s^{-1}
ω_b	Angular velocity of the path-tangent direction	rad s^{-1}
$\omega_{\bar{\mathbf{a}}}$	Angular velocity of the blended guidance direction	rad s^{-1}
$\omega_{\bar{\mathbf{v}}}$	Angular velocity of the estimated velocity direction	rad s^{-1}
$\omega_{\bar{\mathbf{v}}}^{\text{des}}$	Desired angular velocity of the estimated velocity direction	rad s^{-1}
$\boldsymbol{\varepsilon}(t)$	Angular-command approximation error under the time-scale reduction	rad s^{-1}
$J_{\boldsymbol{\varepsilon}}(T)$	Finite-horizon integral of the approximation-error norm	—
ρ_o	Spectral decay rate of the linear observer error system	s^{-1}
σ_o	Observer decay-rate bound used in sufficient estimates	s^{-1}
C_1, C_2	Nonnegative constants in the time-scale approximation bound	—
$\varphi(t)$	Reference/command-rate term in the time-scale approximation test	—
u_{nom}	Nominal surge command	m s^{-1}
u_d	Retained safe feasible surge value	m s^{-1}
u_+, u_-	Positive and negative surge roots	m s^{-1}
Δ	Surge-feasibility discriminant	$\text{m}^2 \text{s}^{-2}$
v_{req}	Inertial speed magnitude required by the selected surge policy	m s^{-1}

Symbol	Description	Unit
Δ_{req}	Runtime feasibility radicand for the selected policy	$\text{m}^2 \text{s}^{-2}$
γ	Surge-modulation factor	–
γ_+	Accelerating surge factor	–
γ_-	Decelerating surge factor	–
γ_{hold}	Surge factor retained in the inner hysteresis mode	–
μ	Effective inertial-to-reference speed ratio	–
Δ_c	Current-dependent term in the Modulation-A speed relation	$\text{m}^2 \text{s}^{-2}$
ρ_{in}	Inner hysteresis radius	m
ρ_{out}	Outer hysteresis radius	m
η	Continuous-policy amplitude	–
α_0	Continuous-policy transition length	m
α	Accelerating surge factor in the two-task formulation	–
β	Decelerating surge factor in the two-task formulation	–
a_{min}	Minimum admissible blend norm	–
c	Directional alignment cosine	–
k	Direction-alignment gain	s^{-1}
k_1, k_2	Directional gains of the two-task formulation	s^{-1}
k_C	Position-observer correction gain	s^{-1}
k_D	Current-observer correction gain	s^{-2}
v_c^{max}	Assumed ocean-current speed bound	m s^{-1}
u_{min}	Minimum admissible surge speed	m s^{-1}
$v_r^{\text{min}}, v_r^{\text{max}}$	Reference-speed bounds	m s^{-1}
a_r^{max}	Reference-acceleration bound	m s^{-2}
κ^{max}	Path-tangent angular-rate bound	rad s^{-1}
V_1, V_2	Task Lyapunov functions of the two-task formulation	–
V_a	Direction-alignment Lyapunov function	–
V_e	Position-error Lyapunov function	m^2
$V_{e,s}$	Scaled position-error storage used in a conditional cascade estimate	–
V_o	Observer-error Lyapunov function	scaled (dim. mixed)
V_{kin}	Conditional composite kinematic Lyapunov candidate	scaled (dim. mixed)
V_{in}	Scaled candidate inner-loop storage for a future cascade proof	–
z_o	Position/current observer-error state	scaled (dim. mixed)
P, Q	Positive-definite Lyapunov-equation matrices	scaled (dim. mixed)
$\psi(s, r)$	Quantitative continuous-policy decay function	–
Ξ	Ideal continuous-policy position-decay expression	–
d_{out}	Combined outer-channel implementation perturbation	m s^{-1}
d_{al}	Velocity error due to imperfect direction alignment and angular-command approximation	m s^{-1}
$\bar{d}_{\text{out}}$	Ultimate bound on the combined outer-channel perturbation	m s^{-1}
m_a	Instantaneous norm of the unnormalised blend vector	–

Symbol	Description	Unit
m_v	Instantaneous heading–velocity alignment margin, distinct from the surge discriminant Δ_v	–
$m_{v,\min}$	Prescribed lower bound on the alignment margin	–
Δ_v	Surge-law discriminant of the smooth policy	$\text{m}^2 \text{s}^{-2}$
T	Nominal simulation horizon of a single run	s
T_0	Onset of the steady reporting window (typically $0.15 T$)	s
T_r	Onset of the steady reporting window (typically $0.15 T$; alias for T_0)	s
T_s	Settling time of the geometric path error to the reference tube	s
α_{conv}	Fitted exponential decay rate of $d_{\mathcal{P}}(t)$ on the convergence window	s^{-1}
s_e^{fin}	Terminal along-track error at run end	m
$J_{\Delta u}$	Integrated surge-command roughness on the steady window	$\text{m}^2 \text{s}^{-3}$
$\text{sat}^{\%}$	Actuator-saturation duty on the steady window	–
$\bar{\delta}$	Eventual uniform bound on the interface velocity error δ	m s^{-1}
$u_{\max}$	Maximum accepted outer-loop surge command	m s^{-1}
χ_{fb}	Runtime feasibility-fallback indicator	–
ξ	Complete inner-loop state, including plant and PI states	scaled (dim. mixed)
w_{in}	Command-rate input to the inner loop	scaled (dim. mixed)
S_{ξ}, S_w, S_{δ}	Scaling matrices for mixed-unit conditional estimates	–
L_e	Position-error scale used in a conditional storage function	m
$\underline{c}_{\text{in}}, \bar{c}_{\text{in}}$	Lower and upper storage-function bounds for a conditional inner proof	–
$a_{\text{in}}, b_{\text{in}}$	Inner-loop dissipation and input constants in a conditional estimate	scaled (dim. mixed)
c_{δ}	Conditional bound from inner state to velocity error	–
c_o	Uniform radial decay constant required by the conditional UUB estimate	–
a_*	Outer decay authority $c_o v_r^{\min}$	m s^{-1}
r_{UUB}	Conditional ultimate bound on the position-error norm	m
κ_o	Local outer-channel quadratic decay constant	s^{-1}
V_{cas}	Conditional composite cascade storage in scaled coordinates	–
$r_{\max}$	Bound used to define a compact position-error set	m
E	Compact-set bound on the position-error norm	m
$\mathcal{H}$	Hysteresis discrete state	–
$\mathcal{D}$	Admissible domain of the smooth guidance law	–
$\mathcal{E}_0$	Instantaneous zero-derivative set of the cruise policy	–
I	Identity matrix of compatible dimension	–
e_1	First Cartesian basis vector	–
A_o	Linear observer-error system matrix	scaled (dim. mixed)
z_{eff}	Effective observer-error state vector	scaled (dim. mixed)

Symbol	Description	Unit
Six-degree-of-freedom vehicle model		
$\mathbf{x}$	Complete vehicle state vector	scaled (dim. mixed)
$\boldsymbol{\eta}$	Inertial pose vector	scaled (dim. mixed)
$\boldsymbol{\nu}$	Body-fixed velocity vector	scaled (dim. mixed)
$\boldsymbol{\nu}_r$	Body-fixed water-relative velocity vector	scaled (dim. mixed)
$\boldsymbol{\nu}_c$	Body-fixed ocean-current velocity vector	scaled (dim. mixed)
$\boldsymbol{\nu}_c^b$	Body-fixed linear ocean-current vector	m s^{-1}
$\boldsymbol{\nu}_{c,\psi}^b$	Yaw-only body-frame current vector	m s^{-1}
$\bar{\boldsymbol{\nu}}_c$	Effective inertial current generated by a body-frame map	m s^{-1}
x, y, z	North, East, and Down coordinates	m
ϕ, θ, ψ	Roll, pitch, and yaw Euler angles	rad
u_b, v_b, w_b	Body-fixed surge, sway, and heave velocities	m s^{-1}
p_b, q_b, r_b	Body-fixed roll, pitch, and yaw rates	rad s^{-1}
$\boldsymbol{\omega}^b$	Body-fixed angular-velocity vector	rad s^{-1}
V_c	Horizontal ocean-current magnitude	m s^{-1}
β_{V_c}	Horizontal ocean-current bearing	rad
w_c	Vertical ocean-current component	m s^{-1}
$\mathbf{J}$	Body-to-inertial kinematic transformation matrix	–
$\mathbf{R}_{zyx}$	Body-to-inertial rotation matrix	–
$\mathbf{T}_{zyx}$	Euler-rate transformation matrix	–
$\mathbf{S}(\cdot)$	Skew-symmetric cross-product matrix operator	–
$\mathbf{M}$	Total inertia matrix	scaled (dim. mixed)
$\mathbf{M}_{RB}$	Rigid-body inertia matrix	scaled (dim. mixed)
$\mathbf{M}_A$	Added-mass inertia matrix	scaled (dim. mixed)
$\mathbf{C}$	Total Coriolis and centripetal matrix	scaled (dim. mixed)
$\mathbf{C}_{RB}$	Rigid-body Coriolis and centripetal matrix	scaled (dim. mixed)
$\mathbf{C}_A$	Added-mass Coriolis and centripetal matrix	scaled (dim. mixed)
$\mathbf{D}$	Hydrodynamic damping matrix	scaled (dim. mixed)
$\mathbf{D}_{\text{lin}}$	Linearised hydrodynamic damping matrix	scaled (dim. mixed)
$\mathbf{g}$	Hydrostatic restoring vector	scaled (dim. mixed)

Symbol	Description	Unit
τ	Generalised control force and moment vector	scaled (dim. mixed)
τ_{lift}	Hydrodynamic lift force and moment vector	scaled (dim. mixed)
τ_{cf}	Cross-flow force and moment vector	scaled (dim. mixed)
m_b	Vehicle mass	kg
L	Vehicle length	m
d_h	Hull diameter	m
I_x, I_y, I_z	Rigid-body moments of inertia	kg m ²
$\kappa_1^A, \kappa_2^A, \kappa_3^A$	Added-mass geometry factors	–
r_{44}	Roll added-inertia coefficient	–
W	Vehicle weight	N
z_g	Vertical centre-of-gravity coordinate	m
T_1, T_2, T_6	Linearised hydrodynamic time constants	s
T_{yaw}	Reduced yaw-channel time constant	s
K_{yaw}	Reduced yaw-channel gain	s ⁻¹
ζ_i	Modal damping ratio	–
ω_i	Modal natural frequency	rad s ⁻¹

Actuators and command interface

δ_r	Rudder angle	rad
δ_s	Stern-plane angle	rad
$\delta_{r,\text{max}}, \delta_{s,\text{max}}$	Rudder and stern-plane saturation limits	rad
n	Propeller rotational speed	rpm
X_{prop}	Propeller surge force	N
K_{prop}	Propeller roll moment	N m
X_r, Y_r	Rudder-induced forces	N
X_s, Z_s	Stern-plane-induced forces	N
x_r, x_s	Longitudinal actuator locations	m
t_{prop}	Propeller thrust-deduction coefficient	–
n_{max}	Maximum propeller rotational speed	rpm
U_{max}	Approximate maximum vehicle speed	m s ⁻¹
ω_{cmd}	Inertial angular-velocity reference	rad s ⁻¹
ω_{cmd}^b	Body-fixed angular-velocity reference	rad s ⁻¹
ω_x^b	Body-axis roll component of the angular reference	rad s ⁻¹
$\dot{\theta}_{\text{max}}$	Interface angular-rate limit	rad s ⁻¹
u_{cmd}	Surge command produced by the guidance layer	m s ⁻¹
u_{target}	Current-corrected surge reference	m s ⁻¹
q_{cmd}	Body pitch-rate command	rad s ⁻¹
r_{cmd}	Body yaw-rate command	rad s ⁻¹
e_u	Surge tracking error	m s ⁻¹
e_q	Pitch-rate tracking error	rad s ⁻¹
e_r	Yaw-rate tracking error	rad s ⁻¹

Symbol	Description	Unit
I_j	Generic PI integrator state for actuator channel j	scaled (dim. mixed)
$K_{p,n}, K_{i,n}$	Propeller-loop proportional and integral gains	scaled (dim. mixed)
$K_{p,q}, K_{i,q}$	Pitch-rate-loop proportional and integral gains	scaled (dim. mixed)
$K_{p,r}, K_{i,r}$	Yaw-rate-loop proportional and integral gains	scaled (dim. mixed)
n_{ff}	Propeller feedforward coefficient	rpm s m^{-1}
$\mathbf{v}_{\text{rel}}$	DVL water-relative inertial velocity	m s^{-1}
δ	Difference between realised and requested water-relative inertial velocity	m s^{-1}

Numerical performance measures

T	Simulation horizon	s
Δt	Numerical integration time step	s
$d\mathcal{P}$	Shortest geometric distance from the vehicle to the path	m
$d\mathcal{P}_{\text{RMS}}$	Root-mean-square geometric path distance	m
$d\mathcal{P}_{\text{max}}$	Maximum geometric path distance	m
$\tilde{\mathbf{v}}_{c,\text{RMS}}$	Root-mean-square current-estimation error	m s^{-1}
f_{fb}	Fraction of the reported horizon spent in fallback mode	–
J_{ω}	Integrated squared angular-command magnitude	$\text{rad}^2 \text{s}^{-1}$

Abbreviations

Abbreviation	Description
ALOS	Adaptive line-of-sight
AUV	Autonomous underwater vehicle
DVL	Doppler velocity log
ILOS	Integral line-of-sight
LOS	Line-of-sight
M-B	Modulation-B surge construction used by the selected continuous thesis law
MED	Mediterranean Conference on Control and Automation; here used to denote the published switching predecessor
MSS	Marine Systems Simulator
NED	North-East-Down
PF	Path following
PI	Proportional-integral
PP	Prescribed progression
REMUS	Remote Environmental Monitoring UnitS

Abbreviation	Description
VRP	Virtual reference point

Chapter 1

Introduction

1.1 Context and motivation

Autonomous underwater vehicles (AUVs) have become important tools for oceanographic survey, environmental monitoring, seabed mapping, infrastructure inspection, and long-duration observation. Many of these missions are naturally expressed as motion along a geometric route rather than arrival at a prescribed position at every instant. The vehicle is expected to converge to the route, remain close to it, and progress with acceptable speed and control effort. Path following is therefore a central problem in marine guidance, navigation, and control [1–3].

The problem is particularly demanding for a slender, torpedo-shaped AUV. Such vehicles are commonly actuated in surge and steered through yaw and pitch, while sway and heave are not independently controlled. Lateral and vertical motion must consequently be generated through the orientation of the forward velocity. This underactuation limits the directions and accelerations that can be realised instantaneously and makes a rigidly time-scheduled reference potentially incompatible with the vehicle motion.




Figure 1.1: Guidance geometry for a current-disturbed AUV: reference path $\mathcal{P}$, moving-target position $\mathbf{r}$, target-relative error $\mathbf{e} = \mathbf{r} - \mathbf{q}$, unit path-tangent vector $\mathbf{b}$, unit body-surge direction vector $\mathbf{h}$, and ocean-current vector $\mathbf{v}_c$.

The guidance geometry and disturbance convention are summarised in Figure 1.1. Ocean current introduces a second fundamental difficulty. The velocity over ground is the vector sum of the water-relative vehicle velocity and the current velocity. An uncompensated current may produce a persistent displacement from the desired route, alter the orientation

required for path convergence, and reduce the set of feasible inertial-velocity commands. The guidance problem must therefore address not only geometric convergence but also current estimation and velocity feasibility.

1.2 Research problem, aim, and objectives

This thesis considers an underactuated AUV moving in three dimensions under an unknown constant ocean current. The guidance model uses the vehicle inertial position, unit surge direction, water-relative surge speed, and a roll-free angular-velocity input. The reference is a moving target constrained to a regular spatial path with externally prescribed progression. Accordingly, the target-relative error and the geometric distance to the path are treated as distinct quantities.

The published MED predecessor uses two hysteretic geometric tasks: target approach and path-tangent alignment [4]. It provides the current observer, the basic direction mechanism, the surge-feasibility relation, and the overshoot treatment. This thesis addresses two remaining issues. First, the discrete switch between these two tasks is replaced by a smooth blended direction field that combines target approach with path-tangent alignment. Under the stated kinematic assumptions, this smooth formulation supports convergence results for the outer-loop guidance law. Second, the resulting ideal kinematic commands are mapped to the inputs accepted by a six-degree-of-freedom REMUS 100 implementation. The thesis also shows that direction alignment with nominal speed matching can still leave a steady along-path target error; removing this remaining error through a signed speed correction is the central kinematic guidance problem.

The aim is to develop a current-aware three-dimensional guidance law that does not require hydrodynamic parameter identification for its outer-loop design, and to evaluate its implementation on a REMUS 100/ Fossen model. The principal analytical contribution remains at the kinematic guidance level, while the dynamic model is used to assess command realisation, actuator limits, and plant-response effects.

1.3 Contributions and scope

The contributions are stated explicitly so that the boundary between the mathematical result and the implementation study remains clear:

Contribution 1: Kinematic formulation. A three-dimensional inertial-vector model is established for an underactuated AUV with a constant unknown current, including the moving-reference error, path tangent, current observer, and roll-free angular-velocity input.

Contribution 2: Along-path speed correction. The thesis identifies a limitation of direction-only guidance: even after the vehicle motion is aligned with the desired path-related direction, the vehicle may remain ahead of or behind the moving target. To address this, a continuous surge-speed correction

based on the signed along-path error is added and analysed in the ideal kinematic model.

Contribution 3: Current estimation and ideal-channel analysis. The current observer is shown to converge exponentially, the ideal direction error converges from every non-antipodal initial direction, and the continuous M-B law conditionally drives the moving-reference position error to zero in the reduced ideal kinematic model. The proof retains the observer and alignment transients explicitly and does not infer a physical-command or dynamic-cascade result from the ideal model.

Contribution 4: Kinematic–dynamic integration. A frame-consistent interface is defined for the REMUS 100/Fossen model, including full-attitude current transformation, inertial-to-body rate conversion, current-aware surge interpretation, PI loops, actuator limits, and the realised-command error used to assess the interface.

1.3.1 Validation programme and scope

The validation has two tracks: an ideal kinematic study that isolates the outer law and a six-degree-of-freedom REMUS 100/PI study that exposes bandwidth and actuator effects. Common paths, initial conditions, and current are used in both. The dynamic benchmark compares the thesis law with its MED predecessor, an exploratory state-dependent progression extension, and two recent three-dimensional methods under the shared thesis PI interface [5, 6]. The progression extension is numerically evaluated to expose its present limitations, but it is not part of the formal contribution or stability analysis.

1.3.2 Relation to the published two-task formulation

The published two-task formulation is the starting point rather than a second, unrelated controller. The division between the established material and the extensions developed in this thesis is:

The analytical results assume a sufficiently smooth reference, nonzero reference speed, constant current, accurate guidance-level position and direction information, and adequate surge and turning authority. The numerical study exercises the closed loop under a single nominal constant current and a single per-path initial condition, with actuator saturation, anti-windup, and rate clamping active; a parametric robustness envelope in the current magnitude and direction, in the initial-offset norm, and in measurement noise is *not* produced in this thesis and is recorded as future work in Chapter 6. Path planning, obstacle avoidance, navigation-filter design, communication architecture, and experimental certification are outside the scope of the thesis. A complete closed-loop stability proof for the six-degree-of-freedom interconnection is also beyond the present contribution; the dynamic study evaluates the command interface and plant response.

Table 1.1: Relationship between the published formulation and the developments in this thesis.

Element	Established two-task formulation	Development and analysis in this thesis	Traceability
Guidance geometry	Approach and path-tangent tasks with hysteresis	Smooth blended target direction and explicit cancellation domain	Theorems 3.3 and 3.4
Current treatment	Constant-current observer and feasible surge root	Time-scale approximation condition, frame-consistent dynamic implementation, and observer frame-consistency test	Theorems 3.2, 3.5 and 4.1
Longitudinal policy	Basic overshoot-aware surge modulation	Parallel analysis of Modulations A and B, hysteretic policy, and continuous policy	Theorems 3.11, 3.13 and 3.14
Validation level	Ideal kinematic guidance	Paired ideal-kinematic and REMUS 100/Fossen cascade studies	Chapter 5

1.4 Thesis organisation

Chapter 2 establishes the theoretical background and critically reviews current-aware path-following methods for underactuated marine vehicles. Chapter 3 develops the three-dimensional kinematic model, reconstructs the initial two-task controller, and presents the smooth direction and surge formulations together with their stability properties. Chapter 4 introduces the REMUS 100 dynamics and derives the interface between the kinematic commands and the dynamic plant. Chapter 5 defines the numerical validation and reports the kinematic and kinematic–dynamic results. Finally, Chapter 6 summarises the main findings and identifies the remaining analytical and experimental developments.

1.5 Chapter summary

This chapter has defined the formulation, contribution boundary, and validation structure. The thesis addresses a moving-reference tangential-equilibrium limitation through a coupled direction–speed law, while treating the REMUS 100 cascade as an implementation study rather than a dynamic-control theorem. Table 1.1 distinguishes these developments from the published predecessor.

Chapter 2

Background and State of the Art

2.1 Review scope and organising concepts

Path following for an autonomous underwater vehicle (AUV) combines geometric guidance, nonlinear control, state estimation, and marine-vehicle dynamics. The literature consequently spans ideal kinematic convergence, guidance coupled to heading and speed autopilots, and force–moment control of a hydrodynamic plant. Ocean current is likewise represented as a bounded bias, an adapted drift angle, or an explicitly estimated three-dimensional vector. These choices determine what a method can guarantee and whether two reported results are genuinely comparable [2, 3, 7].

The review is organised around the decisions that matter to this thesis: the distinction between trajectory tracking and geometric path following; reference progression and error definition; spatial underactuation; current representation and measurement requirements; and the boundary between kinematic guidance and dynamic realisation. Line-of-sight (LOS) and adaptive line-of-sight (ALOS) guidance, virtual-target methods, vector-field methods, explicit current observers, and the closest three-dimensional alternatives receive the most attention. Studies based primarily on full dynamic control, predictive control, finite-time convergence, or trajectory tracking are included only where they clarify this boundary. For example, the survey by Eissa et al. [8] provides useful broader context but is not treated as a direct baseline for the proposed guidance law.

2.2 Foundations of path following

2.2.1 Geometric paths, progression, and error

Trajectory tracking prescribes both the geometry and timing of a signal $\mathbf{r}_d(t)$. A geometric path prescribes only the set

$$\mathcal{P} = \{\mathbf{p}(\sigma) \in \mathbb{R}^3 : \sigma \in \mathcal{I}\}, \quad (2.1)$$

so that the path coordinate σ may be selected independently of clock time. This extra degree of freedom is important for an underactuated AUV because large lateral or vertical recovery need not be forced to follow an incompatible along-path schedule. Early marine formulations introduced path variables, virtual targets, and an explicit separation between geometric guidance and vehicle-specific control [9–13]. The principal milestones are summarised in Table 2.1; the purpose is orientation rather than an exhaustive chronology.

For a moving point $\mathbf{r}(t) = \mathbf{p}(\sigma(t))$, the target-relative error $\mathbf{e} = \mathbf{r} - \mathbf{q}$ measures the

Table 2.1: Condensed development of marine path-following control.

Period	Main development	Relevance here
Foundations	LOS geometry and path-frame error regulation	Practical planar guidance and cross-track regulation.
2000–2009	Path variables, virtual targets, guidance–control separation, and formal distinction from tracking [9–12]	Progression becomes a design variable; underactuated nonminimum-phase limitations are exposed.
2010–2019	ILOS/ALOS, current observers, cascade analysis, and experiments [14–16]	Steady-drift rejection and conditions for guidance–autopilot interconnection.
2020–2026	Spatial ALOS/vector fields/VRP, disturbance observers, and constraint-aware control	Three-dimensional geometry, explicit environment models, and implementation-aware comparison.

displacement from the vehicle to that particular moving point. It therefore depends on the chosen progression of the target along the path, denoted by $\sigma(t)$. By contrast, the geometric distance to the path is the shortest distance from $\mathbf{q}$ to any point on $\mathcal{P}$ and is independent of the chosen moving target. With unit tangent $\mathbf{b}$, a useful local decomposition is

$$\mathbf{e}_{\parallel} = \mathbf{b}^{\top} \mathbf{e}, \quad \mathbf{e}_{\perp} = (\mathbf{I} - \mathbf{b}\mathbf{b}^{\top}) \mathbf{e}. \quad (2.2)$$

Neither component is a substitute for a global closest-point distance on a highly curved path. A fair numerical study should therefore report geometric distance, target-relative error, and progression separately whenever their meanings differ.

2.2.2 Compact classification of formulations

The taxonomy in Table 2.2 condenses the main architectural choices identified in the path-following literature [7]. It also prevents a kinematic direction law, a path-coordinate controller, and a full dynamic controller from being compared as if they solved the same problem.

Table 2.2: Compact classification of path-following formulations.

Axis	Common alternatives	Choice in this thesis
Error and progression	Path/body/inertial error; projection or virtual target	Inertial target and direction vectors; externally supplied moving reference; geometric and target errors reported separately.
Controlled output	Cross-track error; path coordinate; course/attitude; velocity or force	Direction and feasible surge commands; no independent path-coordinate controller.
Architecture	Outer guidance with autopilots; integrated dynamic control	Kinematic outer layer; REMUS/PI cascade used for implementation assessment.
Current treatment	Robust, integral, adaptive-angle, or vector-observer compensation	Explicit constant inertial current-vector observer.

2.2.3 Underactuation and spatial geometry

A torpedo-shaped AUV is normally driven forward by surge propulsion and steered by changing its yaw and pitch. It does not directly command sideways or vertical translation; those motions are produced indirectly by pointing the forward velocity in a suitable direction. For this reason, three-dimensional path following is more difficult than a direct extension of a planar controller. Horizontal and vertical errors are coupled through the vehicle direction, path curvature, sideslip, and angle of attack.

Different geometric descriptions handle this coupling in different ways. Some methods describe errors in a path-attached frame, while others use inertial-vector formulations. Path-attached frames can be intuitive, but they may become ill-conditioned near singular geometric cases such as zero curvature or inflection points. Inertial-vector formulations avoid some of these coordinate issues, but they must state explicitly where the desired direction is well defined and where normalisation or cancellation can fail.

Ocean current adds another difficulty. The AUV produces motion relative to the surrounding water, but its motion over ground also includes current-induced drift. A command that is feasible in still water may therefore produce a different inertial motion when current is present. The guidance law must account for this effect when choosing both the direction of motion and the surge speed.

This distinction is especially important in three dimensions. A scalar crab- or sideslip-angle correction can compensate some horizontal drift, but it does not generally describe the full current vector, especially when vertical flow is present. The present thesis therefore uses an explicit inertial current-vector estimate to support the three-dimensional direction and speed calculations.

2.3 Guidance methods for underactuated marine vehicles

2.3.1 LOS and adaptive LOS

Line-of-sight (LOS) guidance uses a look-ahead point to generate a desired course and is attractive because its output interfaces naturally with conventional heading, depth, and speed autopilots [2]. A characteristic planar relation is

$$\chi_d = \chi_p - \arctan\left(\frac{e_y}{\Delta_{\text{LOS}}}\right), \quad (2.3)$$

where the look-ahead distance trades convergence speed against command smoothness and sensitivity. ILOS adds an internal state to remove a constant-current cross-track bias. Curved-path ILOS, alternative placements of the adaptive term, and combined guidance/low-level cascade results are developed in [14, 15, 17]. The last of these includes rigorous straight-line results in two and three dimensions and experimental validation of selected planar implementations; it also states the required speed–current feasibility conditions. Moe et al. instead combine planar LOS with adaptive feedback linearisation and sliding mode under absolute-velocity measurements [18], illustrating that similar geometric objectives can rest on different sensing and inner-loop assumptions.

Observer-based and adaptive LOS make the disturbance treatment more explicit. Belleter et al. combine a global path-coordinate update, current observer, and LOS-like law on curved planar paths [16]; switching enclosure-based LOS coordinates course, vertical steering, and surge on spatial straight paths [19]. Disturbance-matched ALOS estimates a crab-angle-related quantity and gives strong planar straight-path stability results [20]. Its NED amplitude–phase extension couples horizontal and vertical motion [21], while spherical ALOS removes restrictive vertical crab-angle assumptions and covers a class of curved spatial paths [5]. This makes spherical ALOS the closest recent LOS comparator here. Artificial-physics attraction combined with LOS provides complementary large-curvature experimental evidence, but remains planar and does not estimate a full current vector [22].

2.3.2 Virtual targets, vector fields, and offset points

Virtual-target methods treat progression as a closed-loop variable, allowing a target to slow during recovery and advance after cross-track convergence. State-dependent progression is a classical mechanism rather than a new concept [11, 13, 23]. Vector fields instead prescribe an attractive–tangential velocity direction in a neighbourhood of the path. The three-dimensional kinematic constructions of Indiveri and Zizzari align the available surge direction with a desired translational direction and extend it to constant unknown disturbances [24, 25]; they are the closest conceptual foundation for the direction-alignment mechanism used in this thesis.

Integral vector fields and dynamic inner controllers broaden that idea. ILOS and vector-field guidance can both remove geometric drift, but through different internal variables and tuning assumptions [26]. Yao and Wang combine a three-region spatial field with adaptive sliding-mode control, gaining dynamic disturbance handling at the cost of hydrodynamic dependence and many interacting gains [27]. Bibuli presents a smooth three-dimensional virtual-target law in a Serret–Frenet frame, but external disturbances are outside that study’s stated scope [28]. Xu, Fossen, and Guedes Soares provide a cascade precedent for a time-varying vector field and sliding-mode heading autopilot [29]; this supports the need for a separate interconnection argument and does not prove stability of the PI cascade used here.

Offset-point methods control a virtual point ahead of the vehicle, converting angular actuation into lateral velocity at that point. Degorre et al. apply this VRP principle to three-dimensional AUV path following with integral drift compensation [6]; their observation that particle speed may be adapted motivates, but does not itself design, a state-dependent progression law. The Handy Matrix provides a related allocation framework for underactuated tracking tasks [30]. More recent hand-position/NMPC work combines offset geometry, constraints, and sideslip adaptation [31]. The offset simplifies the controlled output, but affects angular demand, internal motion, and the relation between virtual-point and centre-of-mass error.

2.4 Current compensation and dynamic realisation

2.4.1 Integral adaptation and explicit observation

Current rejection has four distinct meanings in the reviewed literature: robust attenuation of a bounded disturbance; integral learning of a steering bias; adaptation of a crab/sideslip quantity; and explicit reconstruction of a current vector or lumped disturbance. Only the last aims to identify all inertial components. Frame consistency is therefore fundamental: a current that is genuinely constant in the inertial frame ($\dot{\mathbf{v}}_c = \mathbf{0}$) must have a body-frame representation obtained through the full three-axis roll–pitch–yaw rotation, whereas a yaw-only projection cannot represent the same inertial vector when the vehicle rolls or pitches [1]. The formal frame-consistent map is stated in Chapter 4.

Representative observers differ in both model dependence and measured signals. Lekkas and Fossen estimate planar current components with nonlinear adaptive observers [32]; Liang et al. use a kinematics-based NED current-vector observer inside a three-dimensional tracking controller [33]; Kim et al. reconstruct three-dimensional current with a hydrodynamic high-gain observer and report numerical and field evidence [34]; He et al. combine a current observer, Frenet–Serret error dynamics, LOS/backstepping, and a nonlinear disturbance observer [35]; and Ašanin et al. reconstruct sideslip and angle of attack for spatial LOS compensation [36]. These results are complementary rather than interchangeable.

The measurement architecture used in this thesis has two explicit levels. The ideal kinematic observer uses position, surge direction, the known command-model water-relative term $u\mathbf{h}$, and the reference velocity; it does not require a DVL measurement of $\dot{\mathbf{q}}$. In the six-degree-of-freedom study, the observer interface uses the full-attitude water-relative velocity supplied by the noise-free model as an ideal DVL-like signal. Noise, bias, delay, dropout, and bottom-track/water-track switching are outside the present validation. Lei et al.’s disturbance-observer/anti-windup design illustrates why lumped disturbances and actuator saturation must be diagnosed separately in a dynamic AUV implementation [37].

2.4.2 Guidance–dynamics interface

Kinematic guidance is portable because it does not depend on hydrodynamic coefficients, but its convergence proof assumes that commanded motion is realised. A dynamic plant adds inertia, Coriolis and centripetal terms, damping, restoring forces, actuator limits, and environmental loads [1]. Practical marine architectures therefore often separate an outer guidance law from faster inner autopilots [7]. Maurya, Aguiar, and Pascoal formalise this separation through input–output and small-gain arguments for straight paths [38]; their result demonstrates why an interface theorem needs assumptions beyond an outer-loop proof.

Model-based backstepping, sliding-mode, prescribed-performance, and finite-time designs address uncertainty and constraints at a different level. Examples include prescribed-performance tracking in time-varying currents [39] and an experimentally studied guidance–estimation–control cascade for depth tracking [40]. They motivate the separate command–realisation and saturation checks used here, but are not direct baselines for the present outer-loop kinematic spatial law, which does not use hydrodynamic parameters. In general,

transfer of a kinematic result to a specific inner loop requires input-to-state, small-gain, singular-perturbation, or composite-Lyapunov conditions [41–43]; without them, a dynamic simulation establishes implementation behaviour rather than inheritance of the nominal theorem.

2.5 Critical synthesis and thesis positioning

No guidance family dominates every criterion. LOS/ALOS is economical and mature, vector and direct-direction fields are natural for spatial geometry, offset-point methods exploit underactuated coupling, and integrated dynamic or predictive controllers represent constraints most directly. Their evidence is also different: a kinematic theorem, a plant-specific stability proof, a numerical transfer test, and an experiment answer different questions. A defensible comparison must therefore align objective, model level, progression, current representation, measurements, path geometry, and actuator assumptions. Table 2.3 applies that rule to the most relevant methods; other cited studies provide context rather than additional rows.

Table 2.3: Closest current-aware approaches and the position of this thesis.

Reference	Geometry/progression	Current and architecture	Main boundary
Caharija et al. [15]	2-D/3-D straight-path ILOS	Integral rejection; outer guidance plus low-level control	Strong cascade basis; straight-path and speed/current conditions.
Belleter et al. [16]	Planar curves; global path-coordinate update	Explicit current observer with LOS-like guidance	Global progression; horizontal-plane plant.
Yao and Wang [27]	Curved spatial integral vector field	Adaptive dynamic sliding-mode controller	Hydrodynamic dependence and tuning complexity.
Coates and Fossen [5]	Curved 3-D spherical ALOS	Adaptive drift parameters; kinematic guidance	Closest spatial LOS comparator; no full current vector.
Degorre et al. [6]	3-D path and offset virtual point	Integral drift compensation; kinematic VRP	Offset affects attitude and centre error.
He et al. [35]	Spatial Frenet–Serret progression	Current/disturbance observers with backstepping	Integrated but model-dependent design.
Published predecessor [4]	Moving reference; two hysteretic direction tasks	Explicit current-vector observer and feasible surge law	Direct baseline; switch and nonzero tangential equilibrium family.
Present thesis	Moving reference; smooth approach–tangent direction	Same observer family; continuous surge policy; REMUS/PI transfer study	Ideal-channel closure only; no autonomous progression or full dynamic-cascade theorem.

2.5.1 Local research gap

The research gap is local to the prescribed moving-reference formulation used here and in the published predecessor [4]. When progression is imposed externally and the surge command only matches the reference speed, exact direction alignment need not remove a nonzero along-tangent target error. Classical virtual-target laws with stabilising along-track feedback can instead absorb this error through progression [11, 13, 23]; ILOS, ALOS, and current observers address the different problem of disturbance-induced drift [14, 15, 20, 21]. The equilibrium derivation is treated in the controller development, where the required domain restrictions are introduced with the control law; see Section 3.3.

Within that scope, the thesis couples four elements: a full inertial current-vector observer, a smooth approach–tangent direction, an explicitly feasible surge construction with signed tangential-error modulation, and a command-realisation interface to a six-degree-of-freedom plant. The continuous blend alone is not claimed as novel, and the numerical comparison does not establish a method-independent ordering against controllers with different progression laws or inner-loop assumptions. Coates and Fossen [5] and Degorre et al. [6] are the direct guidance comparators; He et al. [35] provide complementary observer/dynamic context. The nominal theorem remains a constant-current, noise-free kinematic result, while the REMUS/PI study evaluates command realisation without claiming a full dynamic-cascade proof.

2.6 Chapter summary

This chapter has reviewed the path-following literature needed to position the thesis. It first distinguished trajectory tracking, moving-target formulations, and geometric path following, because the error to a moving target is not the same as the geometric distance to the path. It then reviewed the main guidance families used for underactuated marine vehicles, including line-of-sight/adaptive line-of-sight methods, virtual-target methods, vector-field methods, and current-aware observer-based approaches. These distinctions motivate the current-aware smooth direction and surge-speed correction developed in the next chapter.

Chapter 3

Three-Dimensional Kinematic Guidance

The guidance contribution is built from a vector kinematic model, target-relative and geometric path-error quantities, a constant-current observer, and a direction–speed control law. The published two-task controller provides the reference formulation. The proposed controller replaces the discrete task switch with a smooth approach–tangent direction and modifies the surge policy so that direction alignment and speed selection act on the same tangential-error channel.

Current estimation, direction alignment, and position convergence are kept as separate analytical layers. The required regularity and feasibility conditions are stated where they enter the design, and the non-ideal terms associated with the reduced angular-command inversion are retained explicitly. Longer algebraic derivations and observer/current-frame calculations are collected in Chapters A and B.

3.1 Problem Formulation and Kinematic Preliminaries

3.1.1 Path-following formulation

This section establishes the kinematic model underlying the guidance design [4]. The overall problem geometry is shown in Figure 1.1. The vehicle is required to follow a reference target that moves on a three-dimensional path. The target may be physical or virtual. Path following, target tracking, and trajectory tracking are distinguished in Chapter 2; the present work uses a moving target on the path as its kinematic representation, while retaining path following as the principal application.

Let $\mathbf{q} \in \mathbb{R}^3$ be the vehicle position and $\mathbf{h} \in \mathbb{S}^2$ its unit body-surge direction, where $\mathbb{S}^2 := \{\mathbf{x} \in \mathbb{R}^3 : \|\mathbf{x}\| = 1\}$. The target position is $\mathbf{r}(t) \in \mathbb{R}^3$, with $\dot{\mathbf{r}}(t) \neq \mathbf{0}$, acceleration $\ddot{\mathbf{r}}(t)$, and unit tangent

$$\mathbf{b} := \frac{\dot{\mathbf{r}}}{\|\dot{\mathbf{r}}\|}. \quad (3.1)$$

The vehicle kinematics are

$$\dot{\mathbf{q}} = u\mathbf{h} + \mathbf{v}_c, \quad (3.2)$$

$$\dot{\mathbf{h}} = \boldsymbol{\omega} \times \mathbf{h}, \quad \boldsymbol{\omega}^\top \mathbf{h} = 0, \quad (3.3)$$

where $u > 0$ is the water-relative surge speed and $\boldsymbol{\omega}$ is the roll-free angular-velocity input. The ocean current is modelled as an unknown constant vector expressed in the inertial frame,

$$\mathbf{v}_c \in \mathbb{R}^3, \quad \dot{\mathbf{v}}_c = \mathbf{0}. \quad (3.4)$$

Thus, the current is not represented by a scalar drift angle; it is a three-dimensional inertial velocity that enters the translational kinematics in (3.2).

The relative position, estimated current, estimated inertial velocity, and their normalized directions are

$$\mathbf{e} := \mathbf{r} - \mathbf{q}, \quad \mathbf{v}_1 := u\mathbf{h} + \hat{\mathbf{v}}_c, \quad \bar{\mathbf{e}} := \frac{\mathbf{e}}{\|\mathbf{e}\|}, \quad \bar{\mathbf{v}}_1 := \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|}, \quad (3.5)$$

whenever the respective vectors are nonzero. The two basic objectives are to approach the target when the displacement is large and to align the inertial velocity with the path tangent near the route. The formulation in Section 3.3 preserves this physical interpretation but replaces the directional switch with a smooth blend.

The exact target-relative error derivative is

$$\dot{\mathbf{e}} = \dot{\mathbf{r}} - u\mathbf{h} - \mathbf{v}_c = \dot{\mathbf{r}} - \mathbf{v}_1 - \tilde{\mathbf{v}}_c, \quad \tilde{\mathbf{v}}_c := \mathbf{v}_c - \hat{\mathbf{v}}_c. \quad (3.6)$$

Whenever $\dot{\mathbf{e}}$ is required by the implementable feedforward calculation, the unknown $\mathbf{v}_c$ in this equation is replaced by its observer estimate $\hat{\mathbf{v}}_c$. No additional control term is introduced by this substitution.

3.1.2 Angular velocity of the geometric directions

For any differentiable nonzero vector $\mathbf{x}$, its normalized direction $\bar{\mathbf{x}} = \mathbf{x}/\|\mathbf{x}\|$ satisfies

$$\dot{\bar{\mathbf{x}}} = \boldsymbol{\omega}_x \times \bar{\mathbf{x}}, \quad \boldsymbol{\omega}_x = \frac{\bar{\mathbf{x}} \times \dot{\bar{\mathbf{x}}}}{\|\bar{\mathbf{x}}\|}. \quad (3.7)$$

Applying (3.7) to the error, estimated inertial velocity, and path tangent gives

$$\boldsymbol{\omega}_e = \frac{\bar{\mathbf{e}} \times \dot{\bar{\mathbf{e}}}}{\|\bar{\mathbf{e}}\|}, \quad \boldsymbol{\omega}_{v_1} = \frac{\bar{\mathbf{v}}_1 \times \dot{\bar{\mathbf{v}}}_1}{\|\bar{\mathbf{v}}_1\|}, \quad \boldsymbol{\omega}_b = \frac{\mathbf{b} \times \dot{\mathbf{r}}}{\|\dot{\mathbf{r}}\|}. \quad (3.8)$$

These vector expressions avoid yaw–pitch parameterisation singularities and give a common way to describe how the relevant directions change in space. The term $\boldsymbol{\omega}_e$ describes the change of the target-relative error direction, $\boldsymbol{\omega}_{v_1}$ describes the change of the estimated inertial-velocity direction, and $\boldsymbol{\omega}_b$ describes the change of the path tangent. This common vector notation is used both in the two-task predecessor and in the smooth-blending formulation, without introducing a local yaw–pitch coordinate frame.

3.1.3 State and measurement conventions

The notation distinguishes physical quantities, observer quantities, and commands so that an estimate is not silently treated as a measurement or as the true environmental state; Table 3.1 lists the categories. The distinction is also preserved when the kinematic law is connected to the dynamic plant.

Table 3.1: State, estimate, command, and realisation conventions used in the thesis.

Category	Principal symbols	Interpretation
Physical/reference	$\mathbf{q}, \mathbf{h}, \mathbf{v}_c, \mathbf{e}, \mathbf{r}, \dot{\mathbf{r}}$	True vehicle, environment, and reference quantities; the constant current $\mathbf{v}_c$ is unknown to the controller.
Estimated/derived	$\hat{\mathbf{e}}, \hat{\mathbf{v}}_c, \mathbf{v}_1, \bar{\mathbf{v}}_1, \bar{\mathbf{a}}$	Quantities formed by the observer or guidance law and used by the outer loop.
Commanded	$u_{\text{cmd}}, \omega_{\text{cmd}}$	Ideal kinematic requests; treated as realised in the kinematic proof and as inputs to the interface in the dynamic study.
Measured/realised	$\mathbf{v}_{\text{rel}}, \delta$	Water-relative velocity supplied by the dynamic model/DVL and the difference between realised and requested motion.

3.1.4 Standing assumptions

The kinematic analysis of this chapter and the next rests on the following standing hypotheses:

- *Kinematic model only.* Hydrodynamic effects are neglected; the analysis operates at the level of the kinematic equations (3.2) and (3.3).
- *Constant unknown current.* The ocean current satisfies $\dot{\mathbf{v}}_c = \mathbf{0}$ with $\|\mathbf{v}_c\| \leq V_c^{\max} < \|\dot{\mathbf{r}}\|$, so the reference moves strictly faster than the maximum current magnitude.
- *Perfect kinematic measurements.* The vehicle position $\mathbf{q}$ and heading $\mathbf{h}$ are available to the outer loop; velocity $\dot{\mathbf{q}}$ is not required.
- *Reference regularity.* The target $\mathbf{r}(t) \in C^2$ has bounded first and second derivatives (Theorem 3.1).
- *Underactuation constraint.* The angular velocity is orthogonal to the heading, $\omega^\top \mathbf{h} = 0$.

The one bound that is used quantitatively in the surge feasibility argument of Section 3.2 is stated as a formal assumption:

Assumption 3.1 (Reference regularity). The target is twice continuously differentiable and satisfies

$$0 < v_r^{\min} \leq \|\dot{\mathbf{r}}(t)\| \leq v_r^{\max}, \quad \|\ddot{\mathbf{r}}(t)\| \leq a_r^{\max}. \quad (3.9)$$

The condition $v_r^{\min} > V_c^{\max}$ means that the prescribed reference motion remains faster than the largest current magnitude allowed by the model. This avoids a loss of feasibility in the current-compensating speed calculation introduced later.

3.1.5 Path representation and error quantities

For a geometric path $\mathbf{r}_\mathcal{P}(\sigma) \in \mathbb{R}^3$, a moving target can be written as

$$\mathbf{r}(t) = \mathbf{r}_\mathcal{P}(\sigma(t)), \quad \dot{\mathbf{r}} = \mathbf{r}'_\mathcal{P}(\sigma)\dot{\sigma}. \quad (3.10)$$

The present study assumes that $\mathbf{r}(t)$, $\dot{\mathbf{r}}(t)$, and $\ddot{\mathbf{r}}(t)$ are supplied to the kinematic controller. It does not yet prescribe the progression law $\dot{\sigma}$; designing a virtual-target planner is a later extension. This choice keeps the present thesis focused on the current-aware guidance and stability problem while preserving a direct connection to geometric path following.

For reporting and interpretation, the target-relative error can be decomposed with respect to the tangent at the moving target,

$$s_e := \mathbf{e}^\top \mathbf{b}, \quad \mathbf{e}_\perp := (\mathbf{I} - \mathbf{b}\mathbf{b}^\top)\mathbf{e}. \quad (3.11)$$

Here s_e measures whether the vehicle is ahead of or behind the moving target along the local tangent, while $\mathbf{e}_\perp$ is the component orthogonal to that tangent at the same target point. These are local target-relative quantities. For a curved path, $\|\mathbf{e}_\perp\|$ is not, in general, the shortest geometric distance from the vehicle to the path; that distance is evaluated separately in the simulation chapter.

The symbols $e_\parallel$ in Equation (2.2), s_e in (3.11), e_ℓ in the initial formulation, and s in the smooth formulation denote the same signed projection $\mathbf{e}^\top \mathbf{b}$ whenever they refer to the common moving target. The different subscripts are retained only where they clarify the source formulation; s is used for the comparative analysis.

3.1.6 Control objective and scope of the proof

The immediate goal is regulation of the target-relative error while advancing along the reference path. A geometric path-following evaluation must additionally report the distance to the path, because a virtual target may intentionally remain ahead of the vehicle. The mathematical results in the following controller sections apply to a constant current and a regular feasible reference. Measurement noise, sampling, saturation, parameter uncertainty, and time-varying currents are robustness tests, not implicit extensions of the constant-current theorem.

3.2 Initial Two-Task Guidance Formulation

The formulation introduced in [4] provides the starting point for the smooth guidance design in Section 3.3. Its architecture has three defining elements: a constant-current observer, two geometrically distinct alignment tasks connected by hysteresis, and a surge policy that handles both approach and longitudinal overshoot.

3.2.1 Constant-current observer

For the observer design, the reference target $\mathbf{r}(t)$ and the vehicle position $\mathbf{q}(t)$ are assumed available in the inertial frame. Hence the target-relative error

$$\mathbf{e} = \mathbf{r} - \mathbf{q}$$

is available as a computed navigation quantity; it is not a separate physical sensor measurement. The observer also uses the known water-relative surge velocity $u\mathbf{h}$ and does not require $\dot{\mathbf{q}}$ or

a direct measurement of the ocean current.

The unknown constant current can be estimated either from the position dynamics or from the target-error dynamics. With outputs denoted by y_q and y_e , the two equivalent design models are

$$\dot{\mathbf{q}} = u\mathbf{h} + \mathbf{v}_c, \quad \dot{\mathbf{v}}_c = \mathbf{0}, \quad y_q = \mathbf{q}, \quad (3.12)$$

$$\dot{\mathbf{e}} = \dot{\mathbf{r}} - u\mathbf{h} - \mathbf{v}_c, \quad \dot{\mathbf{v}}_c = \mathbf{0}, \quad y_e = \mathbf{e}. \quad (3.13)$$

The present thesis follows the second route. This choice is natural because the guidance law, the switching variables, and the stability analysis are written directly in terms of the target-relative error. The “measured- $\mathbf{e}$ ” formulation should therefore be understood as a computed-output formulation based on the measured position $\mathbf{q}$ and the generated reference $\mathbf{r}$, not as an additional physical sensor. The full position-output and target-error-output observer derivations, including the Luenberger error dynamics and gain conditions, are given in Section B.1.

Let $\hat{\mathbf{e}}$ be an observer state associated with the target-relative error and let $\hat{\mathbf{v}}_c$ be the estimated current. Define

$$\tilde{\mathbf{e}} := \mathbf{e} - \hat{\mathbf{e}}, \quad \tilde{\mathbf{v}}_c := \mathbf{v}_c - \hat{\mathbf{v}}_c. \quad (3.14)$$

The Luenberger-type observer is

$$\dot{\hat{\mathbf{e}}} = \dot{\mathbf{r}} - u\mathbf{h} - \hat{\mathbf{v}}_c + k_C \tilde{\mathbf{e}}, \quad (3.15)$$

$$\dot{\hat{\mathbf{v}}}_c = -k_D \tilde{\mathbf{e}}, \quad k_C, k_D > 0. \quad (3.16)$$

Subtracting (3.15)–(3.16) from the error dynamics gives

$$\begin{bmatrix} \dot{\tilde{\mathbf{e}}} \\ \dot{\tilde{\mathbf{v}}}_c \end{bmatrix} = \begin{bmatrix} -k_C \mathbf{I} & -\mathbf{I} \\ k_D \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{e}} \\ \tilde{\mathbf{v}}_c \end{bmatrix}. \quad (3.17)$$

The system contains three identical second-order blocks with characteristic polynomial $s^2 + k_C s + k_D$. Thus the estimation error converges globally and exponentially for $k_C, k_D > 0$ [44]. Critical damping is obtained with $k_D = k_C^2/4$.

Proposition 3.2 (Exponential convergence of the constant-current observer). *Under the kinematic model (3.2), a constant inertial current, and gains $k_C, k_D > 0$, the origin $(\tilde{\mathbf{e}}, \tilde{\mathbf{v}}_c) = (\mathbf{0}, \mathbf{0})$ of (3.17) is globally exponentially stable. In the critically damped case $k_D = k_C^2/4$, every scalar observer channel has a repeated pole at $-k_C/2$.*

The Lyapunov and modal proof is given in Section B.1. The result is independent of vehicle attitude because both the plant and observer use the same inertial representation of the current. The frame-consistency requirement needed to preserve this property in the REMUS model is addressed in Section 4.2.

3.2.2 Two tasks and hysteresis logic

The controller changes objective according to the target-relative distance. In *Task 1* (approach), active when $\|\mathbf{e}\| > \rho_{\text{out}}$, the estimated inertial velocity direction $\bar{\mathbf{v}}_1$ is aligned with $\bar{\mathbf{e}}$. In *Task 2* (path traversal), activated after $\|\mathbf{e}\| \leq \rho_{\text{in}}$ and retained while $\|\mathbf{e}\| \leq \rho_{\text{out}}$, $\bar{\mathbf{v}}_1$ is aligned with the path tangent $\mathbf{b}$. The thresholds satisfy

$$0 < \rho_{\text{in}} < \rho_{\text{out}}. \quad (3.18)$$

This pair of concentric spheres prevents chattering: Task 2 is entered only through the inner sphere and Task 1 is re-entered only outside the outer sphere.

3.2.3 Base surge speed and feasibility

The reference-speed matching condition $\|\mathbf{u}\mathbf{h} + \hat{\mathbf{v}}_c\| = \|\dot{\mathbf{r}}\|$ gives

$$\Delta = (\mathbf{h}^\top \hat{\mathbf{v}}_c)^2 - \|\hat{\mathbf{v}}_c\|^2 + \|\dot{\mathbf{r}}\|^2, \quad (3.19)$$

$$u_{\pm} = -\mathbf{h}^\top \hat{\mathbf{v}}_c \pm \sqrt{\Delta}. \quad (3.20)$$

When $\|\dot{\mathbf{r}}\| > \|\hat{\mathbf{v}}_c\|$, $\Delta > 0$, $u_+ > 0$, and $u_- < 0$. The controller uses u_+ , because the AUV is assumed to maintain positive surge. During the observer transient, a negative discriminant is handled by retaining a safe previously feasible value u_d rather than evaluating an invalid root. The same quadratic calculation, including the modulated-speed version used later, is detailed in Section A.4.

3.2.4 Task 1: target approach

For Task 1, consider

$$V_1 := \frac{1}{2}(1 - \bar{\mathbf{e}}^\top \bar{\mathbf{v}}_1) \in [0, 1]. \quad (3.21)$$

The angular velocity of the error direction is

$$\boldsymbol{\omega}_e = \frac{\bar{\mathbf{e}} \times \dot{\mathbf{e}}}{\|\mathbf{e}\|}. \quad (3.22)$$

In implementation, $\dot{\mathbf{e}}$ is evaluated from (3.6) with $\mathbf{v}_c$ replaced by $\hat{\mathbf{v}}_c$. Task 1 is active outside the hysteresis radius, so the denominator in (3.22) is bounded away from zero.

The desired angular velocity of the estimated inertial-velocity direction is

$$\boldsymbol{\omega}_{\bar{\mathbf{v}}_1}^{\text{des}} = \boldsymbol{\omega}_e + k_1(\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}}), \quad k_1 > 0, \quad (3.23)$$

The derivative of V_1 is obtained by differentiating $\bar{\mathbf{e}}^\top \bar{\mathbf{v}}_1$, substituting (3.23), and using the scalar triple-product identity. The full calculation is given in Section A.5. On the converged-observer manifold, this gives

$$\dot{V}_1 = -\frac{k_1}{2} \|\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}}\|^2 \leq 0. \quad (3.24)$$

On the ideal direction manifold with a converged current estimate, and except for the unstable antipodal direction, $\bar{\mathbf{v}}_1$ converges to $\bar{\mathbf{e}}$. This exceptional set is consistent with the general topological obstruction to global continuous asymptotic stabilisation on the sphere; the cited results do not establish a stronger claim about every possible continuous or time-varying field [45, 46].

The alignment result has a direct implication for target approach. Once the ideal Task 1 alignment manifold has been reached and the current estimate has converged, the kinematic velocity satisfies

$$\dot{\mathbf{q}} = \mathbf{v}_1, \quad \mathbf{v}_1 = \|\mathbf{v}_1\| \bar{\mathbf{e}}.$$

Moreover, $\dot{\mathbf{r}} = \|\dot{\mathbf{r}}\| \mathbf{b}$. Hence the target-distance derivative satisfies

$$\begin{aligned} \frac{d}{dt} \|\mathbf{e}\| &= \bar{\mathbf{e}}^\top \dot{\mathbf{e}} \\ &= \bar{\mathbf{e}}^\top (\dot{\mathbf{r}} - \dot{\mathbf{q}}) \\ &= \bar{\mathbf{e}}^\top (\|\dot{\mathbf{r}}\| \mathbf{b} - \|\mathbf{v}_1\| \bar{\mathbf{e}}) \\ &= \|\dot{\mathbf{r}}\| (\bar{\mathbf{e}}^\top \mathbf{b}) - \|\mathbf{v}_1\| \leq \|\dot{\mathbf{r}}\| - \|\mathbf{v}_1\|. \end{aligned} \quad (3.25)$$

The vehicle therefore approaches whenever the realised inertial speed is larger than the target speed on this ideal aligned manifold.

3.2.5 Task 2: path-tangent alignment

For Task 2, use

$$V_2 := \frac{1}{2}(1 - \bar{\mathbf{v}}_1^\top \mathbf{b}) \in [0, 1] \quad (3.26)$$

and select

$$\omega_{\mathbf{v}_1}^{\text{des}} = \omega_b + k_2(\bar{\mathbf{v}}_1 \times \mathbf{b}), \quad k_2 > 0. \quad (3.27)$$

Then

$$\dot{V}_2 = -\frac{k_2}{2} \|\bar{\mathbf{v}}_1 \times \mathbf{b}\|^2 \leq 0. \quad (3.28)$$

This is the same alignment calculation as in Task 1, with $\bar{\mathbf{e}}$ replaced by $\mathbf{b}$; the full vector-identity derivation is given in Section A.5. This establishes alignment with the tangent from every initial direction except the antipodal configuration, on the ideal direction manifold with a converged observer. It does not, by itself, prove stability of arbitrary switching between the two tasks; the hysteresis mechanism and simulation evidence address switching behaviour operationally.

3.2.6 Physical angular-velocity command

The desired $\omega_{\mathbf{v}_1}^{\text{des}}$ must be converted to the physical input ω . Neglecting the surge-acceleration and current-estimate derivative terms in

$$\dot{\mathbf{v}}_1 = u(\omega \times \mathbf{h}) + \dot{u}\mathbf{h} + \hat{\mathbf{v}}_c, \quad (3.29)$$

the reduced angular-command inversion used by the controller is

$$\boldsymbol{\omega} = \frac{\|\mathbf{v}_1\|}{u(\mathbf{h}^\top \hat{\mathbf{v}}_1)} (\mathbf{I} - \mathbf{h}\mathbf{h}^\top) \boldsymbol{\omega}_{v_1}^{\text{des}}. \quad (3.30)$$

The denominator is positive when $u > \|\hat{\mathbf{v}}_c\|$, and the factor $(\mathbf{I} - \mathbf{h}\mathbf{h}^\top)$ enforces the roll-free condition $\boldsymbol{\omega}^\top \mathbf{h} = 0$. Equation (3.30) is a reduced relation under the stated time-scale approximation, not an exact inversion of the full derivative. The corresponding reduced-inversion derivation and omitted approximation-error terms are reported in Section A.2. Angular-rate saturation is applied after this calculation in simulations.

3.2.7 Unified longitudinal overshoot policy

Define the longitudinal error $e_\ell := \mathbf{e}^\top \mathbf{b}$. A positive value means that the vehicle is behind the reference; a negative value means that it has overtaken it. The initial surge policy is

$$u = \begin{cases} \alpha u_+, & \text{Task 1, } e_\ell \geq 0, \Delta \geq 0, \\ \beta u_+, & \text{Task 1 or Task 2, } e_\ell < 0, \Delta \geq 0, \\ u_+, & \text{Task 2, } e_\ell \geq 0, \Delta \geq 0, \\ u_d, & \Delta < 0, \end{cases} \quad \alpha > 1, \quad 0 < \beta < 1. \quad (3.31)$$

When $e_\ell < 0$, reducing the surge below the matching value lets the target move ahead naturally. This avoids an unnecessary heading reversal in Task 1 and repeated sphere exits in Task 2.

3.2.8 Properties of the initial formulation

The two-task formulation provides exponential current-estimation convergence and nominal directional alignment in each active task from every non-antipodal initial direction, on the ideal direction manifold with a converged observer. It also gives explicit speed feasibility and a practical overshoot mechanism. The smooth formulation retains this architecture while removing the directional switch, exposing the remaining longitudinal equilibrium set, and refining the position and perturbation analysis.

3.3 Smooth Direction Design and Stability Analysis

The smooth formulation builds on the two-task controller in Section 3.2. It retains the same kinematic plant, constant-current observer, and physical angular-velocity input. Its change is geometric: the far-field attraction direction $\bar{\mathbf{e}}$ and near-field path tangent $\mathbf{b}$ are combined into one desired direction. Directional hysteresis is thereby removed; hysteresis is retained only as an optional surge policy.

3.3.1 Smooth blended target direction

Let $\lambda > 0$ be a design length and define

$$p := \frac{\lambda}{\lambda + \|\mathbf{e}\|} \in (0, 1]. \quad (3.32)$$

For $\mathbf{e} \neq \mathbf{0}$, the weight p gives a continuous interpolation between target approach and path-tangent alignment. The unnormalised blended direction is defined as

$$\mathbf{a} := (1 - p)\bar{\mathbf{e}} + p\mathbf{b}, \quad (3.33)$$

$$\bar{\mathbf{a}} := \frac{\mathbf{a}}{\|\mathbf{a}\|}. \quad (3.34)$$

Equivalently, for $\mathbf{e} \neq \mathbf{0}$,

$$\mathbf{a} = \frac{\mathbf{e} + \lambda\mathbf{b}}{\lambda + \|\mathbf{e}\|}. \quad (3.35)$$

As $\|\mathbf{e}\|$ becomes large, $p \rightarrow 0$ and the direction approaches the target-approach direction $\bar{\mathbf{e}}$. As $\mathbf{e} \rightarrow \mathbf{0}$, $p \rightarrow 1$ and the continuous extension is $\mathbf{a} = \mathbf{b}$. Its unit-direction form is

$$\bar{\mathbf{a}} = \frac{\mathbf{e} + \lambda\mathbf{b}}{\|\mathbf{e} + \lambda\mathbf{b}\|}, \quad (3.36)$$

whenever the denominator is nonzero. Geometrically, (3.36) points from the vehicle toward a virtual look-ahead target located a distance λ ahead of the moving target along the path tangent.

Lemma 3.3 (Blend norm and singularity). *For $\mathbf{e} \neq \mathbf{0}$, let $c := \bar{\mathbf{e}}^\top \mathbf{b} \in [-1, 1]$. Then*

$$\|\mathbf{a}\|^2 = 1 - 2p(1 - p)(1 - c), \quad (3.37)$$

and $\mathbf{a} = \mathbf{0}$ if and only if $\mathbf{e} = -\lambda\mathbf{b}$. At $\mathbf{e} = \mathbf{0}$, (3.37) is understood by continuity and gives $\|\mathbf{a}\| = 1$.

The complete algebra and the continuous virtual-target extension are given in Section A.3.

Assumption 3.4 (Blend regularity). There is a constant $a_{\min} \in (0, 1]$ such that $\|\mathbf{a}(t)\| \geq a_{\min}$ along the closed-loop trajectory. This assumption is a trajectory-regularity condition, not a consequence of the curvature bound alone. The algebraic result in Theorem 3.3 identifies the only cancellation configuration, but excluding that configuration for a complete trajectory requires a separately certified positively invariant set (or a direct numerical check for each run). The implementation therefore logs $\|\mathbf{a}(t)\|$ and must define a bounded fallback command if the prescribed margin is violated; the nominal results below do not apply during such an interval. No implicit regularisation of the normalisation is used.

This assumption makes the isolated virtual-target cancellation explicit rather than hiding it in numerical regularisation. The smooth law is therefore defined on the admissible domain

$$\mathcal{D} := \{(\mathbf{e}, \mathbf{h}, \hat{\mathbf{v}}_c) : \mathbf{e} + \lambda\mathbf{b} \neq \mathbf{0}, \|\mathbf{v}_1\| > 0, \mathbf{h}^\top \bar{\mathbf{v}}_1 > 0\}. \quad (3.38)$$

All smoothness and convergence statements in this section are restricted to trajectories that remain in $\mathcal{D}$; no global avoidance claim is made for the isolated cancellation configuration. Since $\|\mathbf{a}\| \leq 1$, the normalisation factor

$$\kappa_a := \frac{1}{\|\mathbf{a}\|} \in [1, a_{\min}^{-1}] \quad (3.39)$$

is bounded and converges to one on the tangential set $\bar{\mathbf{e}} = \mathbf{b}$.

The rotation rate of the blended target is obtained from the quantities already defined above. For $\mathbf{e} \neq \mathbf{0}$,

$$\dot{p} = -\frac{\lambda}{(\lambda + \|\mathbf{e}\|)^2} \bar{\mathbf{e}}^\top \dot{\mathbf{e}}, \quad (3.40)$$

$$\dot{\mathbf{a}} = (1 - p)(\boldsymbol{\omega}_e \times \bar{\mathbf{e}}) + p(\boldsymbol{\omega}_b \times \mathbf{b}) + \dot{p}(\mathbf{b} - \bar{\mathbf{e}}), \quad (3.41)$$

$$\boldsymbol{\omega}_{\bar{\mathbf{a}}} = \frac{\bar{\mathbf{a}} \times \dot{\mathbf{a}}}{\|\mathbf{a}\|}. \quad (3.42)$$

The equivalent virtual-target expression is used for the continuous evaluation at $\mathbf{e} = \mathbf{0}$; it fails only at the cancellation point excluded from $\mathcal{D}$. In implementation, the unknown $\mathbf{v}_c$ inside $\dot{\mathbf{e}}$ is replaced by $\hat{\mathbf{v}}_c$, exactly as stated after (3.6). This is an observer-based evaluation of the same feedforward formula, not an additional controller or a separately named error channel.

3.3.2 Surge speed and direction command

The estimated inertial velocity generated by the kinematic command is

$$\mathbf{v}_1 = u\mathbf{h} + \hat{\mathbf{v}}_c. \quad (3.43)$$

The nominal surge is chosen so that this vector has the same magnitude as the moving reference velocity,

$$\|\mathbf{v}_1\| = \|\dot{\mathbf{r}}\|. \quad (3.44)$$

Substituting (3.43) into (3.44) gives a scalar quadratic in u . Since $\|\mathbf{h}\| = 1$,

$$u^2 + 2(\mathbf{h}^\top \hat{\mathbf{v}}_c)u + \|\hat{\mathbf{v}}_c\|^2 - \|\dot{\mathbf{r}}\|^2 = 0. \quad (3.45)$$

The physically admissible solution is the positive root

$$u_{\text{nom}} = -\mathbf{h}^\top \hat{\mathbf{v}}_c + \sqrt{(\mathbf{h}^\top \hat{\mathbf{v}}_c)^2 - \|\hat{\mathbf{v}}_c\|^2 + \|\dot{\mathbf{r}}\|^2}. \quad (3.46)$$

The stronger condition $\|\dot{\mathbf{r}}\| > \|\hat{\mathbf{v}}_c\|$ guarantees that the positive root exists for every heading. The same calculation, including the modulated-speed case, is reported in Section A.4.

The desired angular velocity of the estimated inertial-velocity direction is

$$\boldsymbol{\omega}_{\bar{\mathbf{v}}}^{\text{des}} = \boldsymbol{\omega}_{\bar{\mathbf{a}}} + k(\bar{\mathbf{v}}_1 \times \bar{\mathbf{a}}), \quad k > 0. \quad (3.47)$$

This is the desired rate of change of the unit vector $\bar{\mathbf{v}}_1 = \mathbf{v}_1 / \|\mathbf{v}_1\|$; it is not yet the physical

angular-velocity input of the vehicle. Indeed, differentiating (3.43) gives

$$\dot{\mathbf{v}}_1 = \dot{u}\mathbf{h} + u(\boldsymbol{\omega} \times \mathbf{h}) + \dot{\hat{\mathbf{v}}}_c. \quad (3.48)$$

The implemented command uses a reduced direction-rate inversion in which the terms $\dot{u}\mathbf{h}$ and $\dot{\hat{\mathbf{v}}}_c$ are not used while computing $\boldsymbol{\omega}$. Thus the inversion is based on

$$\dot{\mathbf{v}}_1 \simeq u(\boldsymbol{\omega} \times \mathbf{h}). \quad (3.49)$$

Solving the projected unit-direction relation associated with (3.49) gives the roll-free physical angular command

$$\boldsymbol{\omega} = \frac{\|\mathbf{v}_1\|}{u(\mathbf{h}^\top \bar{\mathbf{v}}_1)} (\mathbf{I} - \mathbf{h}\mathbf{h}^\top) \boldsymbol{\omega}_{\bar{\mathbf{v}}}^{\text{des}}. \quad (3.50)$$

The projection $(\mathbf{I} - \mathbf{h}\mathbf{h}^\top)$ removes the component parallel to $\mathbf{h}$ and therefore enforces the roll-free constraint $\boldsymbol{\omega}^\top \mathbf{h} = 0$. The denominator is positive when $u > \|\hat{\mathbf{v}}_c\|$, because then the commanded inertial velocity has a positive component along the surge direction. The derivation and the corresponding approximation terms are separated in the next subsection and detailed in Section A.2. Therefore (3.50) should be read as the adopted reduced angular-command inversion, not as an exact inversion of the full derivative (3.48).

3.3.3 Reduced mapping from desired direction rate to vehicle angular rate

The command (3.50) is the implementable roll-free approximation of the desired direction-rate law. The only point of the present subsection is to make explicit what this approximation leaves outside the ideal direction theorem: a projection term caused by the constraint $\boldsymbol{\omega}^\top \mathbf{h} = 0$, and a time-scale term caused by omitting $\dot{u}\mathbf{h} + \dot{\hat{\mathbf{v}}}_c$ while computing the command. Direct substitution of (3.50) into the reduced direction kinematics gives

$$\boldsymbol{\omega}_{\bar{\mathbf{v}}} \simeq \boldsymbol{\omega}_{\bar{\mathbf{v}}}^{\text{des}} - \mathbf{h} \frac{\bar{\mathbf{v}}_1^\top \boldsymbol{\omega}_{\bar{\mathbf{v}}}^{\text{des}}}{\mathbf{h}^\top \bar{\mathbf{v}}_1} + \boldsymbol{\varepsilon}(t), \quad (3.51)$$

where the projection term and the time-scale approximation term are, respectively,

$$\mathbf{d}_{\text{proj}} := -\mathbf{h} \frac{\bar{\mathbf{v}}_1^\top \boldsymbol{\omega}_{\bar{\mathbf{v}}}^{\text{des}}}{\mathbf{h}^\top \bar{\mathbf{v}}_1}, \quad (3.52)$$

$$\boldsymbol{\varepsilon}(t) := \frac{1}{\|\mathbf{v}_1\|} \bar{\mathbf{v}}_1 \times (\dot{u}\mathbf{h} + \dot{\hat{\mathbf{v}}}_c). \quad (3.53)$$

The term $\mathbf{d}_{\text{proj}}$ appears because the physical command cannot contain a component parallel to $\mathbf{h}$; it vanishes when $\bar{\mathbf{v}}_1^\top \boldsymbol{\omega}_{\bar{\mathbf{v}}}^{\text{des}} = 0$. The term $\boldsymbol{\varepsilon}(t)$ is the time-scale approximation error associated with the omitted derivative terms in (3.48). It is a quantity to be checked in simulation, not an additional control input. The observer contribution $\dot{\hat{\mathbf{v}}}_c = -k_D \tilde{\mathbf{e}}$ decays exponentially; integrability of the part due to $\dot{u}$ is imposed as a time-scale condition and is

not claimed as a consequence of the controller. Accordingly, the analysis assumes

$$\int_0^\infty \|\boldsymbol{\varepsilon}(t)\| dt < \infty, \quad \boldsymbol{\varepsilon}(t) \rightarrow \mathbf{0}. \quad (3.54)$$

Lemma 3.5 (Sufficient time-scale approximation condition). *The time-scale condition in (3.54) is implied by the following explicit bound:*

$$\|\boldsymbol{\varepsilon}(t)\| \leq \frac{C_1 e^{-\sigma_o t} + C_2 \varphi(t)}{v_r^{\min} \gamma_-}, \quad \varphi \in L_1([0, \infty)), \quad \varphi(t) \rightarrow 0, \quad (3.55)$$

for finite nonnegative constants C_1, C_2 and $\sigma_o > 0$. This is a sufficient verification condition, not a consequence claimed for every reference or actuator implementation.

The estimate isolates the exponentially decaying observer contribution from the reference and command-rate contribution. Its proof is given in Section A.11; the numerical study reports the integral of the measured approximation error rather than assuming that the condition is automatically satisfied. The derivation of the reduced mapping and of the two approximation terms is given in Section A.2.

3.3.4 Direction-alignment analysis

Define

$$V_a := \frac{1}{2}(1 - \bar{\mathbf{a}}^\top \bar{\mathbf{v}}_1) \in [0, 1]. \quad (3.56)$$

Lemma 3.6 (Logistic identity). *For unit vectors $\bar{\mathbf{v}}_1$ and $\bar{\mathbf{a}}$,*

$$\|\bar{\mathbf{v}}_1 \times \bar{\mathbf{a}}\|^2 = 4V_a(1 - V_a). \quad (3.57)$$

This follows from $\|\bar{\mathbf{v}}_1 \times \bar{\mathbf{a}}\|^2 = 1 - (\bar{\mathbf{v}}_1^\top \bar{\mathbf{a}})^2$.

Theorem 3.7 (Nominal alignment except the antipodal direction). *If $\tilde{\mathbf{v}}_c = \mathbf{0}$, the ideal direction kinematics*

$$= \boldsymbol{\omega}_{\tilde{\mathbf{v}}}^{\text{des}} \times \bar{\mathbf{v}}_1 \quad (3.58)$$

hold, and $V_a(0) < 1$, then

$$\dot{V}_a = -2kV_a(1 - V_a), \quad (3.59)$$

$$V_a(t) = \frac{V_a(0)}{V_a(0) + (1 - V_a(0))e^{2kt}}. \quad (3.60)$$

Hence $V_a(t) \rightarrow 0$. The antipodal state $V_a = 1$ is the exceptional unstable equilibrium of the present field. This is consistent with the general topological obstruction to global continuous asymptotic stabilisation on $\mathbb{S}^2$; the cited results do not imply that every continuous or time-varying field must have exactly one antipodal repeller [45, 46].

The proof and intermediate vector identities are given in Section A.5. Condition (3.54) states the modelling requirement used when the reduced command is assessed. The

theorem remains an ideal-direction result; transfer to (3.50) additionally requires direction-rate matching or a separate perturbation argument. The internal signal flow is shown in Figure 3.1.




Figure 3.1: Internal block structure of the kinematic guidance layer.

3.3.5 Position error under nominal surge

Let

$$V_e := \frac{1}{2} \|\mathbf{e}\|^2, \quad s := \mathbf{e}^\top \mathbf{b}, \quad \Pi_e := \mathbf{e}^\top \bar{\mathbf{a}}. \quad (3.61)$$

The scalar s is the tangential target-relative error, whereas Π_e is the projection on the *normalised* blended direction. Retaining this normalisation is essential:

$$\Pi_e = \kappa_a [(1 - p) \|\mathbf{e}\| + ps]. \quad (3.62)$$

Lemma 3.8 (Bounded derivative of the position Lyapunov function). *Assume bounded reference derivatives, a uniform blend bound $\|\mathbf{a}\| \geq a_{\min} > 0$, feasible bounded policy factors, bounded realised guidance commands, and a bounded target-relative error $\|\mathbf{e}(t)\| \leq E$ on the considered invariant set. Then $\dot{V}_e$ is bounded and $\ddot{V}_e$ is uniformly continuous along every admissible trajectory in that set.*

The condition is used only when a Barbalat-type argument is invoked. In particular, boundedness of $\ddot{V}_e$ is not inferred from the kinematic geometry alone when rate saturation, sampling, or the dynamic plant is active; those effects are checked through the realised-command error in Section 4.3. The calculation is given in Section A.12.

Lemma 3.9 (Target-projection inequality). *Under Theorem 3.4,*

$$\Pi_e \geq s, \quad \Pi_e = s \iff \mathbf{e} = \ell \mathbf{b}, \quad \ell > -\lambda. \quad (3.63)$$

The equality at $\ell = 0$ is understood through the continuous virtual-target form (3.36); $\ell = -\lambda$ is the excluded blend singularity.

The proof, including the equality-set discussion, is given in Section A.6.

Under perfect alignment, converged current observer, and nominal speed matching,

$$\dot{V}_e = \|\dot{\mathbf{r}}\| (s - \Pi_e) \leq 0. \quad (3.64)$$

Thus nominal speed matching leaves a tangential equilibrium family $\mathbf{e} = \ell \mathbf{b}$, $\ell > -\lambda$, rather than forcing the target-relative error to zero. Positive ℓ denotes a vehicle behind the target; $-\lambda < \ell < 0$ denotes a vehicle ahead of the moving target but still behind the virtual look-ahead target. The hysteretic and continuous policies below remove this latter stationary behaviour through their $s < 0$ speed reduction.

3.3.6 Modified surge policies

Let $0 < \rho_{\text{in}} < \rho_{\text{out}}$ and use a hysteresis state $\mathcal{H} \in \{A, B\}$: Mode A changes to B at $\|\mathbf{e}\| \leq \rho_{\text{in}}$, and Mode B returns to A at $\|\mathbf{e}\| > \rho_{\text{out}}$. The piecewise policy factor is

$$\gamma(s, \mathcal{H}) = \begin{cases} \gamma_-, & \mathcal{H} = A, s < 0, \\ \gamma_+, & \mathcal{H} = A, s \geq 0, \\ \gamma_{\text{hold}}, & \mathcal{H} = B, \end{cases} \quad 0 < \gamma_- < 1 < \gamma_+, \quad \gamma_{\text{hold}} \geq 1. \quad (3.65)$$

Modulation A and the effective speed ratio

The multiplicative policy $u = \gamma u_{\text{nom}}$ does not, in general, produce an inertial speed equal to $\gamma \|\dot{\mathbf{r}}\|$ because the current is added after the surge command. Define

$$\mathbf{v}_{1,A} := \gamma u_{\text{nom}} \mathbf{h} + \hat{\mathbf{v}}_c, \quad \mu := \frac{\|\mathbf{v}_{1,A}\|}{\|\dot{\mathbf{r}}\|}. \quad (3.66)$$

Using the nominal quadratic relation gives

$$\|\mathbf{v}_{1,A}\|^2 = \gamma^2 \|\dot{\mathbf{r}}\|^2 + (1 - \gamma) \Delta_c, \quad (3.67)$$

$$\Delta_c := (1 + \gamma) \|\hat{\mathbf{v}}_c\|^2 + 2\gamma u_{\text{nom}} (\mathbf{h}^\top \hat{\mathbf{v}}_c), \quad (3.68)$$

and therefore

$$\mu = \sqrt{\gamma^2 + (1 - \gamma) \Delta_c / \|\dot{\mathbf{r}}\|^2} > 0. \quad (3.69)$$

Lemma 3.10 (Position derivative under Modulation A). *On the ideal alignment manifold,*

$$\dot{V}_e = \|\dot{\mathbf{r}}\| s - \mu \|\dot{\mathbf{r}}\| \Pi_e - \mathbf{e}^\top \tilde{\mathbf{v}}_c, \quad \Pi_e := \bar{\mathbf{a}}^\top \mathbf{e} \geq s. \quad (3.70)$$

Thus the sign of the longitudinal correction is determined by the effective ratio μ , not directly by the commanded factor γ .

The derivation of the effective ratio and of (3.70) is given in Section A.13.

Proposition 3.11 (Mode-A decrease under Modulation A). *Under Modulation A, ideal alignment $\tilde{\mathbf{v}}_1 = \bar{\mathbf{a}}$, and $\tilde{\mathbf{v}}_c \equiv \mathbf{0}$, provided the effective ratio satisfies $\mu > 1$ in Mode A with*

$s \geq 0$ and $\mu < 1$ in Mode A with $s < 0$ (which holds whenever $\|\hat{\mathbf{v}}_c\| \ll \|\dot{\mathbf{r}}\|$; when $\hat{\mathbf{v}}_c = \mathbf{0}$ one has $\mu = \gamma$ exactly), $\dot{V}_e < 0$ throughout Mode A for every $\mathbf{e} \neq \mathbf{0}$. This mode-wise result does not by itself establish uniform ultimate boundedness or an attracting set for the complete hysteretic closed loop; such conclusions require a separate hybrid invariance analysis of the Mode B flow and switching boundaries.

The conditional sign argument is collected in Section A.13.

Two mathematically distinct surge modulations are considered:

$$\text{Modulation A: } u = \gamma u_{\text{nom}}, \quad (3.71)$$

$$\text{Modulation B: } u = -\mathbf{h}^T \hat{\mathbf{v}}_c + \sqrt{(\mathbf{h}^T \hat{\mathbf{v}}_c)^2 - \|\hat{\mathbf{v}}_c\|^2 + \gamma^2 \|\dot{\mathbf{r}}\|^2}. \quad (3.72)$$

For Modulation B, the exact feasibility requirement is a nonnegative radicand together with positivity of the selected root. The transparent sufficient condition used in the analysis is

$$\gamma(t) \|\dot{\mathbf{r}}(t)\| > \|\hat{\mathbf{v}}_c(t)\|. \quad (3.73)$$

For the continuous policy below, a uniform sufficient condition is $(1 - \eta) v_r^{\min} > \sup_t \|\hat{\mathbf{v}}_c(t)\|$. If feasibility is lost during an observer transient, the implementation applies a bounded previously feasible surge command; the nominal position theorems do not apply on that fallback interval. Modulation B enforces $\|\mathbf{v}_1\| = \gamma \|\dot{\mathbf{r}}\|$ exactly. For Modulation A, define the effective ratio $\mu := \|\mathbf{v}_1\| / \|\dot{\mathbf{r}}\|$; it is current dependent and need not equal γ .

Under alignment, the general position derivative is

$$\dot{V}_e = \|\dot{\mathbf{r}}\| s - \|\mathbf{v}_1\| \Pi_e - \mathbf{e}^T \tilde{\mathbf{v}}_c. \quad (3.74)$$

For Modulation B it becomes

$$\dot{V}_e = \|\dot{\mathbf{r}}\| (s - \gamma \Pi_e) - \mathbf{e}^T \tilde{\mathbf{v}}_c. \quad (3.75)$$

Theorem 3.12 (Mode-wise properties of the hysteresis policy). *Assume ideal alignment, $\tilde{\mathbf{v}}_c = \mathbf{0}$, feasible positive-root operation according to (3.73), and Modulation B. In Mode A, V_e decreases strictly for every $\mathbf{e} \neq \mathbf{0}$. In Mode B with $\gamma_{\text{hold}} = 1$, it is non-increasing. Within the Mode B flow domain $\|\mathbf{e}\| \leq \rho_{\text{out}}$, its instantaneous zero-derivative set is*

$$\mathcal{E}_0 = \{\mathbf{e} = \ell \mathbf{b} : \ell > -\lambda, |\ell| \leq \rho_{\text{out}}\}. \quad (3.76)$$

For $\gamma_{\text{hold}} > 1$, strict decrease also holds in the part of Mode B with $s \geq 0$ and $\mathbf{e} \neq \mathbf{0}$; no monotonicity claim is made here for Mode B states with $s < 0$. For Modulation A, the corresponding Mode A conclusions require $\mu > 1$ for $s \geq 0$ and $\mu < 1$ for $s < 0$.

The mode-wise sign analysis is given in Section A.7. The set above is an instantaneous flow set, not an asserted attracting or ultimate set. Reachability depends on switching history, and attraction to a time-varying set for curved $\mathbf{b}(t)$ would require a separate hybrid invariance argument.

Corollary 3.13 (Zero-radius ideal policy). *If $\rho_{\text{in}} = \rho_{\text{out}} = 0$, the ideal Modulation B policy has no nonzero Mode B enclosure. Under the same alignment, observer, and feasibility assumptions as Theorem 3.12, the origin is the only remaining tangential equilibrium. The inequalities*

$$\dot{V}_e \leq -\|\dot{\mathbf{r}}\| (\gamma_+ - 1) \|\mathbf{e}\|, \quad s \geq 0, \quad (3.77)$$

$$\dot{V}_e \leq -\|\dot{\mathbf{r}}\| \gamma_- (1 - p) \|\mathbf{e}\|, \quad s < 0, \quad (3.78)$$

are rate bounds for the ideal channel; they are not asserted as exponential rates for the physical dynamic cascade.

This corollary is obtained by applying the mode-wise inequalities of Section A.7 with a zero-width enclosure.

3.3.7 Continuous policy and zero-error convergence

The smooth alternative removes the cruise band:

$$\gamma(s) = 1 + \eta \tanh(s/\alpha_0), \quad 0 < \eta < 1, \quad \alpha_0 > 0. \quad (3.79)$$

Theorem 3.14 (Negative-definite nominal position derivative). *Under ideal alignment, Modulation B, $\tilde{\mathbf{v}}_c = 0$, (3.73), and Theorem 3.4,*

$$\dot{V}_e < 0 \text{ for every } \mathbf{e} \neq 0, \quad \dot{V}_e = 0 \iff \mathbf{e} = 0. \quad (3.80)$$

The three-case sign proof and the compact-uniform argument that upgrades pointwise negativity to asymptotic convergence are given in Sections A.8 and A.9.

Corollary 3.15 (Conditional asymptotic convergence). *If, in addition to the hypotheses of Theorem 3.14, the pair $(\mathbf{e}(t), \mathbf{b}(t))$ remains in a bounded set with the uniform blend bound $\|\mathbf{a}\| \geq a_{\min} > 0$, then $\mathbf{e}(t) \rightarrow \mathbf{0}$. The required compact uniform-decay argument is given in Section A.9; pointwise negativity alone is not used as a Barbalat conclusion.*

Lemma 3.16 (Quantitative continuous-policy decay). *On the ideal alignment and converged-observer manifold, the continuous policy admits the explicit two-branch estimate*

$$\dot{V}_e \leq -\|\dot{\mathbf{r}}\| \psi(s, \|\mathbf{e}\|), \quad (3.81)$$

where, with $\gamma_- := 1 - \eta$,

$$\psi(s, r) := \begin{cases} (\gamma(s) - 1)r, & s > 0 \text{ and } \gamma(s)p \leq 1, \\ \gamma_- \frac{r^2}{\lambda + r}, & \text{otherwise.} \end{cases} \quad (3.82)$$

Near the target, the second branch gives $\dot{V}_e \lesssim -2\gamma_- \|\dot{\mathbf{r}}\| V_e/\lambda$.

The two branches identify the practical mechanism of the smooth policy: the far-side

correction is proportional to the distance, whereas the near-side bound is quadratic in the distance. The proof is collected in Section A.14.

3.3.8 Layered convergence statement for the ideal kinematic model

The preceding results were stated in layers: observer convergence, direction alignment, and position-error decrease under ideal alignment. For reporting these layers together, define

$$V_o = \mathbf{z}_o^\top P \mathbf{z}_o, \quad A_o^\top P + P A_o = -Q, \quad P = P^\top \succ 0, \quad Q = Q^\top \succ 0, \quad (3.83)$$

where $\mathbf{z}_o = [\tilde{\mathbf{e}}^\top \tilde{\mathbf{v}}_c^\top]^\top$, together with

$$V_{\text{kin}} := c_o V_o + c_a V_a + c_e V_e, \quad c_o, c_a, c_e > 0, \quad (3.84)$$

where the positive constants provide the dimensional scaling required when observer, directional, and position quantities are displayed in one storage measure. The convergence proof below does not infer that $\dot{V}_{\text{kin}}$ is negative term by term. Instead, it retains the observer and alignment errors explicitly as exponentially vanishing inputs to the position channel. The complete argument is given in Section A.15.

Remark 3.17 (Conditional scope of the layered result). The result below is only for the reduced ideal kinematic model $\dot{\mathbf{q}} = \mathbf{u}\mathbf{h} + \mathbf{v}_c$. It uses the following explicit hypotheses: (i) the ideal direction dynamics $\dot{\mathbf{v}} = \omega_{\mathbf{v}}^{\text{des}} \times \mathbf{v}_1$ are realised exactly; (ii) the initial directions are not antipodal, $V_a(0) < 1$; (iii) the admissible domain $\mathcal{D}$ of (3.38) is forward invariant; (iv) the reference derivatives are bounded; and (v) Modulation B remains feasible. The result is *not* claimed to extend to the six-degree-of-freedom REMUS/PI closed loop of Chapter 4; conditional routes for such an extension are collected in Appendix E.

Theorem 3.18 (Conditional convergence for the reduced ideal kinematic model). *Under Theorem 3.4, Theorem 3.1, and the standing assumptions of Section 3.1, consider a closed-loop trajectory of the continuous Modulation B law for which the hypotheses of Theorem 3.17 hold. Then the observer error, alignment error, and target-relative position error converge according to*

$$\lim_{t \rightarrow \infty} \mathbf{z}_o(t) = \mathbf{0}, \quad \lim_{t \rightarrow \infty} V_a(t) = 0, \quad \forall \varepsilon > 0 \exists T_\varepsilon : \|\mathbf{e}(t)\| \leq \varepsilon \quad \forall t \geq T_\varepsilon. \quad (3.85)$$

For every $\varepsilon > 0$, the finite entry time T_ε depends on the initial observer, alignment, and position errors. The proof first bounds the position error using the integrability of the observer and alignment transients. It then uses uniform positivity of the nominal decay function on each compact annulus $\varepsilon \leq \|\mathbf{e}\| \leq E$ to establish entry into, and retention inside, every prescribed neighbourhood. This avoids imposing a linear-in- $\|\mathbf{e}\|$ decay bound near the origin, where the nominal decay is generally quadratic. The detailed calculation is given in Section A.15.

The projected physical command used in this thesis satisfies the exact matching condition only when $\tilde{\mathbf{v}}_1^\top \omega_{\mathbf{v}}^{\text{des}} = 0$; otherwise the explicit discrepancy in (3.51) remains. Each of the three

ideal-model requirements is stated in Section A.15, together with the exact point at which the physical mapping leaves the proof. Consequently, Theorem 3.18 is a reduced-model statement rather than an implementation guarantee. The integrability condition in Theorem 3.5 is a useful requirement for a future perturbation analysis, but it is not used here to claim automatic transfer of the theorem to the projected command.

The time-scale approximation error (3.53), the projection discrepancy in (3.51), sensor noise, sampling, saturation, and the REMUS/Fossen inner model are therefore not absorbed silently into the nominal theorems. They are logged and evaluated in the two simulation tracks of Chapter 5; a complete stability proof for the dynamic cascade is deferred.

3.4 Chapter summary

This chapter established the kinematic guidance layer in three steps. First, the constant-current observer estimates the inertial current from the computed target-relative error. Second, the direction law aligns the estimated inertial-velocity direction with a smooth approach–tangent target direction on the admissible domain. Third, the surge policy determines how the signed along-path error evolves: nominal speed matching can leave a tangential equilibrium, whereas the continuous Modulation B policy removes that equilibrium in the ideal aligned kinematic channel. The angular-rate command used by the vehicle is a reduced mapping, and the associated approximation terms are derived in Section A.2. These are layered kinematic results; they do not imply stability of the simultaneous observer–guidance–vehicle interconnection considered in the next chapter.

Chapter 4

Dynamic Model and Controller Integration

The kinematic controller produces a water-relative surge command and an inertial angular-velocity vector, whereas the REMUS 100 plant evolves body-fixed velocities and rates under hydrodynamic and actuator forces. This chapter defines the six-degree-of-freedom model, the full-attitude current representation, and the command interface connecting those descriptions. The plant is an implementation and validation environment; the formal controller contribution remains kinematic.

4.1 Six-Degree-of-Freedom REMUS 100 Model

4.1.1 Purpose and modelling scope

The kinematic development assumes ideal execution of the commanded surge and angular velocity. To examine the physical implications of this assumption, the guidance law is connected to a six-degree-of-freedom REMUS 100 plant (Figure 4.1) with hydrodynamic forces and three physical actuators. The model follows the formulation implemented in the Fossen MSS toolbox [1]. Its role is to define the plant, environmental terms, actuator mapping, and physical limits used in the numerical study, rather than to introduce a new dynamic-control contribution.



Figure 4.1: The REMUS 100 autonomous underwater vehicle, modelled as the reference platform for the dynamic simulations of Chapter 5. Complete geometric, hydrodynamic and actuator specifications are given in Section D.6. Image reproduced from [6].

4.1.2 Frames and state

The inertial frame is North-East-Down (NED). The body-fixed frame has its origin at the control origin, with the x axis forward, y starboard, and z downward. The state is

$$\mathbf{x} = \begin{bmatrix} \mathbf{v} \\ \boldsymbol{\eta} \end{bmatrix} = \begin{bmatrix} u_b & v_b & w_b & p_b & q_b & r_b & x & y & z & \phi & \theta & \psi \end{bmatrix}^\top, \quad (4.1)$$

where

$$\mathbf{v} = \begin{bmatrix} u_b & v_b & w_b & p_b & q_b & r_b \end{bmatrix}^\top, \quad \boldsymbol{\eta} = \begin{bmatrix} x & y & z & \phi & \theta & \psi \end{bmatrix}^\top. \quad (4.2)$$

The subscript b distinguishes measured body surge and rates from the ideal guidance input u .

4.1.3 Ocean current and relative velocity

The hydrodynamic forces depend on velocity relative to the surrounding water:

$$\mathbf{v}_r = \mathbf{v} - \mathbf{v}_c. \quad (4.3)$$

For a horizontal current magnitude V_c , bearing β_{V_c} , and vertical component w_c , define the constant NED current

$$\mathbf{v}_c = \begin{bmatrix} V_c \cos \beta_{V_c} & V_c \sin \beta_{V_c} & w_c \end{bmatrix}^\top. \quad (4.4)$$

Its frame-consistent body representation is

$$\mathbf{v}_c = \begin{bmatrix} \mathbf{R}_{zyx}^\top \mathbf{v}_c \\ \mathbf{0}_{3 \times 1} \end{bmatrix}. \quad (4.5)$$

Since the hydrodynamic equations are written in relative velocity, the corresponding derivative is retained:

$$\dot{\mathbf{v}}_c = \begin{bmatrix} -\mathbf{S}(\boldsymbol{\omega}^b) \mathbf{R}_{zyx}^\top \mathbf{v}_c \\ \mathbf{0}_{3 \times 1} \end{bmatrix}, \quad \boldsymbol{\omega}^b := \begin{bmatrix} p_b & q_b & r_b \end{bmatrix}^\top. \quad (4.6)$$

The distinction between ground velocity $\mathbf{v}$ and water-relative velocity $\mathbf{v}_r$ is essential. Added mass, lift, drag, and fin forces use $\mathbf{v}_r$, while inertial position is propagated using $\mathbf{v}$. The relation between (4.5) and the yaw-only current parametrisation commonly used in the baseline REMUS implementation is analysed in Section 4.2.

4.1.4 Kinematics

The pose dynamics are

$$\dot{\boldsymbol{\eta}} = \mathbf{J}(\boldsymbol{\eta}) \mathbf{v}, \quad \mathbf{J} = \begin{bmatrix} \mathbf{R}_{zyx}(\phi, \theta, \psi) & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_{zyx}(\phi, \theta) \end{bmatrix}. \quad (4.7)$$

The first column of $\mathbf{R}_{zyx}$ is the heading direction used by the outer controller:

$$\mathbf{h} = \mathbf{R}_{zyx} \mathbf{e}_1. \quad (4.8)$$

The Euler-rate matrix is singular at $\theta = \pm 90^\circ$, but the nominal REMUS operating envelope remains far from this set.

4.1.5 Equations of motion

The relative-velocity form is

$$\boxed{\mathbf{M}\dot{\mathbf{v}}_r + \mathbf{C}(\mathbf{v}_r)\mathbf{v}_r + \mathbf{D}(\mathbf{v}_r)\mathbf{v}_r + \mathbf{g}(\boldsymbol{\eta}) = \boldsymbol{\tau} + \boldsymbol{\tau}_{\text{lift}} + \boldsymbol{\tau}_{\text{cf}}} \quad (4.9)$$

with

$$\dot{\mathbf{v}} = \dot{\mathbf{v}}_c + \mathbf{M}^{-1}(\boldsymbol{\tau} + \boldsymbol{\tau}_{\text{lift}} + \boldsymbol{\tau}_{\text{cf}} - \mathbf{C}(\mathbf{v}_r)\mathbf{v}_r - \mathbf{D}(\mathbf{v}_r)\mathbf{v}_r - \mathbf{g}(\boldsymbol{\eta})). \quad (4.10)$$

The inertia and Coriolis matrices decompose as

$$\mathbf{M} = \mathbf{M}_{RB} + \mathbf{M}_A, \quad \mathbf{C} = \mathbf{C}_{RB} + \mathbf{C}_A. \quad (4.11)$$

The Coriolis terms redistribute energy across degrees of freedom, damping dissipates energy, and the restoring vector provides passive roll and pitch stability.

4.1.6 Rigid-body and added-mass properties

The hull is approximated by a prolate spheroid with length $L = 1.6$ m and diameter $d_h = 0.19$ m. The model mass is $m_b \approx 31$ kg, with dry-body moments approximately $I_x = 0.112$ kg m² and $I_y = I_z = 4.02$ kg m². The centre of gravity lies 0.02 m below the control origin, while the centre of buoyancy coincides with the control origin.

The elongated geometry produces strongly anisotropic added mass. The Lamb factors are approximately $\kappa_1^A = 0.027$ in surge, $\kappa_2^A = 0.949$ in sway and heave, and $\kappa_3^A = 0.851$ in pitch and yaw. The adopted axisymmetric approximation uses $I_y = I_z$; the yaw expression is nevertheless written with I_z to preserve the axis convention. The resulting diagonal effective inertias are summarized in Table 4.1.

Table 4.1: Approximate effective inertia of the REMUS 100 model.

Degree of freedom	Expression	Approximate value
Surge	$m_b(1 + \kappa_1^A)$	31.8 kg
Sway	$m_b(1 + \kappa_2^A)$	60.4 kg
Heave	$m_b(1 + \kappa_2^A)$	60.4 kg
Roll	$I_x(1 + r_{44})$	0.146 kg m ²
Pitch	$I_y(1 + \kappa_3^A)$	7.44 kg m ²
Yaw	$I_z(1 + \kappa_3^A)$	7.44 kg m ²

The nearly doubled lateral effective mass explains why a torpedo-shaped vehicle cannot reproduce arbitrary transverse velocity commands. The outer controller must rotate the vehicle and allow the hydrodynamic response to develop.

4.1.7 Damping and restoring forces

Three damping mechanisms are represented: low-speed linear damping, nonlinear hull lift and drag, and cross-flow drag. In the adopted REMUS 100 implementation, the low-speed surge and sway terms are tapered with relative speed and the cross-flow contribution is evaluated by strip theory over the hull sections. The linear term has the form

$$\mathbf{D}_{\text{lin}} = \text{diag}\left(\frac{M_{11}}{T_1}, \frac{M_{22}}{T_2}, \frac{M_{33}}{T_2}, 2\zeta_4\omega_4M_{44}, 2\zeta_5\omega_5M_{55}, \frac{M_{66}}{T_6}\right), \quad (4.12)$$

with $T_1 = T_2 = 20$ s and $T_6 = 1$ s. At cruise, quadratic forces dominate. Cross-flow strip theory produces sway, heave, pitch, and yaw forces proportional to local transverse speed squared.

Neutral buoyancy removes translational hydrostatic forces. With the NED and Euler-angle conventions used here, the restoring vector is

$$\mathbf{g}(\boldsymbol{\eta}) = [0 \quad 0 \quad 0 \quad z_g W \cos \theta \sin \phi \quad z_g W \sin \theta \quad 0]^\top, \quad (4.13)$$

where $W = m_b g$. Because the centre of gravity lies below the centre of buoyancy, the restoring vector contains roll and pitch moments that return the vehicle toward level attitude. There is no hydrostatic yaw restoring moment.

4.1.8 Actuators and generalized input

The vehicle has three commands:

1. rudder angle δ_r for yaw, with an associated sway force;
2. stern-plane angle δ_s for pitch, with an associated heave force; and
3. propeller speed n for surge, with a small roll torque.

The generalized input is

$$\boldsymbol{\tau} = \begin{bmatrix} (1 - t_{\text{prop}})X_{\text{prop}} + X_r + X_s \\ Y_r \\ Z_s \\ K_{\text{prop}}/10 \\ -x_s Z_s \\ x_r Y_r \end{bmatrix}. \quad (4.14)$$

The fin forces scale approximately with water-relative speed squared. Consequently, yaw and pitch authority collapse at low forward speed. This constraint is incorporated in the interpretation of the surge policy: a mathematically admissible low surge may leave the dynamic vehicle unable to realise the requested angular motion.

The input map in (4.14) is a model-specific interface, not a universal actuator law. In particular, the factor $K_{\text{prop}}/10$ is a calibration or scaling factor inherited from the adopted REMUS implementation; its source, sign, and numerical value must be recorded with the

simulation model. It must not be interpreted as a physical relation without that model traceability. The same requirement applies to $X_{\text{prop}}(n)$, the fin-force coefficients, and the signs of X_r , Y_r , and Z_s .

4.1.9 Physical constraints

Table 4.2: Control-relevant REMUS 100 limits and nominal parameters.

Quantity	Symbol	Value
Fin limit	$ \delta_r , \delta_s $	20°
Propeller limit	n_{max}	1525 rpm
Approximate top speed	U_{max}	2.5 m s^{-1}
Sustainable yaw rate	$ r_b $	$0.065\text{--}0.10 \text{ rad s}^{-1}$
Guidance-interface rate clamp	$ q_{\text{cmd}} , r_{\text{cmd}} $	15° s^{-1}
Surge time constant	T_1	20 s
Yaw time constant	T_6	1 s

The difference between the 15° s^{-1} command clamp and the smaller sustainable steady yaw rate is intentional: the first limits requests sent to the PI loop, while the latter describes what the hydrodynamic plant can maintain. The numerical analysis therefore distinguishes commanded from achieved rates and records the corresponding saturation intervals.

4.1.10 Model traceability and coefficient declaration

The model equations and effective inertias above define the plant structure; the numerical coefficient set, actuator limits, controller gains, solver settings, and initial conditions used for the reported runs are consolidated in Chapter D. This separates the physical discussion in the main chapter from the reproducibility record in the appendix.

4.1.11 Reduced models for controller design

Near straight-line cruise, the dominant channels can be approximated by

$$M_{11}\dot{u}_b \approx (1 - t_{\text{prop}})X_{\text{prop}}(n) - D_{11}u_r + X_r + X_s, \quad (4.15)$$

$$T_{\text{yaw}}\dot{r}_b + r_b = K_{\text{yaw}}\delta_r, \quad (4.16)$$

$$M_{55}\dot{q}_b \approx -x_s Z_s - D_{55}q_b - z_g W \sin \theta. \quad (4.17)$$

These reductions support initial gain selection, whereas the reported numerical results are obtained from the complete model in (4.9). They also identify a possible starting point for a future backstepping or nonlinear inner-loop design with an explicit dynamic Lyapunov function.

4.2 Frame-Consistent Ocean-Current Representation

The observer of Theorem 3.2 estimates a current that is constant in the inertial frame. Preserving that property in a six-degree-of-freedom simulation requires more than using the same numerical current magnitude: the body-frame current supplied to the hydrodynamic model must transform back to the same inertial vector for every vehicle attitude. This section identifies the inconsistency that arises from a yaw-only current map and defines the full-rotation correction used in the coupled simulations.

The issue is isolated analytically below by comparing the effective inertial current generated by a yaw-only map with the intended constant NED vector. The controlled numerical comparison is specified in Table D.11; exploratory runs are excluded from the reported evidence.

4.2.1 Yaw-only mapping and attitude-dependent leakage

A common horizontal-current parametrisation forms the body-linear current as

$$\mathbf{v}_{c,\psi}^b = \begin{bmatrix} V_c \cos(\beta_{V_c} - \psi) \\ V_c \sin(\beta_{V_c} - \psi) \\ w_c \end{bmatrix}. \quad (4.18)$$

This representation is compatible with planar motion and remains a close approximation when roll and pitch are small. If it is combined with the full pose kinematics, however, the current actually injected into the NED position rate becomes

$$\bar{\mathbf{v}}_c(t) := \mathbf{R}_{zyx}(\phi, \theta, \psi) \mathbf{v}_{c,\psi}^b, \quad (4.19)$$

which is attitude dependent. For $\phi = 0$, its vertical component is

$$\bar{v}_{c,z}(t) = -\sin \theta V_c \cos(\beta_{V_c} - \psi) + \cos \theta w_c. \quad (4.20)$$

Thus finite pitch rotates part of the nominal horizontal current into the vertical NED channel. The effect is small on a gentle helix but can become comparable with the prescribed vertical current during terrain-following motion with large pitch.

The observer of (3.17) was derived under the assumption that the current is a constant inertial vector. Under the yaw-only map, the current the observer actually sees along the trajectory is $\bar{\mathbf{v}}_c(t)$ of (4.19), which is attitude dependent. The observer error therefore satisfies the *forced* system

$$\dot{\mathbf{z}}_o = \mathbf{A}_o \mathbf{z}_o + \begin{bmatrix} \mathbf{0} \\ \dot{\bar{\mathbf{v}}}_c(t) \end{bmatrix}, \quad \mathbf{z}_o = [\tilde{\mathbf{e}}^\top \quad \tilde{\mathbf{v}}_c^\top]^\top, \quad (4.21)$$

in place of the homogeneous system (3.17). Although $\mathbf{A}_o$ is Hurwitz, the time-varying forcing $\dot{\bar{\mathbf{v}}}_c(t)$ prevents convergence to a constant vector in general. The attitude-rate dependence of $\dot{\bar{\mathbf{v}}}_c$ is given in Section B.2.

4.2.2 Full-rotation correction

The simulation instead starts from the intended constant NED vector (4.4) and maps it into the body frame with the full attitude matrix,

$$\mathbf{v}_c^b = \mathbf{R}_{zyx}^\top \mathbf{v}_c, \quad \dot{\mathbf{v}}_c^b = -\mathbf{S}(\boldsymbol{\omega}^b) \mathbf{v}_c^b. \quad (4.22)$$

Both the algebraic map and its derivative are required: the former defines the water-relative velocity, while the latter enters the relative-velocity equations of motion through $\dot{\mathbf{v}}_c$.

Proposition 4.1 (Recovery of the nominal observer model). *If the plant uses (4.22), attitude is known exactly, and the noise-free DVL-aiding term is $\mathbf{v}_{\text{rel}} = \mathbf{R}_{zyx} \mathbf{v}_{r,1:3}$, then the current contribution to the NED position kinematics is exactly $\mathbf{v}_c$ for every attitude. Consequently the observer error reduces to the homogeneous system (3.17) and, by Theorem 3.2, $\hat{\mathbf{v}}_c \rightarrow \mathbf{v}_c$ exponentially with time constant $2/k_C$ in the critically damped case.*

The proof, including verification of (4.6), is given in Section B.3.

4.2.3 Implementation boundary

The correction is applied as a current-interface wrapper around the REMUS plant at every integration stage; the validated rigid-body, added-mass, damping, restoring, lift, cross-flow, and actuator models are otherwise unchanged. This is a frame-consistency correction between the simulation plant and the guidance assumption, not a new hydrodynamic model. Remaining depth error caused by stern-plane or pitch-rate saturation is therefore analysed separately from observer error in the numerical-validation chapter.

4.3 Kinematic-to-Dynamic Command Interface

4.3.1 Purpose and scope

The kinematic guidance law produces an ideal surge command and an inertial angular-velocity command, whereas the REMUS 100 model accepts rudder, stern-plane, and propeller inputs. This section defines the intervening command interface, making the transition from $(u_{\text{cmd}}, \boldsymbol{\omega}_{\text{cmd}})$ to (δ_r, δ_s, n) explicit and consistent with the body and inertial frames. The resulting architecture is used to assess dynamic realisability; a full stability claim for the cascade lies outside the scope of the thesis.

4.3.2 Closed-loop data flow

The block-diagram of the kinematic-to-dynamic interface adopted in the sequel is shown in Figure 4.2.




Figure 4.2: Complete closed-loop cascade of the outer kinematic guidance, coupling interface, and inner six-degree-of-freedom REMUS 100 dynamics.

4.3.3 Conversion of the angular command

The guidance layer returns an inertial angular-velocity command ω_{cmd} satisfying $\omega_{\text{cmd}}^T \mathbf{h} = 0$. The interface rotates it into the body frame,

$$\omega_{\text{cmd}}^b = \mathbf{R}_{zyx}^T \omega_{\text{cmd}} = [\omega_x^b \quad q_{\text{cmd}} \quad r_{\text{cmd}}]^T. \quad (4.23)$$

Since $\omega_{\text{cmd}}^T \mathbf{h} = 0$ and $\mathbf{h} = \mathbf{R}_{zyx} \mathbf{e}_1$, the ideal body-axis roll component is zero; the interface projects out any numerical roundoff component there, then rate-limits the pitch and yaw references before the fin loops:

$$q_{\text{cmd}}, r_{\text{cmd}} \leftarrow \text{clamp}(q_{\text{cmd}}, r_{\text{cmd}}; \pm \dot{\theta}_{\text{max}}), \quad \dot{\theta}_{\text{max}} = 15^\circ \text{s}^{-1}. \quad (4.24)$$

Consistency of the rotation convention, the signs of q and r , and the discarded roll component is evaluated against the adopted Fossen-model implementation.

4.3.4 Water-relative surge and ground-frame measurement

The guidance command u_{cmd} is a water-relative surge, whereas the dynamic state u_b is the body- x component of ground velocity: the ideal kinematics give $\dot{\mathbf{q}} = u_{\text{cmd}} \mathbf{h} + \mathbf{v}_c$, and the exact pose relation $\dot{\mathbf{q}} = \mathbf{R}_{zyx} \mathbf{v}$ projected onto the body x axis gives $u_b = (\mathbf{v}_r)_1 + \mathbf{h}^T \mathbf{v}_c$. The ideal model replaces $(\mathbf{v}_r)_1$ by u_{cmd} , so on the dynamic plant this equality is a tracking objective rather than a plant identity, and the surge reference is corrected with the current estimate:

$$u_{\text{target}} = u_{\text{cmd}} + \mathbf{h}^T \hat{\mathbf{v}}_c. \quad (4.25)$$

The propeller feedforward continues to use u_{cmd} (thrust depends on water-relative speed); without (4.25) the rotating current projection would appear as a false surge-tracking error.

4.3.5 Low-level actuator interface

The rate-limited commands are tracked by independent PI loops for the stern plane, rudder, and propeller. The unsaturated commands are

$$\delta_s^{\text{raw}} = - \left(K_{p,q} e_q + K_{i,q} \int e_q dt \right), \quad e_q = q_{\text{cmd}} - q_b, \quad (4.26)$$

$$\delta_r^{\text{raw}} = - \left(K_{p,r} e_r + K_{i,r} \int e_r dt \right), \quad e_r = r_{\text{cmd}} - r_b, \quad (4.27)$$

$$n^{\text{raw}} = n_{\text{ff}} u_{\text{cmd}} + K_{p,n} e_u + K_{i,n} \int e_u dt, \quad e_u = u_{\text{target}} - u_b. \quad (4.28)$$

The actuator commands applied to the plant are

$$\delta_s = \text{sat}_{[-\delta_{s,\text{max}}, \delta_{s,\text{max}}]}(\delta_s^{\text{raw}}), \quad \delta_r = \text{sat}_{[-\delta_{r,\text{max}}, \delta_{r,\text{max}}]}(\delta_r^{\text{raw}}), \quad n = \text{sat}_{[0, n_{\text{max}}]}(n^{\text{raw}}). \quad (4.29)$$

The limits are given in Table 4.2. During saturation each PI integrator is clamped when further integration would grow the saturation error, with I_j the integrator state and $y_j^{\text{raw}}, y_j^{\text{sat}}$ the raw and applied commands:

$$I_j = \begin{cases} e_j, & y_j^{\text{raw}} = y_j^{\text{sat}} \text{ or } e_j(y_j^{\text{sat}} - y_j^{\text{raw}}) \geq 0, \\ 0, & \text{otherwise.} \end{cases} \quad (4.30)$$

Nominal values are $K_{p,q} = K_{p,r} = 6$, $K_{i,q} = K_{i,r} = 1.5$, $K_{p,n} = 200$, $K_{i,n} = 40$, $n_{\text{ff}} = 610 \text{ RPM}/(\text{m/s})$; fin and propeller signs follow Equation (4.14).

4.3.6 DVL-aided current observer

On the ideal kinematic plant $u\mathbf{h}$ is the entire non-current velocity, whereas a dynamic REMUS vehicle develops sway and heave during turns, so a command-model-only observer would interpret sideslip as a time-varying current. Given a DVL water-track measurement, the observed non-current velocity is

$$\mathbf{v}_{\text{rel}} = \mathbf{R}_{zyx} \mathbf{v}_{r,1:3}, \quad (4.31)$$

here assumed body-frame and rotated into NED with the measured attitude. Bottom-track, attitude error, and sensor noise contribute a separate measurement perturbation and must not be silently identified with the constant-current estimation error; each test case therefore labels $\mathbf{v}_{\text{rel}}$ as model-derived, noise-free, or sensor-corrupted. This separation follows the observer-based marine-guidance distinction between directly estimated current components and model-dependent reconstruction of water-relative velocity [17, 34]. The observer is then

$$\dot{\hat{\mathbf{e}}} = \dot{\mathbf{r}} - \mathbf{v}_{\text{rel}} - \hat{\mathbf{v}}_c + k_C(\mathbf{e} - \hat{\mathbf{e}}), \quad (4.32)$$

$$\dot{\hat{\mathbf{v}}}_c = -k_D(\mathbf{e} - \hat{\mathbf{e}}). \quad (4.33)$$

Under the full-rotation current map (4.22), $\dot{\mathbf{q}} = \mathbf{v}_{\text{rel}} + \mathbf{v}_c$ and the observer retains the homogeneous error system of Theorem 3.2; a yaw-only current map instead yields the forced effective-current error system (4.21). The two cases are separated explicitly in the numerical-validation chapter.

4.3.7 Realised-command error and analytical implications

Let the water-relative inertial velocity realised by the dynamic plant be $\mathbf{v}_{\text{rel}}$. The difference between this velocity and the outer-loop request is

$$\boldsymbol{\delta} := \mathbf{v}_{\text{rel}} - u_{\text{cmd}} \mathbf{h}. \quad (4.34)$$

Under the frame-consistent current map, $\tilde{\mathbf{v}}_c = \mathbf{v}_c - \hat{\mathbf{v}}_c$ and the realised inertial velocity and target-error dynamics satisfy

$$\dot{\mathbf{q}} = \mathbf{v}_1 + \tilde{\mathbf{v}}_c + \boldsymbol{\delta}, \quad \dot{\mathbf{e}} = \dot{\mathbf{r}} - \mathbf{v}_1 - \tilde{\mathbf{v}}_c - \boldsymbol{\delta}. \quad (4.35)$$

The term $\boldsymbol{\delta}$ collects surge lag, rate-tracking error, saturation, sideslip, and plant-model differences, while $\tilde{\mathbf{v}}_c$ retains the distinct contribution of the current-estimation error; under the yaw-only map $\tilde{\mathbf{v}}_c$ is replaced by $\tilde{\mathbf{v}}_{c,\text{eff}}$ to prevent plant-induced frame leakage from being misattributed to the observer. This decomposition links the ideal kinematic analysis to the six-degree-of-freedom simulations and is interpreted here as an implementation check, not as the basis for an ISS, UUB, or composite-Lyapunov claim.

The interface thus converts $(u_{\text{cmd}}, \omega_{\text{cmd}})$ into (δ_r, δ_s, n) through body-rate conversion, current-corrected surge targeting, bounded PI loops, and the full-rotation DVL-aided observer. The error $\boldsymbol{\delta}$ in (4.34) is the perturbation channel that any future cascade theorem must bound; the present thesis reports it only as an implementation measure.

4.4 Chapter summary

The model, current map, and command interface preserve the physical meaning of the kinematic signals during dynamic simulation. They provide the implementation tested in Chapter 5, not a new dynamic controller or cascade-stability theorem.

Chapter 5

Numerical Results and Comparison

This chapter tests three aspects of the thesis law: current estimation, outer direction–position behaviour, and realisation of the resulting commands by a six-degree-of-freedom AUV. In the ideal kinematic cases the observer uses position, surge direction, and the known command-model water-relative term; in the dynamic cases it uses the noise-free model-derived water-relative velocity through the DVL-like interface of Section 4.3.

Six analysis cases, five reference geometries, and two plant fidelities separate guidance behaviour from actuator and model effects (Table 5.2). The proposed law is first studied in its four analysed internal configurations; the plots also carry an exploratory path-following variant as a fifth supplementary trace. The main comparison then includes the thesis law, its switching predecessor, the exploratory variant, and two published three-dimensional baselines under a controlled shared-plant protocol.

The numerical study is structured in two complementary layers. The first layer assesses the proposed guidance law in isolation and after dynamic realisation, using respectively the ideal kinematic plant and the REMUS 100/PI implementation. The second layer evaluates the selected M-B continuous-PP configuration against the MED switching predecessor, the exploratory M-B continuous-PF variant, and two external three-dimensional guidance baselines under a common transfer protocol. Here PP denotes the prescribed-progression formulation analysed in this thesis, whereas PF denotes the supplementary autonomous path-progression variant.

5.1 Scenarios and comparison protocol

The five paths exercise complementary conditions: continuous spatial curvature (helix), pure translation (inclined line), repeated curvature transitions (lawnmower), closed spatial motion with a crossing only in horizontal projection (Gerono lemniscate), and terrain following at a desired 25 m clearance (oceanscan). The last two survey paths are C^1 at leg–turn junctions; the oceanscan vertical demand follows a seabed whose maximum reported slope is approximately 18° . Table 5.1 gives the principal parameters and Section D.1 the analytical definitions.

The plant is either the ideal kinematic integrator $\dot{\mathbf{q}} = \mathbf{u}\mathbf{h} + \mathbf{v}_c$ (label prefix K) or the six-degree-of-freedom REMUS 100 model driven by the PI cascade of Section 4.3 (label prefix D). The proposed law is evaluated in four self-configurations (surge form A or B $\times$ policy hard or smooth) whenever the run is designed to isolate the proposed method, and in a single main configuration (surge B with smooth γ) whenever the run is designed to compare the proposed law against baselines. Cases and their purpose are collected in Table 5.2.

Case K2, the ideal-plant five-method comparison, is reported in Section C.3 as supple-

Table 5.1: Reference paths used in the validation programme.

Key	Geometry	Mean $\ \dot{\mathbf{r}}\ $ [m/s]	T [s]	h [s]
			300 (K1) 600 (D1)	
Helix	descending 3-D helix, radius 50 m	1.256	400 (comparison)	0.05
Line	straight inclined 3-D line	1.166	300	0.05
Lawnmower	5-leg 2-D boustrophedon, U-turn radius 200 m	2.000	7255	0.10
Lemniscate3D	closed 3-D Geronon lemniscate, half-width 500 m	1.500	2035	0.05
Oceanscan	terrain-following survey pattern, altitude 25 m above seabed	2.009 (range 2.000–2.052)	7255	0.10

Table 5.2: Simulation case map.

Case	Plant	Methods	Purpose
K1	Ideal	Own (4 analysed + 1 exploratory)	Ideal-plant verification for the four analysed configurations; supplementary progression trace
K2	Ideal	Comparative	Outer-law transfer comparison, isolated from actuator dynamics
D1	REMUS	Own (4 analysed + 1 exploratory)	Command realisation plus a supplementary progression trace
D2	REMUS	5 methods (helix)	Proposed law, MED predecessor, path-following extension, Coates ALOS, and Degorre VRP
D3	REMUS	5 methods (others)	Repeated benchmark on line, lawnmower, lemniscate, and oceanscan
D4	REMUS	Own (4 analysed + 1 exploratory)	Separate terrain-focused own-method evaluation (oceanscan)

mentary guidance-level evidence. The main text emphasises the REMUS 100/PI transfer comparison, since it reflects the implemented vehicle study. The complete supplementary numerical material is collected in Chapter C.

The controlled-transfer protocol of Section D.10.1 governs every comparative run: identical vehicle, identical inner PI cascade, identical reference path, identical initial condition, identical integrator settings, and identical steady-window $[T_r, T]$ with $T_r = 0.15 T$. The constant NED ocean current used across the reported runs is $\mathbf{v}_c = [0.35, 0.22, 0.06]^\top$ m/s, oriented at approximately 32° east of north, with $\|\mathbf{v}_c\| \approx 0.42$ m s $^{-1}$; per-path initial-condition composition is reported in Section D.3. The numerical simulations were carried out in MATLAB 2025. The ideal cases use the reduced kinematic model, whereas the REMUS 100 dynamic cases use the Fossen MSS Toolbox implementation with the integration settings reported in Chapter D. Seven common metrics are used throughout the comparison, as summarised in Table 5.3. Scheduled-reference error, internal particle-relative error, and progress are retained only as architecture-specific quantities and are not ranked in one common column, because their definitions differ between prescribed- and autonomous-progression formulations. Two companion measures—the fitted exponential decay rate α_{conv} and the fraction of steady time inside fixed tolerance tubes—are reported only where the main table cannot express transient shape or tolerance persistence. The full metric definitions, projection convention, steady-window convention, and simulation-parameter record are collected in Section D.4 and Chapter D.

Table 5.3: Performance metrics used in Chapter 5.

Symbol	Meaning	Interpretation
d_φ^{RMS}	RMS geometric distance to the path	Main accuracy metric; lower means the vehicle remains closer to the geometric path, independently of target timing.
d_φ^{max}	Peak geometric distance to the path	Worst reported path-distance excursion on the reporting window.
T_s	Settling time to the path-dependent tube	First time after which d_φ remains inside $d_\varphi \leq \max(0.10 d_\varphi(0), 0.05 \text{ m})$.
J_ω	Integrated squared angular command	Outer-guidance angular effort over the steady reporting window.
$J_{\Delta u}$	Integrated surge-command roughness	Measures how rapidly the surge command changes; smaller values indicate smoother speed demands.
sat. [%]	Any-actuator saturation duty	Percentage of the steady window in which at least one actuator or rate limit is active.
$\ \tilde{\mathbf{v}}_c(T_r)\ $	Current-estimation error at the reporting-window onset	Reported only for implementations that include the thesis current observer.

The RMS and peak of d_φ are computed by projecting the vehicle position onto the reference over a search window of $\pm \min(0.25 T, 120 \text{ s})$ around each report instant; values below the reference-path discretisation spacing ($\approx 0.2 \text{ m}$) should not be interpreted as accurate beyond that resolution. This finite-window estimator can also overestimate the distance to the complete curve when an implementation develops a large along-path lag. The

tables therefore report the values produced by the common estimator, but sub-decimetre differences and large-lag cases are not used to claim an intrinsic ordering of the guidance laws. The external-benchmark protocol is described in greater detail in Sections D.10 and D.10.1.

5.2 Kinematic simulation results

Case K1 uses the ideal integrator $\dot{\mathbf{q}} = \mathbf{u}\mathbf{h} + \mathbf{v}_c$ to isolate the four analysed proposed-law configurations from actuator dynamics. The exploratory path-progression trace is shown only as contextual numerical information and is not part of the formal kinematic result. The trajectory and target-error results are followed by observer, alignment, and Lyapunov checks.

5.2.1 Trajectory and position error

Figure 5.1 shows the vehicle trajectory in three dimensions under the four analysed self-configurations of the proposed law, together with the exploratory progression variant, on the ideal kinematic plant, with the helix reference drawn as a reference curve. All five numerical traces converged onto the helix from the common stress-testing initial offset and remained within a small neighbourhood of the reference for the rest of the horizon. Figure 5.2 reports the norm of the outer position error $\|\mathbf{e}(t)\|$ against the reporting horizon; the four configurations reached a steady-state error norm in the range 0.013–0.055 m, and the two smooth- γ configurations recorded the tightest steady error. These two figures are the primary outer-loop evidence of the kinematic case: they show that the outer guidance law drives the target-relative error to a small neighbourhood of zero on the ideal integrator plant. Geometric path distance is evaluated separately and is not inferred from $\|\mathbf{e}\|$.




Figure 5.1: Vehicle trajectory on the helix under the four formally analysed self-configurations and the exploratory path-following variant (case K1).




Figure 5.2: Position-error norm on the helix (case K1).

5.2.2 Observer convergence

The current observer is the six-state Luenberger realisation (two vector states $\tilde{e}, \tilde{v}_c$) of Chapter 3 with $k_C = 1$, $k_D = 0.25$; the target is exponential convergence of the current-estimate error $\tilde{v}_c$ under bounded reference derivatives. Figure 5.3 reports the norm of the observer state against the critically damped envelope predicted by the observer lemma. Across the four self-configurations the post-transient current-estimate error at T_r settled to $\|\tilde{v}_c\| \approx 5 \times 10^{-10}$ m/s — negligible both against the true current magnitude and against the numerical reporting floor on the finite integration horizon — and the alignment Lyapunov

reached $V_a \approx 3.7 \times 10^{-7}$ before the start of the steady window. The observer therefore converges much faster than the outer loop in this run. Its transient is nevertheless retained explicitly in the layered proof of Section A.15; the small finite-horizon value is numerical evidence, not a replacement for that argument.




Figure 5.3: Current-observer convergence on the helix (case K1).

5.2.3 Alignment law

The alignment law drives the alignment Lyapunov V_a to zero according to the closed-form logistic solution of Chapter 3. The initial transient reproduced the closed-form logistic decay: the first-crossing alignment time was $t_{\text{align}} \approx 0.6$ s, corresponding to only 0.2% of the horizon, so the guidance layer entered the aligned regime long before any statistic on the steady window was computed. Subsequent changes of the target direction along the reference (curvature transients, blend–tangent switches) re-excite V_a above the closed-form envelope in localised pulses that remain small in absolute terms and decay back to the aligned regime; the quoted t_{align} is therefore an initial-crossing time rather than a global-in-time envelope.

The alignment channel is not reported here as a standalone figure: its contribution V_a is drawn as one of the three components of the composite Lyapunov panel of Figure 5.4 below.

5.2.4 Outer position consistency checks

The outer analysis depends on blend regularity, the target-projection inequality, and the sign quantity Ξ . These quantities are shown directly in Figure 5.4. The numerical sign check $\Xi \leq 0$ was satisfied on the complete reporting window for every internal configuration. The near-target decay rates were 0.126 s^{-1} for the hysteretic policy and 0.151 s^{-1} for the continuous policy. Together with the observer and alignment traces, these results are numerically consistent with the kinematic preconditions. They are not presented as a finite-horizon proof of the layered convergence theorem. The complete argument and its ideal-model boundary are given in Section A.15 and Theorem A.1; the supplementary numerical consistency trace retained outside the main chapter is reported in Section C.1.




Figure 5.4: Blend regularity, target projection, projection gap, and outer sign checks on the helix (case K1).

5.2.5 Empirical check of the conditional bound

The conditional ultimate-boundedness expression derived in Theorem E.17 is $r_{\text{UUB}}(\theta) = \lambda \bar{\delta}_{95} / [(1 - \theta) \underline{\gamma} v_r^{\min} - \bar{\delta}_{95}]$ with a margin $\theta \in [0, 1)$. In the recorded simulation data, $\bar{\delta}_{95}$ is the 95th percentile of the observed velocity difference rather than the uniform bound required by Theorem E.17. Consequently, substituting it in the formula is only a sensitivity check and cannot verify the theorem. Table 5.4 reports this check for the two Modulation B configurations. At $\theta = 0.05$, the observed maximum steady-window error is about 2.4 times the computed radius for the continuous policy and 8.4 times it for the hysteretic policy. Increasing θ enlarges the computed radius, but doing so after observing the response does not turn the calculation into a predictive stability claim. Verification would require a uniform bound chosen independently of the reported trajectory and satisfaction of all hypotheses of Theorem E.17, including the explicit outer-decay dominance assumption of Theorem E.3.

Table 5.4: Empirical UUB-radius sensitivity on the helix.

Configuration	Plant	$\rho_{\text{emp}}(0.05)$	$\rho_{\text{emp}}(0.30)$	$\rho_{\text{emp}}(0.50)$	$\rho_{\text{emp}}(0.90)$
M-B continuous	Kinematic	2.4	1.8	1.3	0.25
M-B hysteretic	Kinematic	8.4	6.2	4.4	0.87

Note. The ratio is $\rho_{\text{emp}} = e_{\text{ss,max}} / r_{\text{UUB}}(\theta)$ for case K1. Since $\bar{\delta}_{95}$ is a percentile rather than a uniform bound, values below unity are not stability proofs.

5.3 Kinematic–dynamic simulation results

This section reports the runs on the REMUS 100 plant driven by the PI cascade — the closed-loop system that the simulation realisation of the vehicle exercises. All results are obtained with the frame-consistent full-rotation current representation of (D.11). The helix own-method study (case D1) and the five-method comparison (case D2) both use the per-path offset $[10, -24, 12]$ m, for which $\|e_0\| \approx 29.03$ m. Case D1 runs for 600 s and reports five internal traces, whereas case D2 runs for 400 s and reports the common steady-window benchmark. Their metrics therefore answer different questions even though the helix initial condition is shared.

5.3.1 Trajectory following on the helix

Figure 5.5 shows the vehicle following the helical reference under the four analysed self-configurations and the exploratory progression variant with the REMUS 100 plant. All five traces converged onto the reference from the common initial offset and remained within a small neighbourhood of the reference for the rest of the horizon; the terminal position-error norms are collected in Table 5.5. The smooth- γ configurations (rows 3 and 4) recorded the smallest angular effort and no policy-switch events, in contrast to the hysteretic configurations, which showed frequent switches and a higher angular effort. Table 5.5 reports the terminal position-error norm, the average position-error norm over $[0.15T, T]$, the full-horizon angular effort, and the number of γ -switching events for each configuration.




Figure 5.5: Vehicle following the helix under four analysed self-configurations and the exploratory path-following variant (case D1).

Table 5.5: Own-method dynamic helix summary.

Configuration	$\ e(T)\ $ [m]	$\ e\ _{\text{avg}}$ [m]	J_{ω}	γ -switches
1 (M-A, hysteretic)	4.017	2.331	0.891	7
2 (M-B, hysteretic)	4.074	2.349	0.880	7
3 (M-A, continuous)	0.738	0.635	0.819	0
4 (M-B, continuous)	0.739	0.642	0.813	0
5 (M-B, continuous; autonomous progression)	0.686	0.573	1.408	0

5.3.2 Position error under the dynamic plant

Figure 5.6 reports the position-error norm $\|e(t)\|$ for the four analysed self-configurations and the exploratory progression variant under the REMUS/PI plant on the helix. The two smooth- γ configurations (3, 4) reached a steady-window neighbourhood of approximately 0.63–0.64 m; the exploratory trace averaged approximately 0.57 m. These values are those that Table 5.5 summarises; the two hysteretic configurations (1, 2) recorded average errors of approximately 2.3 m, consistent with the repeated switching visible in the trace. The peak error occurs during the initial approach and is bounded by $\|e_0\| \approx 29.03$ m.




Figure 5.6: Position-error norm under the dynamic plant (case D1, helix).

5.3.3 Command–response tracking on the inner loop

Because the outer guidance law delivers a surge command u_{cmd} and an angular-velocity command ω_{cmd} to the PI cascade of Section 4.3, the closed-loop realisation of the guidance intent has to be quantified through the difference between commanded and achieved body-frame quantities and through the actuator commands themselves. Additional helix response histories, including attitude, body velocity, body rates, and cumulative command activity, are collected in Section C.4. Figure 5.7 overlays commanded and plant-realised surge on the helix (case D1); the underlying figure reports the four analysed self-configurations and the exploratory progression variant of case D1 side by side, and the discussion below refers primarily to the M-B continuous trace, which is the reference variant taken forward to the external comparison. The surge tracking error decays to a small steady error after the transient window. Because a scalar command-realisation metric was not computed for this run, this observation is retained as a qualitative trace-level result and is listed as a limitation in Section 5.5. Figure 5.8 reports the corresponding rudder, stern-plane and propeller command traces against their saturation limits. On case D1, all five configurations have zero saturation duty on the reporting window $[0.15T, T]$ in the own-method data record. A separate post-alignment ISS audit, which uses a different window, records limited any-channel activity: 3.43% and 3.33% for the hysteretic configurations, 0.65% and 0.60% for the continuous configurations, and 1.76% for the exploratory variant. This distinction prevents a transient audit from being conflated with the benchmark saturation metric. Case D2 likewise reports zero steady-window saturation for the M-B continuous law (Table 5.8). Attribution of the remaining steady error observed in Table 5.8 to the guidance geometry rather than to inner-loop lag or actuator clipping would require a direct analysis of commanded versus realised motion, which is beyond the reporting scope of this chapter.




Figure 5.7: Surge command–response tracking (case D1, helix).




Figure 5.8: Actuator commands against saturation limits (case D1, helix).

5.3.4 Observer convergence under the dynamic plant

Under the REMUS plant the observer is driven by the DVL-aided water-relative velocity $\mathbf{v}_{\text{rel}}$ (Section 4.3.6); with the full-rotation current map of Section 4.2 the resulting error system is homogeneous in $\tilde{\mathbf{v}}_c$, so sway and heave enter the measurement $\mathbf{v}_{\text{rel}}$ rather than the observer error. Any nonzero observer error on the dynamic plant therefore reflects numerical integration, sample-and-hold, measurement construction, or plant–observer implementation differences rather than uncompensated sway/heave or the inner-loop motion itself. Figure 5.9 shows the observer converged under this architecture. Three complementary numbers are useful here and must not be confused: (i) in case D1 (M-B continuous), the terminal observer error was $\|\tilde{\mathbf{v}}_c(T)\| \approx 4.05 \times 10^{-6}$ m/s; (ii) its steady-window mean was 1.56×10^{-5} m/s, whereas the full-horizon RMS was 2.69×10^{-2} m/s because it includes the observer transient; (iii) in case D2 (M-B continuous, per-path benchmark), the observer error at the start of the steady window was $\|\tilde{\mathbf{v}}_c(T_r)\| \approx 3.76 \times 10^{-6}$ m/s. The terminal and steady-window values are more than four orders of magnitude smaller than the true current magnitude $\|\mathbf{v}_c\| \approx 0.42 \text{ m s}^{-1}$; the full-horizon RMS is not a steady-state measure. The transient of Figure 5.9 contains small departures from the ideal critically damped envelope during the initial approach and curvature transitions. Because full water-relative velocity is supplied, these departures cannot be attributed to unmeasured sway or heave; they are treated as finite-step and implementation errors without a stronger causal claim. The observer error decays to the values listed in (i)–(iii), so the observer output is usable as a current estimate throughout the reporting window. The observer is one of the two capabilities that distinguishes the proposed law from the two baselines considered in Section 5.4; the second is the runtime feasibility protocol of Section D.5, which was triggered on none of the reported runs.




Figure 5.9: Current-observer convergence under the dynamic plant (case D1, helix).

5.3.5 Terrain-following outcome on the survey reference

The terrain-following survey reference of Section D.1 combines the boustrophedon geometry of the lawnmower with a vertical demand set by the smoothed seabed at a nominal clearance of 25 m and exercises the horizontal and vertical channels simultaneously under the actuator envelope of the REMUS/PI plant. From a mission-application standpoint, it is the most representative single scenario in this thesis for the survey-class duty cycle in which current-aware guidance is intended to operate. The common-metric values reported here are taken from the M-B continuous execution in case D3. It recorded an RMS geometric distance of 0.099 m, a peak of 0.411 m, a settling time of 146.9 s, an integrated angular effort of 0.042, and 3.9% actuator saturation on the steady window (Table 5.6). Case D4 is a separate own-method run with the large-offset initial condition; its figures compare the four analysed internal configurations and the exploratory progression variant and are used for the spatial clearance history, not for the case D3 benchmark numbers. Figures 5.10 and 5.11 report the four analysed internal configurations and the exploratory progression variant: the vehicle trajectories over the seabed and the seabed clearance $h_{\text{clr}}(t) := z_{\text{bed}}(x(t), y(t)) - z_{\text{AUV}}(t)$. For the proposed method, the post-transient trace remains close to the commanded 25 m clearance; the approximate plotted range is 24.8–25.2 m. This plotted range is reported as such and is not presented as a clearance RMS, peak-error, or minimum-clearance statistic. The corresponding five-method oceanscan path-distance comparison is reported later in Section 5.4.4, with the full set of path-wise comparison plots collected in Section C.2.

Table 5.6: Terrain-following assessment.

Metric	Value
Path-distance RMS $d_{\mathcal{P}}^{\text{RMS}}$ [m]	0.099
Path-distance peak $d_{\mathcal{P}}^{\text{max}}$ [m]	0.411
Settling time T_s [s]	146.9
Angular effort J_{ω}	0.042
Any-channel saturation duty sat%	3.9
Post-transient observer error $\ \tilde{\mathbf{v}}_c(T_r)\ $ [m/s]	$\approx 2.0 \times 10^{-4}$
Desired seabed clearance [m]	25.0
Approximate post-transient clearance range [m]	24.8–25.2

Note. Path-distance, effort, saturation, and observer values are from case D3. The clearance interval is read from the separate case D4 post-transient trace and is reported as an estimate from the plotted clearance history rather than as a tabulated clearance statistic.




Figure 5.10: Trajectory over the seabed for the four analysed internal configurations and the exploratory progression variant on the survey reference (case D4).




Figure 5.11: Altitude, seabed cross-section, and depth histories for the four analysed internal configurations and the exploratory progression variant on the terrain-following survey (case D4).

5.4 Comparative evaluation

Five implementations are evaluated through the shared REMUS 100/PI interface: the continuous M-B law developed in this thesis, referred to as *M-B continuous-PP* for the *prescribed-progression formulation*; its switching MED predecessor [4]; an exploratory state-dependent path-progression extension referred to as *M-B continuous-PF* for the *autonomous path-progression formulation*; spherical 3-D ALOS [5]; and 3-D VRP [6]. Plant, path, current, initial condition, integration, and reporting window are common. The external methods retain their published outer-guidance gains where available. Consequently, the experiment is a controlled transfer study of the complete outer-law–adapter implementations, not a reproduction or intrinsic ranking of the native closed loops. The progression extension M-B continuous-PF is included to expose its present numerical limits; it is not part of the formal contribution established in Chapter 3, which concerns the prescribed-progression formulation M-B continuous-PP. The main text reports the REMUS 100/PI comparison tables and selected path-distance histories; the ideal-plant comparison and additional path-wise plots are provided in Appendices C.3 and C.2.

Table 5.7: Compared implementations.

Label	Meaning	Role in the comparison
M-B continuous-PP	Smooth Modulation-B law with prescribed moving-reference progression	Main thesis method; this is the formulation analysed in Chapter 3.
MED switching predecessor	Published two-task controller with hysteretic switching [4]	Internal predecessor used to assess the benefit of replacing the switch by a continuous blended law.
M-B continuous-PF	Exploratory autonomous path-progression variant of the same framework	Supplementary thesis-lineage comparator; not claimed as a validated contribution.
Coates 3-D ALOS	Spherical adaptive line-of-sight baseline [5]	External three-dimensional guidance baseline under the shared transfer interface.
Degorre 3-D VRP	Three-dimensional virtual-reference-point baseline [6]	External virtual-target baseline under the shared transfer interface.

5.4.1 Main benchmark on the helix

Table 5.8 reports the common geometric and implementation metrics for case D2. The architecture-dependent along-track and coverage variables are intentionally excluded from the main comparison: their definitions differ between prescribed and autonomous progression and therefore do not support a like-for-like ranking.

On the helix, the continuous thesis law settles in a sub-metre neighbourhood with low angular effort and zero steady-window saturation. Degorre VRP gives the smallest reported RMS and settling time on this path. The exploratory M-B continuous-PF trace remains close in geometric accuracy, but with lower coverage of the prescribed progression. The

Table 5.8: Five-method helix benchmark.

Method	$d_{\mathcal{P}}^{\text{RMS}}$ [m]	$d_{\mathcal{P}}^{\text{max}}$ [m]	T_s [s]
M-B continuous-PP	<i>0.431</i>	<i>0.708</i>	27.8
MED switching predecessor	2.740	4.415	390.4
M-B continuous-PF	<i>0.645</i>	0.941	31.1
Coates 3-D ALOS	—	1.811	∞
Degorre 3-D VRP	0.169	0.337	17.8

Method	J_{ω}	$J_{\Delta u}$	sat. [%]	$\ \tilde{\mathbf{v}}_c(T_r)\ $ [m/s]
M-B continuous-PP	0.217	<i>0.013</i>	0.0	3.2×10^{-3}
MED switching predecessor	6.973	31.734	5.7	2.1×10^{-3}
M-B continuous-PF	0.212	0.014	0.0	—
Coates 3-D ALOS	0.177	15.769	23.8	—
Degorre 3-D VRP	0.219	0.001	0.0	—

Note. Case D2. An em-dash denotes a non-finite RMS entry or an implementation without the thesis current observer, depending on the column. Bold marks the best value in each column; italic marks the best thesis-lineage value when the overall best is a non-thesis-lineage method.

MED predecessor is less accurate and much rougher, while the ALOS transfer does not enter the settling tube within the simulated horizon. For the proposed law, the observer error at the start of the reporting window is $\|\tilde{\mathbf{v}}_c(T_r)\| \approx 3.2 \times 10^{-3} \text{ m s}^{-1}$, compared with $\|\mathbf{v}_c\| \approx 0.42 \text{ m s}^{-1}$.

5.4.2 Trajectory overlay and path-distance time history

The trajectory and path-distance histories in Figures 5.12 and 5.13 show all five implementations. The thesis law, path-following extension, and VRP remain in a sub-metre neighbourhood after the transient. The MED and ALOS traces expose, respectively, the cost of the switching predecessor and the limitations of the shared ALOS transfer on this case.




Figure 5.12: Five-method trajectory overlay on the helix (case D2).




Figure 5.13: Five-method geometric path distance on the helix, shown on linear and logarithmic scales (case D2).

5.4.3 Path-distance behaviour on the lawnmower

On the lawnmower, the thesis implementation gives the smallest reported RMS path distance and finite settling under the shared REMUS/PI transfer. The exploratory M-B continuous-PF trace is also sub-metre in RMS, but it progresses more slowly along the reference. The MED predecessor remains within a few metres, while the ALOS and VRP transfers are less favourable in this particular survey geometry (Figure 5.14).




Figure 5.14: Five-method geometric path distance on the lawnmower (case D3).

5.4.4 Path-distance behaviour on the terrain-following survey

The oceanscan reference is the mission-representative survey scenario of Section 5.3.5, evaluated here under the same five-method comparison protocol as the other geometries. Figure 5.15 reports the geometric path-distance history for the five implementations. The thesis law M-B continuous-PP gives the lowest reported RMS distance in this case ($d_p^{\text{RMS}} = 0.099$ m, from Table 5.9). The exploratory M-B continuous-PF trace also remains sub-metre in RMS but with lower coverage of the prescribed progression. The MED predecessor remains near the few-metre range, while the ALOS and VRP transfers are less accurate on this terrain-following survey under the shared interface. The pattern is close to that of the lawnmower (Figure 5.14), consistent with the oceanscan reference sharing the boustrophedon

geometry of the lawnmower with an additional vertical demand set by the seabed profile.




Figure 5.15: Five-method geometric path distance on the terrain-following survey (case D3).

5.4.5 Performance across reference geometries

Repeating the five-method benchmark across the remaining reference geometries (case D3) exposes both the strengths and the limitations of the transferred implementations. To keep the tables readable, Tables 5.9 and 5.10 report RMS distance and settling time separately. In both tables, bold marks the best value in each column, and italic marks the best thesis-lineage value when the overall best is a non-thesis-lineage method. The full path-distance histories for the line and lemniscate, which are not reproduced in the main text, are included with the other path-wise comparison plots in Appendix C.2.

Four observations are supported by the common metrics. First, the continuous thesis law remains below 0.45 m RMS on the four curved and survey references and at 0.970 m on the line. Second, Degorre VRP gives the lower reported RMS on the helix and lemniscate, whereas the thesis law gives the lower reported RMS on the line, lawnmower, and terrain-following survey under the same transfer interface. Third, the MED predecessor is consistently less accurate and much rougher than the continuous thesis law, showing the benefit of the smooth architecture in this implementation. Fourth, the autonomous-progression extension remains a useful exploratory comparator: it is close to the thesis law on the helix and remains sub-metre

Table 5.9: RMS geometric path distance on the dynamic plant [m].

Method	Helix	Line	Lawnmower	Lemniscate3D	Oceanscan
M-B continuous-PP	<i>0.431</i>	0.970	0.095	<i>0.118</i>	0.099
MED switching predecessor	2.740	2.368	2.809	2.987	3.028
M-B continuous-PF	0.645	5.371	0.185	0.292	0.193
Coates 3-D ALOS	–	4.714	1.049	0.215	0.976
Degorre 3-D VRP	0.169	9.295	2.728	0.081	2.746

Table 5.10: Settling time on the dynamic plant [s].

Method	Helix	Line	Lawnmower	Lemniscate3D	Oceanscan
M-B continuous-PP	27.8	48.5	144.8	125.3	146.9
MED switching predecessor	390.4	43.9	141.7	125.2	143.6
M-B continuous-PF	31.1	115.4	171.4	203.2	172.6
Coates 3-D ALOS	∞	67.2	177.8	195.4	178.3
Degorre 3-D VRP	17.8	∞	156.8	95.7	159.9

Note. The symbol ∞ means that the implementation did not enter and remain in the prescribed path-dependent tube within the simulated horizon. An em-dash in the RMS table denotes a non-finite reported RMS value under the common geometric-distance estimator.

on several curved/survey cases, but its coverage and settling behaviour are less consistent, especially on the line and long missions. The non-ranked coverage indicator supports this reading: M-B continuous-PP remains essentially at the prescribed progression on all five paths (99.9–100.0%), whereas M-B continuous-PF covers only 89.3–95.0% of the reference schedule in the corresponding dynamic runs. ALOS also shows a path-dependent transfer outcome. Values below the declared projection resolution are not used for sub-decimetre ranking.

Table 5.11: Cross-method interpretation.

Method	Strongest observed behaviour	Main limitation in this study
M-B continuous-PP	Finite settling on all five dynamic paths, lowest RMS on three of the five paths, and low command roughness	It is not the lowest-RMS method on the helix or lemniscate, and moderate steady-window saturation appears on the line and survey paths.
MED switching predecessor	Finite dynamic settling on all five paths	Larger geometric error and substantially larger command roughness than the continuous thesis law.
M-B continuous-PF	Comparable to M-B continuous-PP on the helix and sub-metre in RMS on several curved/survey paths	Lower progression coverage and slower settling than the prescribed-progression law; high effort and saturation on the line.
Coates 3-D ALOS	Moderate RMS on lawnmower, lemniscate, and oceanscan under the shared transfer interface	Non-finite reported RMS on the helix, high saturation on the line, and no explicit current-vector observer in this implementation.
Degorre 3-D VRP	Lowest reported RMS on the helix and lemniscate	Larger RMS or command effort on the line, lawnmower, and oceanscan under the shared transfer interface.

Note. The table summarises the reported transfer experiments only; it is not a claim about the native closed-loop tuning of the external methods.

5.4.6 Note on the straight-line dynamic behaviour

On the line, the thesis law achieves $T_s = 48.5$ s, $d_\phi^{\text{RMS}} = 0.970$ m, and 8.2% any-channel saturation duty. The MED predecessor settles slightly faster but with larger RMS and much higher surge-command roughness. The path-following extension is less accurate and uses considerably greater angular effort. ALOS settles but spends most of the reporting window at an actuator or rate limit, whereas VRP does not enter the tube and demands very large angular effort. This is a trade-off comparison under the shared adapter, not a universal ordering of the native methods.

5.5 Discussion, limitations, and interpretation

Balanced interpretation. The results support observer convergence, nominal alignment, and the recorded outer sign condition on both plant levels. They also show sub-metre RMS for the thesis law on all five paths, a substantially smoother response than the MED predecessor, and a favourable straight-line trade-off for the thesis implementation. The path-following extension M-B continuous-PF is informative as an autonomous path-progression variant, but its coverage and settling behaviour are less consistent than those of M-B continuous-PP. The comparison is a shared-interface transfer study with fixed gains, so the results are interpreted for the implemented adapters rather than as a method-independent ordering of the native controllers. The feasibility flag was not triggered in the reported runs.

Improvement over the MED predecessor. The strongest defensible improvement supported by the numerical evidence is that of M-B continuous-PP over its own switching MED predecessor [4] under the shared inner loop. The RMS geometric path distance drops from 2.740 m to 0.431 m on the helix ($\approx 84\%$ reduction), from 2.368 m to 0.970 m on the line ($\approx 59\%$), from 2.809 m to 0.095 m on the lawnmower ($\approx 97\%$), from 2.987 m to 0.118 m on the lemniscate ($\approx 96\%$), and from 3.028 m to 0.099 m on the terrain-following survey ($\approx 97\%$). Replacing the switching policy with the continuous law also removes all discrete γ -switch events, drops the surge-command roughness $J_{\Delta u}$ by factors ranging from approximately 1.6×10^2 to 3.8×10^4 across the five paths (e.g. $31.734 \rightarrow 0.013$ on the helix and $95.607 \rightarrow 0.048$ on the survey), and lowers the integrated angular effort J_{ω} on all curved and survey references. The saturation comparison is path-dependent: the continuous law removes steady-window saturation on the helix and lemniscate, reduces it on the lawnmower and survey, and incurs a larger line saturation duty than the predecessor. Against the two external baselines, the outcome is path dependent: Degorre VRP records lower RMS on the helix and lemniscate but is less favourable on the line, lawnmower, and survey under the shared interface, whereas Coates ALOS is strongly path-dependent. The five-path evidence therefore supports *a clear and consistent improvement over the MED switching predecessor, together with a balanced accuracy–smoothness–feasibility trade-off under the shared REMUS/PI adapter*; it does not imply lower geometric RMS than every external method on every path.

On the ISS/UUB status under the REMUS plant. On the REMUS 100 plant, the observer and alignment preconditions continue to hold and the outer sign condition $\Xi \leq 0$ remains satisfied on 100% of the reporting window, but the *a priori* Nomoto-based inner dissipativity assumption underpinning the composite proof of Appendix E is violated on 98.9–100.0% of the off-saturation window. The empirical percentile-based check of Table 5.4 places the observed steady-window maximum inside the closed-form radius for both continuous-policy configurations (≈ 0.75 – 0.76 m against ≈ 1.20 – 1.21 m at $\theta = 0.05$) and outside it for both hysteretic configurations (≈ 4.12 – 4.18 m against ≈ 1.63 m). Since the underlying dissipativity assumption is not empirically supported on the full six-degree-of-freedom plant, this empirical satisfaction is interpreted as a numerical consistency check rather than as a consequence of the theoretical formula; a complete cascade stability proof for the REMUS/PI implementation is not established here (Appendix E).

Scope. The runs use one constant current, one initial condition per path, and noise-free sensors. No equal-budget retuning of the external baselines, parametric envelope in current/initial offset, or data-driven comparison is provided [5, 6]. These limits, together with the unresolved cascade condition above, define the scope of the comparative statements in this chapter.

The command–response figures establish qualitative tracking of the outer requests, but the reported data do not provide one common scalar measure of water-relative velocity and body-rate realisation. Likewise, the oceanscan clearance interval is estimated from the plotted clearance history rather than from a dedicated clearance statistics table. These two results are therefore implementation evidence; quantitative interface error and clearance

RMS, peak error, minimum clearance, and time-in-band would be needed for a more detailed engineering assessment.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

This thesis developed a current-aware three-dimensional direction–speed architecture for an underactuated AUV. The formal result concerns a moving reference with externally prescribed path progression. On the admissible domain, a smooth blended direction replaces the predecessor’s task switch, while signed surge modulation removes the identified tangential equilibrium in the ideal feasible channel. In the reduced ideal kinematic model, the current-observer error converges exponentially, the direction error converges from every non-antipodal initial direction, and the continuous M-B law conditionally drives the moving-reference position error to zero. These conclusions require forward invariance of the admissible domain, surge feasibility, bounded reference derivatives, and exact ideal direction dynamics. A full-attitude current map and command interface then connect the kinematic requests to a six-degree-of-freedom REMUS 100/PI model without interpreting that implementation as a dynamic-cascade proof.

The numerical study supports the intended distinction between the analytical and implementation levels. The current observer converges to a negligible error on the ideal plant and to a small error on the REMUS model. Across the five dynamic paths, the selected prescribed-progression implementation achieves RMS geometric path distances of 0.095–0.431 m on the four curved or survey references and 0.970 m on the inclined line. The clearest supported numerical conclusion is the improvement over the published switching predecessor under the shared inner loop: the RMS geometric error drops by approximately 59–97% across the five paths, the discrete γ -switch is removed, the surge-command roughness $J_{\Delta u}$ falls by factors from approximately 1.6×10^2 to 3.8×10^4 , and steady-window saturation is zero on the helix and lemniscate and remains below 8.2% on the other paths. The comparison with the two external three-dimensional baselines is path-dependent. The virtual-reference-point baseline gives lower reported RMS on the helix and lemniscate, whereas the thesis law gives lower RMS on the line, lawnmower, and terrain-following survey under the shared REMUS/PI interface. The spherical ALOS baseline shows a more path-dependent transfer in the same interface.

The exploratory autonomous path-progression variant is useful as a supplementary comparison, but it is not a formal contribution of the thesis. It matches the prescribed-progression implementation most closely on the helix and remains sub-metre in RMS on several curved and survey references, but its progression coverage and settling behaviour are less consistent, especially on the line and long missions. The oceanscan case confirms sub-metre geometric tracking of the prescribed-progression law with low steady-window saturation; the plotted clearance history remains close to the nominal 25 m terrain-clearance

command, but the present data do not yet provide RMS, peak-error, or minimum-clearance statistics. The results support the smooth current-aware architecture within these limits: the comparison against external baselines is implementation-specific, the projection resolution bounds sub-decimetres claims, the clearance range is estimated from the plotted clearance history, and a complete REMUS/PI cascade theorem is not established here.

6.2 Future work

A natural extension is a rigorous autonomous path-coordinate version of the prescribed-progression law, denoted M-B continuous-PF. In this thesis it is treated as an extra progression layer built on the same direction–speed structure, not as the main contribution. Its progression law has a valid ideal-channel Lyapunov motivation and gives useful numerical behaviour on several paths, but it does not yet provide the same consistency as M-B continuous-PP across all reference geometries, especially on the inclined line and long survey missions. Before it can be claimed as a formal contribution, the formulation must include the observer and alignment transients, the reduced angular-command approximation, forward-progress and endpoint conventions, closed-path lap handling, and the limited regularity of the survey paths. The formal result of this thesis therefore remains the prescribed-progression law.

Further work should also establish a verified inner-loop stability condition for the REMUS/PI cascade and repeat the study with continuous path projection, multiple initial conditions and current directions, realistic DVL noise and delay, and quantitative command-realisation and terrain-clearance metrics. Hardware-in-the-loop and experimental AUV trials are then required to determine whether the simulated accuracy–effort trade-offs persist in practice.

Appendix A

Supplementary Derivations

A.1 Derivative of a normalized vector

For a differentiable nonzero vector $\mathbf{x}$ and $\bar{\mathbf{x}} = \mathbf{x}/\|\mathbf{x}\|$,

$$\dot{\bar{\mathbf{x}}} = \frac{1}{\|\mathbf{x}\|}(\mathbf{I} - \bar{\mathbf{x}}\bar{\mathbf{x}}^\top)\dot{\mathbf{x}}. \quad (\text{A.1})$$

Because $(\bar{\mathbf{x}} \times \dot{\mathbf{x}}) \times \bar{\mathbf{x}} = (\mathbf{I} - \bar{\mathbf{x}}\bar{\mathbf{x}}^\top)\dot{\mathbf{x}}$, this can be written as

$$\dot{\bar{\mathbf{x}}} = \boldsymbol{\omega}_x \times \bar{\mathbf{x}}, \quad \boldsymbol{\omega}_x = \frac{\bar{\mathbf{x}} \times \dot{\mathbf{x}}}{\|\mathbf{x}\|}. \quad (\text{A.2})$$

The angular velocity is not unique because any component parallel to $\bar{\mathbf{x}}$ produces no change. The expression above selects the perpendicular, minimum-norm representative used in (3.7) and (3.42).

A.2 Reduced angular-command inversion and approximation errors

The full derivative of the estimated inertial velocity is

$$\dot{\mathbf{v}}_1 = \dot{u}\mathbf{h} + u(\boldsymbol{\omega} \times \mathbf{h}) + \dot{\mathbf{v}}_c. \quad (\text{A.3})$$

Under the stated time-scale reduction, the first and third terms are neglected while the command is calculated, so that

$$\dot{\mathbf{v}}_1 \simeq u(\boldsymbol{\omega} \times \mathbf{h}), \quad \boldsymbol{\omega}_{v_1} \simeq \frac{u}{\|\mathbf{v}_1\|} \bar{\mathbf{v}}_1 \times (\boldsymbol{\omega} \times \mathbf{h}). \quad (\text{A.4})$$

The command $\boldsymbol{\omega}$ is constrained to be roll free, so only the component perpendicular to $\mathbf{h}$ can be assigned. Therefore the desired direction-rate vector is first projected onto the admissible plane orthogonal to $\mathbf{h}$. Using

$$\bar{\mathbf{v}}_1 \times (\boldsymbol{\omega} \times \mathbf{h}) = \boldsymbol{\omega}(\bar{\mathbf{v}}_1^\top \mathbf{h}) - \mathbf{h}(\bar{\mathbf{v}}_1^\top \boldsymbol{\omega}) \quad (\text{A.5})$$

and matching the projected reduced direction rate gives

$$\frac{u(\mathbf{h}^\top \bar{\mathbf{v}}_1)}{\|\mathbf{v}_1\|} \boldsymbol{\omega} = (\mathbf{I} - \mathbf{h}\mathbf{h}^\top) \boldsymbol{\omega}_{\bar{\mathbf{v}}}^{\text{des}}, \quad (\text{A.6})$$

because the physical command is roll free, $\omega^\top \mathbf{h} = 0$. For $u > 0$ and $\mathbf{h}^\top \bar{\mathbf{v}}_1 > 0$, the adopted projected command is therefore

$$\omega = \frac{\|\mathbf{v}_1\|}{u(\mathbf{h}^\top \bar{\mathbf{v}}_1)} (\mathbf{I} - \mathbf{h}\mathbf{h}^\top) \omega_{\bar{\mathbf{v}}}^{\text{des}}. \quad (\text{A.7})$$

Substituting this command back into the reduced unit-direction kinematics gives, within that reduced model,

$$\omega_{\mathbf{v}_1} = \omega_{\bar{\mathbf{v}}}^{\text{des}} - \mathbf{h} \frac{\bar{\mathbf{v}}_1^\top \omega_{\bar{\mathbf{v}}}^{\text{des}}}{\mathbf{h}^\top \bar{\mathbf{v}}_1}. \quad (\text{A.8})$$

Thus the adopted physical command matches the desired direction rate exactly only when $\bar{\mathbf{v}}_1^\top \omega_{\bar{\mathbf{v}}}^{\text{des}} = 0$. This condition is automatic for the corrective term $k(\bar{\mathbf{v}}_1 \times \bar{\mathbf{a}})$ but not generally for the transported feedforward $\omega_{\bar{\mathbf{a}}}$ away from alignment.

Restoring the two explicitly omitted derivative terms identifies the time-scale approximation error retained in the adopted controller derivation,

$$\boldsymbol{\varepsilon}(t) = \frac{1}{\|\mathbf{v}_1\|} \bar{\mathbf{v}}_1 \times (\dot{u}\mathbf{h} + \dot{\mathbf{p}}_c). \quad (\text{A.9})$$

Combining (A.8) with this time-scale term gives (3.51). The projection discrepancy and the time-scale approximation error are neither additional control inputs nor independently designed error channels. The integrability of the $\dot{u}$ contribution in $\boldsymbol{\varepsilon}(t)$ is an explicit assumption in (3.54), not a proved closed-loop consequence.

A.3 Blend simplification and singularity

For $\mathbf{e} \neq \mathbf{0}$, the weighted direction suggested by $p = \lambda/(\lambda + \|\mathbf{e}\|)$ satisfies

$$(1 - p)\bar{\mathbf{e}} + p\mathbf{b} = \frac{\|\mathbf{e}\|}{\lambda + \|\mathbf{e}\|} \frac{\mathbf{e}}{\|\mathbf{e}\|} + \frac{\lambda}{\lambda + \|\mathbf{e}\|} \mathbf{b} \quad (\text{A.10})$$

$$= \frac{\mathbf{e} + \lambda\mathbf{b}}{\lambda + \|\mathbf{e}\|}. \quad (\text{A.11})$$

The last expression is adopted as the definition of $\bar{\mathbf{a}}$ and extends continuously to $\mathbf{a} = \mathbf{b}$ at $\mathbf{e} = \mathbf{0}$. Its scalar denominator is strictly positive. Therefore $\mathbf{a} = \mathbf{0}$ if and only if $\mathbf{e} = -\lambda\mathbf{b}$, and away from this point

$$\bar{\mathbf{a}} = \frac{\mathbf{e} + \lambda\mathbf{b}}{\lambda + \|\mathbf{e}\|}. \quad (\text{A.12})$$

This identity is preferable in software because it avoids separately normalizing $\mathbf{e}$ near the origin.

A.4 Positive surge root

Enforcing $\|u\mathbf{h} + \hat{\mathbf{v}}_c\| = \gamma \|\dot{\mathbf{r}}\|$ gives

$$u^2 + 2(\mathbf{h}^\top \hat{\mathbf{v}}_c)u + \|\hat{\mathbf{v}}_c\|^2 - \gamma^2 \|\dot{\mathbf{r}}\|^2 = 0. \quad (\text{A.13})$$

Hence

$$u_{\pm} = -\mathbf{h}^\top \hat{\mathbf{v}}_c \pm \sqrt{(\mathbf{h}^\top \hat{\mathbf{v}}_c)^2 - \|\hat{\mathbf{v}}_c\|^2 + \gamma^2 \|\dot{\mathbf{r}}\|^2}. \quad (\text{A.14})$$

When $\gamma \|\dot{\mathbf{r}}\| > \|\hat{\mathbf{v}}_c\|$, the radicand is positive and its square root exceeds $|\mathbf{h}^\top \hat{\mathbf{v}}_c|$. It follows that $u_+ > 0$ and $u_- < 0$, so (3.72) is the unique positive solution under this sufficient feasibility condition. The exact feasibility condition is non-negativity of the radicand together with positivity of the selected root.

A.5 Alignment derivative

Let $c = \bar{\mathbf{a}}^\top \bar{\mathbf{v}}_1$ and $V = (1 - c)/2$. Under

$$= [\omega_{\bar{\mathbf{a}}} + k(\bar{\mathbf{v}}_1 \times \bar{\mathbf{a}})] \times \bar{\mathbf{v}}_1, \quad = \omega_{\bar{\mathbf{a}}} \times \bar{\mathbf{a}}, \quad (\text{A.15})$$

the common transport term cancels in $\dot{c}$. The remaining term is

$$\dot{c} = k\bar{\mathbf{a}}^\top [(\bar{\mathbf{v}}_1 \times \bar{\mathbf{a}}) \times \bar{\mathbf{v}}_1] \quad (\text{A.16})$$

$$= k(1 - c^2). \quad (\text{A.17})$$

Thus

$$\dot{V} = -\frac{k}{2}(1 - c^2) = -2kV(1 - V). \quad (\text{A.18})$$

For $V(0) < 1$, separation of variables yields

$$V(t) = \frac{V(0)e^{-2kt}}{1 - V(0) + V(0)e^{-2kt}}. \quad (\text{A.19})$$

The exceptional value $V(0) = 1$ corresponds to $\bar{\mathbf{v}}_1(0) = -\bar{\mathbf{a}}(0)$ and remains invariant in the ideal alignment model.

Specialisation to Task 1

For Task 1 in Section 3.2, the desired direction is the target-relative direction $\bar{\mathbf{e}}$ and the Lyapunov function is

$$V_1 = \frac{1}{2}(1 - \bar{\mathbf{e}}^\top \bar{\mathbf{v}}_1). \quad (\text{A.20})$$

Define

$$c_1 := \bar{\mathbf{e}}^\top \bar{\mathbf{v}}_1, \quad V_1 = \frac{1}{2}(1 - c_1). \quad (\text{A.21})$$

On the ideal direction manifold with a converged current estimate, the implemented direction dynamics are

$$= [\omega_e + k_1(\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}})] \times \bar{\mathbf{v}}_1, \quad \dot{\bar{\mathbf{e}}} = \omega_e \times \bar{\mathbf{e}}. \quad (\text{A.22})$$

Differentiating c_1 gives

$$\begin{aligned} \dot{c}_1 &= \dot{\bar{\mathbf{e}}}^\top \bar{\mathbf{v}}_1 + \bar{\mathbf{e}}^\top \\ &= (\omega_e \times \bar{\mathbf{e}})^\top \bar{\mathbf{v}}_1 + \bar{\mathbf{e}}^\top ([\omega_e + k_1(\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}})] \times \bar{\mathbf{v}}_1) \\ &= (\omega_e \times \bar{\mathbf{e}})^\top \bar{\mathbf{v}}_1 + \bar{\mathbf{e}}^\top (\omega_e \times \bar{\mathbf{v}}_1) + k_1 \bar{\mathbf{e}}^\top ((\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}}) \times \bar{\mathbf{v}}_1). \end{aligned} \quad (\text{A.23})$$

The first two terms cancel by the scalar triple-product identity. For the remaining term, using $\|\bar{\mathbf{e}}\| = \|\bar{\mathbf{v}}_1\| = 1$ gives

$$\begin{aligned} \bar{\mathbf{e}}^\top ((\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}}) \times \bar{\mathbf{v}}_1) &= \bar{\mathbf{e}}^\top (\bar{\mathbf{e}} - (\bar{\mathbf{e}}^\top \bar{\mathbf{v}}_1) \bar{\mathbf{v}}_1) \\ &= 1 - c_1^2 \\ &= \|\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}}\|^2. \end{aligned} \quad (\text{A.24})$$

Therefore

$$\dot{c}_1 = k_1 \|\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}}\|^2, \quad (\text{A.25})$$

and hence

$$\dot{V}_1 = -\frac{1}{2} \dot{c}_1 = -\frac{k_1}{2} \|\bar{\mathbf{v}}_1 \times \bar{\mathbf{e}}\|^2 \leq 0. \quad (\text{A.26})$$

A.6 Proof of the target-projection inequality

Using the virtual look-ahead target representation and $s = \mathbf{e}^\top \mathbf{b}$ gives

$$\Pi_e = \frac{\|\mathbf{e}\|^2 + \lambda s}{\|\mathbf{e} + \lambda \mathbf{b}\|}. \quad (\text{A.27})$$

Set $r := \|\mathbf{e}\|$ and $D := \|\mathbf{e} + \lambda \mathbf{b}\|$. If $s < 0$, then $\Pi_e - s = (r^2 + s(\lambda - D))/D$. Since $D \geq |r - \lambda|$, one has $\lambda - D \leq r$; multiplying by $s < 0$ reverses the inequality and gives $s(\lambda - D) \geq sr \geq -r^2$. Hence $\Pi_e - s \geq 0$.

If $s \geq 0$, both $r^2 + \lambda s$ and sD are nonnegative. Squaring their comparison and using $D^2 = r^2 + 2\lambda s + \lambda^2$ yields

$$(r^2 + \lambda s)^2 - s^2 D^2 = (r^2 - s^2)(r^2 + 2\lambda s) \geq 0, \quad (\text{A.28})$$

because $|s| \leq r$. Therefore $\Pi_e \geq s$ in both cases. Equality requires $r^2 = s^2$, so $\mathbf{e}$ is collinear with $\mathbf{b}$; substitution in the nonsingular virtual-target expression gives $\mathbf{e} = \ell \mathbf{b}$ with $\ell > -\lambda$. Conversely, every such vector gives $\Pi_e = s$. The value $\ell = -\lambda$ is precisely the excluded cancellation point.

A.7 Mode-wise proof for the hysteresis policy

Under ideal alignment, a converged current estimate, and Modulation B, (3.75) reduces to

$$\dot{V}_e = \|\dot{\mathbf{r}}\| (s - \gamma \Pi_e), \quad \Pi_e \geq s. \quad (\text{A.29})$$

In Mode A with $s > 0$, $\gamma = \gamma_+ > 1$, and hence

$$s - \gamma_+ \Pi_e \leq (1 - \gamma_+)s < 0. \quad (\text{A.30})$$

If $s = 0$ and $\mathbf{e} \neq \mathbf{0}$, the strict part of the projection inequality gives $\Pi_e > 0$, so the derivative remains negative. In Mode A with $s < 0$, $\gamma = \gamma_- < 1$. Since multiplication by the positive γ_- preserves $\Pi_e \geq s$,

$$s - \gamma_- \Pi_e \leq (1 - \gamma_-)s < 0. \quad (\text{A.31})$$

Thus Mode A is strictly decreasing away from the origin.

In Mode B with $\gamma_{\text{hold}} = 1$, (A.29) becomes $\dot{V}_e = \|\dot{\mathbf{r}}\| (s - \Pi_e) \leq 0$. Equality holds exactly for $\mathbf{e} = \ell \mathbf{b}$, $\ell > -\lambda$. Intersecting this equality family with the Mode B flow domain $\|\mathbf{e}\| \leq \rho_{\text{out}}$ gives the set stated in Theorem 3.12. If $\gamma_{\text{hold}} > 1$ and $s \geq 0$, the same calculation used for γ_+ gives strict decrease for $\mathbf{e} \neq \mathbf{0}$. For $s < 0$, the sign is not fixed by $\gamma_{\text{hold}} \geq 1$, which is why no broader Mode B claim is made.

A.8 Continuous-policy position decay

For the continuous law $\gamma(s) = 1 + \eta \tanh(s/\alpha_0)$, one has $\gamma > 1$ when $s > 0$, $\gamma < 1$ when $s < 0$, and $\gamma = 1$ when $s = 0$. Under the hypotheses of Theorem 3.14,

$$\dot{V}_e = \|\dot{\mathbf{r}}\| (s - \gamma(s) \Pi_e). \quad (\text{A.32})$$

If $s > 0$, the projection inequality yields

$$s - \gamma \Pi_e \leq (1 - \gamma)s < 0. \quad (\text{A.33})$$

If $s < 0$, then $\gamma \Pi_e \geq \gamma s$ and again

$$s - \gamma \Pi_e \leq (1 - \gamma)s < 0. \quad (\text{A.34})$$

Finally, if $s = 0$ and $\mathbf{e} \neq \mathbf{0}$, the strict part of Theorem 3.9 gives $\Pi_e > 0$, so $s - \gamma \Pi_e = -\Pi_e < 0$. At $\mathbf{e} = \mathbf{0}$, the virtual-target extension gives $V_e = 0$ and $\dot{V}_e = 0$. This proves the negative-definite statement without asserting a uniform decay rate.

A.9 Uniform-decay condition for asymptotic position convergence

Let $r_{\max} > 0$ bound $\|\mathbf{e}\|$ along the trajectory and define the augmented compact set

$$\mathcal{K}_e := \{(\mathbf{e}, \mathbf{b}) : \|\mathbf{e}\| \leq r_{\max}, \|\mathbf{b}\| = 1, \|\mathbf{a}(\mathbf{e}, \mathbf{b})\| \geq a_{\min}\}. \quad (\text{A.35})$$

This set includes the time-varying tangent explicitly and excludes the blend singularity uniformly. Suppose also that the reference speed is bounded below by $v_r^{\min} > 0$. Define

$$G(\mathbf{e}, \mathbf{b}) := \gamma(\mathbf{e}^\top \mathbf{b}) \Pi_e - \mathbf{e}^\top \mathbf{b}. \quad (\text{A.36})$$

By Section A.8, G is continuous on $\mathcal{K}_e$, is nonnegative, and vanishes only when $\mathbf{e} = \mathbf{0}$. For any fixed $\epsilon > 0$, compactness therefore gives

$$m_\epsilon := \min_{\mathcal{K}_e \cap \{\|\mathbf{e}\| \geq \epsilon\}} G(\mathbf{e}, \mathbf{b}) > 0. \quad (\text{A.37})$$

Consequently, whenever $\|\mathbf{e}\| \geq \epsilon$,

$$\dot{V}_e \leq -v_r^{\min} m_\epsilon < 0. \quad (\text{A.38})$$

The assumed reference bounds and the bounded closed-loop guidance signals make $\dot{\mathbf{e}}$ bounded, so $\mathbf{e}$ is uniformly continuous. If $\mathbf{e}$ did not converge to zero, there would be an $\epsilon > 0$ and infinitely many time intervals of a uniform positive duration on which $\|\mathbf{e}\| \geq \epsilon/2$. The preceding bound would then force the nonnegative function V_e to decrease by a fixed positive amount infinitely many times, a contradiction. Hence $\mathbf{e}(t) \rightarrow \mathbf{0}$. This argument supplies the uniformity missing from a conclusion based only on pointwise negativity of $\dot{V}_e$.

A.10 Dynamic interface error identity

For the kinematic–dynamic simulation interface, denote the difference between the realised and requested water-relative inertial velocities by

$$\boldsymbol{\delta} := \mathbf{v}_{\text{rel}} - u_{\text{cmd}} \mathbf{h}. \quad (\text{A.39})$$

Using $\mathbf{v}_1 = u_{\text{cmd}} \mathbf{h} + \hat{\mathbf{v}}_c$ and $\tilde{\mathbf{v}}_c = \mathbf{v}_c - \hat{\mathbf{v}}_c$, the realised inertial velocity and the target-error derivative can be written as

$$\dot{\mathbf{q}} = \mathbf{v}_1 + \tilde{\mathbf{v}}_c + \boldsymbol{\delta}, \quad \dot{\mathbf{e}} = \dot{\mathbf{r}} - \mathbf{v}_1 - \tilde{\mathbf{v}}_c - \boldsymbol{\delta}. \quad (\text{A.40})$$

This identity is used to interpret the simulations: the results report the magnitude of $\boldsymbol{\delta}$, the observer error, and any rate or actuator saturation. It is not used here to infer input-to-state stability or uniform ultimate boundedness of the dynamic plant. Such a conclusion would require a separately specified dynamic controller and a certified bound on $\boldsymbol{\delta}$.

A.11 Sufficient time-scale approximation condition

Let the right-hand side of (3.55) be denoted by $g(t)$. Since $\sigma_o > 0$,

$$\int_0^\infty C_1 e^{-\sigma_o t} dt = \frac{C_1}{\sigma_o} < \infty. \quad (\text{A.41})$$

The second term is integrable by the stated condition $\varphi \in L_1([0, \infty))$. Hence $g \in L_1$, and comparison gives $\varepsilon \in L_1$. Because the same bound also assumes $\varphi(t) \rightarrow 0$, it gives $\varepsilon(t) \rightarrow 0$ pointwise. Without that additional limit condition, integrability alone would not justify the pointwise conclusion.

A.12 Bounded derivative of the position Lyapunov function

From $V_e = \frac{1}{2} \mathbf{e}^\top \mathbf{e}$,

$$\dot{V}_e = \mathbf{e}^\top \dot{\mathbf{e}}, \quad \ddot{V}_e = \|\dot{\mathbf{e}}\|^2 + \mathbf{e}^\top. \quad (\text{A.42})$$

The error dynamics give

$$= \ddot{\mathbf{r}} - \dot{\mathbf{v}}_1 - \dot{\hat{\mathbf{v}}}_c. \quad (\text{A.43})$$

On an admissible trajectory, the reference bounds make $\ddot{\mathbf{r}}$ bounded. The observer estimate satisfies $\dot{\hat{\mathbf{v}}}_c = -\hat{\mathbf{v}}_c = k_D \tilde{\mathbf{e}}$, and the bounded observer state therefore supplies a bounded contribution. The uniform blend bound keeps κ_a bounded, while bounded policy factors and bounded realised commands keep $\mathbf{u}$, $\mathbf{v}_1$, and the direction-rate terms bounded. Together with the bounded-error condition $\|\mathbf{e}\| \leq E$ in Theorem 3.8, this makes $\dot{\mathbf{e}}$ and bounded in (A.42), which proves boundedness of $\dot{V}_e$. A continuously differentiable $\dot{V}_e$ with bounded derivative is uniformly continuous.

A.13 Modulation A and its effective speed ratio

The nominal speed equation is

$$u_{\text{nom}}^2 + 2u_{\text{nom}}(\mathbf{h}^\top \hat{\mathbf{v}}_c) + \|\hat{\mathbf{v}}_c\|^2 = \|\dot{\mathbf{r}}\|^2. \quad (\text{A.44})$$

For $u = \gamma u_{\text{nom}}$,

$$\begin{aligned} \|\mathbf{v}_{1,A}\|^2 &= \gamma^2 u_{\text{nom}}^2 + 2\gamma u_{\text{nom}}(\mathbf{h}^\top \hat{\mathbf{v}}_c) + \|\hat{\mathbf{v}}_c\|^2 \\ &= \gamma^2 \|\dot{\mathbf{r}}\|^2 + (1 - \gamma) \left[(1 + \gamma) \|\hat{\mathbf{v}}_c\|^2 + 2\gamma u_{\text{nom}}(\mathbf{h}^\top \hat{\mathbf{v}}_c) \right]. \end{aligned} \quad (\text{A.45})$$

This is (3.67)–(3.69). On the ideal alignment manifold, $\mathbf{v}_1 = \|\mathbf{v}_1\| \bar{\mathbf{a}}$ and $\mathbf{e}^\top \mathbf{v}_1 = \|\mathbf{v}_1\| \Pi_e$. Substitution into $\dot{V}_e = \mathbf{e}^\top (\dot{\mathbf{r}} - \mathbf{v}_1 - \hat{\mathbf{v}}_c)$ gives (3.70). If the stated effective-ratio inequalities hold, the same sign split as in Section A.7 gives strict decrease in Mode A; no enclosure of the complete switching trajectory follows unless the Mode B flow and the switching boundaries are also shown to be invariant. This is why Theorem 3.11 is stated only as a mode-wise decrease result.

A.14 Quantitative decay for the continuous policy

On the ideal alignment and converged-observer manifold,

$$\dot{V}_e = \|\dot{\mathbf{r}}\| \Xi, \quad \Xi = (1 - \gamma p)s - \gamma(1 - p) \|\mathbf{e}\|. \quad (\text{A.46})$$

If $s > 0$ and $\gamma p \leq 1$, use $s \leq \|\mathbf{e}\|$ to obtain $\Xi \leq (1 - \gamma) \|\mathbf{e}\| = -(\gamma - 1) \|\mathbf{e}\|$. In all other cases the first term is nonpositive. Indeed, if $s < 0$, then $\gamma(s) < 1$ and $p \leq 1$, so $(1 - \gamma p)s \leq 0$; if $s > 0$ and $\gamma p > 1$, the factor $1 - \gamma p$ is negative. Therefore,

$$\Xi \leq -\gamma(1 - p) \|\mathbf{e}\| \leq -\gamma \frac{\|\mathbf{e}\|^2}{\lambda + \|\mathbf{e}\|}. \quad (\text{A.47})$$

Multiplication by the positive reference speed gives (3.81)–(3.82).

A.15 Layered kinematic convergence proof

This section proves Theorem 3.18 for the reduced ideal kinematic model. The proof keeps the observer, alignment, and position channels separate until their coupling terms have been identified. In particular, it does not require a uniform linear-in- $\|\mathbf{e}\|$ decay estimate near the origin.

Observer and alignment estimates

Let

$$\mathbf{z}_o = [\tilde{\mathbf{e}}^\top \quad \tilde{\mathbf{v}}_c^\top]^\top, \quad V_o = \mathbf{z}_o^\top P \mathbf{z}_o, \quad (\text{A.48})$$

where $P = P^\top \succ 0$ satisfies

$$A_o^\top P + P A_o = -Q, \quad Q = Q^\top \succ 0. \quad (\text{A.49})$$

By Theorem 3.2, there are constants $C_o > 0$ and $\sigma_o > 0$ such that

$$\|\mathbf{z}_o(t)\| \leq C_o e^{-\sigma_o t} \|\mathbf{z}_o(0)\|, \quad \|\tilde{\mathbf{v}}_c(t)\| \leq C_o e^{-\sigma_o t} \|\mathbf{z}_o(0)\|. \quad (\text{A.50})$$

Under the exact ideal direction dynamics and $V_a(0) < 1$, Theorem 3.7 gives

$$V_a(t) = \frac{V_a(0)}{V_a(0) + [1 - V_a(0)]e^{2kt}} \leq C_a e^{-2kt}, \quad C_a := \frac{V_a(0)}{1 - V_a(0)}. \quad (\text{A.51})$$

For two unit vectors, $\|\tilde{\mathbf{v}}_1 - \bar{\mathbf{a}}\|^2 = 2(1 - \tilde{\mathbf{v}}_1^\top \bar{\mathbf{a}}) = 4V_a$; hence

$$\|\tilde{\mathbf{v}}_1(t) - \bar{\mathbf{a}}(t)\| \leq 2\sqrt{C_a} e^{-kt}. \quad (\text{A.52})$$

Position dynamics with both vanishing inputs retained

For Modulation B, $\|\mathbf{v}_1\| = \gamma(s) \|\dot{\mathbf{r}}\|$ and $\gamma \in [\gamma_-, \gamma_+]$. Without imposing exact alignment in the position equation,

$$\begin{aligned}\dot{V}_e &= \mathbf{e}^\top (\dot{\mathbf{r}} - \mathbf{v}_1 - \tilde{\mathbf{v}}_c) \\ &= \|\dot{\mathbf{r}}\| s - \gamma \|\dot{\mathbf{r}}\| \mathbf{e}^\top \tilde{\mathbf{v}}_1 - \mathbf{e}^\top \tilde{\mathbf{v}}_c \\ &= -\|\dot{\mathbf{r}}\| G(\mathbf{e}, \mathbf{b}) - \mathbf{e}^\top \mathbf{d}(t),\end{aligned}\tag{A.53}$$

where

$$G(\mathbf{e}, \mathbf{b}) := \gamma(s)\Pi_e - s, \quad s = \mathbf{e}^\top \mathbf{b}, \quad \Pi_e = \mathbf{e}^\top \bar{\mathbf{a}},\tag{A.54}$$

$$\mathbf{d}(t) := \gamma(s) \|\dot{\mathbf{r}}\| [\tilde{\mathbf{v}}_1 - \bar{\mathbf{a}}] + \tilde{\mathbf{v}}_c.\tag{A.55}$$

The first term in $\mathbf{d}$ is the effect of imperfect direction alignment on the position channel; the second is the current-estimation error. From Theorem 3.1, (A.50), and (A.52),

$$\|\mathbf{d}(t)\| \leq 2\gamma_+ v_r^{\max} \sqrt{C_a} e^{-kt} + C_o \|\mathbf{z}_o(0)\| e^{-\sigma_o t} =: \beta(t),\tag{A.56}$$

where $\beta(t) \rightarrow 0$ and $\beta \in L_1([0, \infty))$.

By Theorem 3.14, the continuous policy and blend regularity imply

$$G(\mathbf{e}, \mathbf{b}) > 0 \quad \text{for every } \mathbf{e} \neq \mathbf{0}, \quad G(\mathbf{0}, \mathbf{b}) = 0.\tag{A.57}$$

Equation (A.53) therefore gives

$$\dot{V}_e \leq -v_r^{\min} G(\mathbf{e}, \mathbf{b}) + \beta(t) \|\mathbf{e}\|.\tag{A.58}$$

Boundedness and convergence of the position error

Let $r(t) := \|\mathbf{e}(t)\|$. For $r > 0$, since $\dot{V}_e = r\dot{r}$, dropping the non-positive nominal term in (A.58) yields

$$\dot{r}(t) \leq \beta(t).\tag{A.59}$$

The same inequality holds in the upper-Dini-derivative sense at $r = 0$. Consequently,

$$r(t) \leq r(0) + \int_0^t \beta(\tau) d\tau \leq r(0) + \|\beta\|_{L_1} =: E < \infty.\tag{A.60}$$

Thus the position error is bounded without assuming a bounded sublevel set in advance.

Fix any $\varepsilon \in (0, E]$ and define

$$\mathcal{K}_\varepsilon := \{(\mathbf{e}, \mathbf{b}) : \varepsilon \leq \|\mathbf{e}\| \leq E, \quad \|\mathbf{b}\| = 1, \quad \|\mathbf{a}(\mathbf{e}, \mathbf{b})\| \geq a_{\min}\}.\tag{A.61}$$

This set is compact. The function G is continuous on it and is strictly positive by (A.57). If

$\mathcal{K}_\varepsilon$ is empty, the trajectory is already inside the ε -ball. Otherwise, compactness gives

$$m_\varepsilon := \min_{\mathcal{K}_\varepsilon} G(\mathbf{e}, \mathbf{b}) > 0. \quad (\text{A.62})$$

Because $\beta(t) \rightarrow 0$, there is a finite T_ε such that

$$\beta(t)E \leq \frac{1}{2}v_r^{\min}m_\varepsilon, \quad t \geq T_\varepsilon. \quad (\text{A.63})$$

Whenever $t \geq T_\varepsilon$ and $\|\mathbf{e}(t)\| \geq \varepsilon$, equation (A.58) then gives

$$\dot{V}_e(t) \leq -\frac{1}{2}v_r^{\min}m_\varepsilon < 0. \quad (\text{A.64})$$

If the trajectory never entered the ball $\|\mathbf{e}\| < \varepsilon$, the non-negative function V_e would decrease at the fixed rate in (A.64), which is impossible. After entry, the same strict inward sign on the boundary $\|\mathbf{e}\| = \varepsilon$ prevents a later exit. Hence, for every $\varepsilon > 0$, the trajectory enters and remains in that ball after a finite time. This proves

$$\lim_{t \rightarrow \infty} \mathbf{e}(t) = \mathbf{0}. \quad (\text{A.65})$$

Together with (A.50) and (A.51), this establishes all three limits in (3.85).

Remark A.1 (What is, and is not, included). The proof includes the transient current-observer error and the transient difference $\bar{\mathbf{v}}_1 - \bar{\mathbf{a}}$ generated by the ideal alignment dynamics. It does not include the projection term or the time-scale approximation term of (3.51). Once the projected physical angular command is used, the logistic estimate (A.51) is no longer automatic and a separate perturbation analysis is required. The proof also does not include sampling, rate limits, actuator saturation, or the REMUS 100 dynamics. Those effects are evaluated numerically and are not part of Theorem 3.18.

Remark A.2 (Hysteretic surge policy). The proof above uses the strict nominal position decrease of the continuous policy. For the hysteretic policy, Theorem 3.12 provides mode-wise sign statements and an instantaneous zero-derivative set, but not a complete hybrid invariance proof for a curved time-varying tangent. Therefore the asymptotic conclusion of Theorem 3.18 is not claimed for the hysteretic policy.

A.16 Unit-direction integration

Forward Euler integration of (3.3) does not preserve $\|\mathbf{h}\| = 1$ exactly. A practical first-order update is

$$\mathbf{h}_{k+1}^{\text{raw}} = \mathbf{h}_k + \Delta t(\boldsymbol{\omega}_k \times \mathbf{h}_k), \quad \mathbf{h}_{k+1} = \frac{\mathbf{h}_{k+1}^{\text{raw}}}{\|\mathbf{h}_{k+1}^{\text{raw}}\|}. \quad (\text{A.66})$$

A Rodrigues or matrix-exponential update preserves the sphere geometrically and is preferred when the angular increment is not small.

Appendix B

Observer and Current-Frame Derivations

B.1 Exponential convergence of the constant-current observer

The kinematic current observer can be derived either from the vehicle position dynamics or from the target-relative error dynamics. This appendix gives both constructions explicitly and then specialises to the error-dynamics observer used in the thesis.

Position-output formulation

Starting from

$$\dot{\mathbf{q}} = \mathbf{u}\mathbf{h} + \mathbf{v}_c, \quad \dot{\mathbf{v}}_c = \mathbf{0}, \quad y_q = \mathbf{q}, \quad (\text{B.1})$$

a Luenberger observer with estimated position $\hat{\mathbf{q}}$ and estimated current $\hat{\mathbf{v}}_c$ can be written as

$$\dot{\hat{\mathbf{q}}} = \mathbf{u}\mathbf{h} + \hat{\mathbf{v}}_c + k_q(\mathbf{q} - \hat{\mathbf{q}}), \quad (\text{B.2})$$

$$\dot{\hat{\mathbf{v}}}_c = k_v(\mathbf{q} - \hat{\mathbf{q}}), \quad k_q, k_v > 0, \quad (\text{B.3})$$

where scalar gains are used so that the three Cartesian channels are identical. Define the position and current estimation errors

$$\tilde{\mathbf{q}} := \mathbf{q} - \hat{\mathbf{q}}, \quad \tilde{\mathbf{v}}_c := \mathbf{v}_c - \hat{\mathbf{v}}_c. \quad (\text{B.4})$$

The first error equation is obtained by direct subtraction:

$$\begin{aligned} \dot{\tilde{\mathbf{q}}} &= \dot{\mathbf{q}} - \dot{\hat{\mathbf{q}}} \\ &= (\mathbf{u}\mathbf{h} + \mathbf{v}_c) - (\mathbf{u}\mathbf{h} + \hat{\mathbf{v}}_c + k_q(\mathbf{q} - \hat{\mathbf{q}})) \\ &= \mathbf{v}_c - \hat{\mathbf{v}}_c - k_q\tilde{\mathbf{q}} \\ &= \tilde{\mathbf{v}}_c - k_q\tilde{\mathbf{q}}. \end{aligned} \quad (\text{B.5})$$

The current-error equation is

$$\begin{aligned} \dot{\tilde{\mathbf{v}}}_c &= \dot{\mathbf{v}}_c - \dot{\hat{\mathbf{v}}}_c \\ &= \mathbf{0} - k_v(\mathbf{q} - \hat{\mathbf{q}}) \\ &= -k_v\tilde{\mathbf{q}}. \end{aligned} \quad (\text{B.6})$$

Thus the position-output observer has error dynamics

$$\frac{d}{dt} \begin{bmatrix} \tilde{\mathbf{q}} \\ \tilde{\mathbf{v}}_c \end{bmatrix} = \begin{bmatrix} -k_q \mathbf{I} & \mathbf{I} \\ -k_v \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{q}} \\ \tilde{\mathbf{v}}_c \end{bmatrix}. \quad (\text{B.7})$$

Target-error-output formulation

The target-relative error is $\mathbf{e} = \mathbf{r} - \mathbf{q}$. Because the reference target is generated by the guidance system and $\mathbf{q}$ is the navigation measurement, $\mathbf{e}$ can be used as a computed output. Differentiating $\mathbf{e} = \mathbf{r} - \mathbf{q}$ and using (3.2) gives

$$\dot{\mathbf{e}} = \dot{\mathbf{r}} - \dot{\mathbf{q}} = \dot{\mathbf{r}} - u\mathbf{h} - \mathbf{v}_c, \quad \dot{\mathbf{v}}_c = \mathbf{0}, \quad y_e = \mathbf{e}. \quad (\text{B.8})$$

The thesis uses the following Luenberger observer:

$$\dot{\hat{\mathbf{e}}} = \dot{\mathbf{r}} - u\mathbf{h} - \hat{\mathbf{v}}_c + k_C(\mathbf{e} - \hat{\mathbf{e}}), \quad (\text{B.9})$$

$$\dot{\hat{\mathbf{v}}}_c = -k_D(\mathbf{e} - \hat{\mathbf{e}}), \quad k_C, k_D > 0. \quad (\text{B.10})$$

Define

$$\tilde{\mathbf{e}} := \mathbf{e} - \hat{\mathbf{e}}, \quad \tilde{\mathbf{v}}_c := \mathbf{v}_c - \hat{\mathbf{v}}_c. \quad (\text{B.11})$$

The target-error observer error satisfies

$$\begin{aligned} \dot{\tilde{\mathbf{e}}} &= \dot{\mathbf{e}} - \dot{\hat{\mathbf{e}}} \\ &= (\dot{\mathbf{r}} - u\mathbf{h} - \mathbf{v}_c) - (\dot{\mathbf{r}} - u\mathbf{h} - \hat{\mathbf{v}}_c + k_C(\mathbf{e} - \hat{\mathbf{e}})) \\ &= -\mathbf{v}_c + \hat{\mathbf{v}}_c - k_C\tilde{\mathbf{e}} \\ &= -\tilde{\mathbf{v}}_c - k_C\tilde{\mathbf{e}}. \end{aligned} \quad (\text{B.12})$$

The current-estimation error satisfies

$$\begin{aligned} \dot{\tilde{\mathbf{v}}}_c &= \dot{\mathbf{v}}_c - \dot{\hat{\mathbf{v}}}_c \\ &= \mathbf{0} - (-k_D(\mathbf{e} - \hat{\mathbf{e}})) \\ &= k_D\tilde{\mathbf{e}}. \end{aligned} \quad (\text{B.13})$$

Combining (B.12) and (B.13) gives the observer-error dynamics used in the main text:

$$\frac{d}{dt} \begin{bmatrix} \tilde{\mathbf{e}} \\ \tilde{\mathbf{v}}_c \end{bmatrix} = \underbrace{\begin{bmatrix} -k_C \mathbf{I} & -\mathbf{I} \\ k_D \mathbf{I} & \mathbf{0} \end{bmatrix}}_{A_o} \begin{bmatrix} \tilde{\mathbf{e}} \\ \tilde{\mathbf{v}}_c \end{bmatrix}. \quad (\text{B.14})$$

The three Cartesian modes of both (B.7) and (B.14) are identical. For the target-error observer, each scalar channel satisfies

$$\dot{\tilde{e}}_i = -k_C \tilde{e}_i - \tilde{v}_{c,i}, \quad \dot{\tilde{v}}_{c,i} = k_D \tilde{e}_i. \quad (\text{B.15})$$

Eliminating $\tilde{v}_{c,i}$ from the first equation gives

$$\begin{aligned}\ddot{\tilde{e}}_i &= -k_C \dot{\tilde{e}}_i - \dot{\tilde{v}}_{c,i} \\ &= -k_C \dot{\tilde{e}}_i - k_D \tilde{e}_i,\end{aligned}\tag{B.16}$$

or

$$\ddot{\tilde{e}}_i + k_C \dot{\tilde{e}}_i + k_D \tilde{e}_i = 0.\tag{B.17}$$

The characteristic polynomial is therefore

$$\lambda_o^2 + k_C \lambda_o + k_D.\tag{B.18}$$

For $k_C > 0$ and $k_D > 0$, this polynomial is Hurwitz. Consequently, A_o is Hurwitz and the origin of (B.14) is globally exponentially stable. Equivalently, for every symmetric $Q \succ 0$, there is a unique $P \succ 0$ satisfying $A_o^T P + P A_o = -Q$, which gives a quadratic exponential bound. Critical damping is obtained when $k_C^2 = 4k_D$.

B.2 Mismatch created by a yaw-only current transformation

Suppose that a constant current intended in the NED frame is transformed to the body frame using yaw only,

$$\mathbf{v}_{c,\psi}^b = \mathbf{R}_z(\psi)^T \mathbf{v}_c,\tag{B.19}$$

while the six-degree-of-freedom kinematics reconstruct inertial velocity with the full attitude matrix $\mathbf{R}_{zyx}$. The current actually injected in the inertial kinematics is then

$$\bar{\mathbf{v}}_c(t) = \mathbf{R}_{zyx}(t) \mathbf{v}_{c,\psi}^b(t).\tag{B.20}$$

Unless roll and pitch vanish, $\bar{\mathbf{v}}_c$ differs from the intended constant vector $\mathbf{v}_c$. Its time derivative is

$$\dot{\bar{\mathbf{v}}}_c = \mathbf{R}_{zyx} \mathbf{S}(\omega^b) \mathbf{v}_{c,\psi}^b + \mathbf{R}_{zyx} \dot{\mathbf{v}}_{c,\psi}^b,\tag{B.21}$$

where $\mathbf{S}(\cdot)$ is the skew-symmetric cross-product matrix. Thus, even though the parameters used to define the horizontal current are constant, the effective inertial current seen by the observer is generally time varying.

The relevant error is therefore not $\mathbf{v}_c - \hat{\mathbf{v}}_c$ but

$$\tilde{\mathbf{v}}_{c,\text{eff}} := \bar{\mathbf{v}}_c(t) - \hat{\mathbf{v}}_c.\tag{B.22}$$

With $\mathbf{z}_{\text{eff}} = [\tilde{\mathbf{e}}^T, \tilde{\mathbf{v}}_{c,\text{eff}}^T]^T$, direct differentiation gives the forced system

$$\dot{\mathbf{z}}_{\text{eff}} = \mathbf{A}_o \mathbf{z}_{\text{eff}} + \begin{bmatrix} \mathbf{0} \\ \dot{\bar{\mathbf{v}}}_c \end{bmatrix}.\tag{B.23}$$

The nominal observer is therefore a stable filter driven by a modelling-induced input. For a harmonic component of frequency Ω , the scalar transfer magnitude from the effective-current derivative to the current-estimation error is determined by $(s + k_C)/(s^2 + k_C s + k_D)$; any

simple low-frequency ripple estimate based on this expression is a scaling argument, not a certified nonlinear bound for an arbitrary manoeuvre.

B.3 Full-rotation correction and recovery of the nominal observer

For a current that is constant in the NED frame, the frame-consistent body representation is

$$\mathbf{v}_c^b = \mathbf{R}_{zyx}^\top \mathbf{v}_c. \quad (\text{B.24})$$

Since $\dot{\mathbf{R}}_{zyx} = \mathbf{R}_{zyx} \mathbf{S}(\boldsymbol{\omega}^b)$,

$$\dot{\mathbf{v}}_c^b = -\mathbf{S}(\boldsymbol{\omega}^b) \mathbf{v}_c^b. \quad (\text{B.25})$$

Equations (B.24)–(B.25) imply $\mathbf{R}_{zyx} \mathbf{v}_c^b = \mathbf{v}_c$ for every nonsingular attitude representation. Hence $\bar{\mathbf{v}}_c = \mathbf{v}_c$, $\dot{\bar{\mathbf{v}}}_c = \mathbf{0}$, and the forcing term in (B.23) vanishes. The observer-error dynamics reduce exactly to the homogeneous exponentially stable system (B.14).

Appendix C

Supplementary Numerical Results

This appendix collects numerical evidence that complements the main figures and tables of Chapter 5. It contains extended numerical consistency checks, path-wise comparison plots, the ideal-plant benchmark, and additional dynamic-response histories.

C.1 Extended numerical consistency checks

The main sign and regularity checks are reported once in Figure 5.4. The additional dissipation trace for the REMUS/PI loop is not reproduced here because its candidate inequality is violated through most of the tested window; it therefore cannot serve as a stability proof. The remaining alignment-storage history below is retained as a supplementary numerical check and does not replace the assumptions of the conditional cascade routes.

C.2 Five-method comparison on all reference paths

Figures C.2 to C.5 show the geometric path-distance histories for the five transferred implementations. The path-following implementation is an exploratory comparator and is not covered by the thesis stability result. The oceanscan panel is accompanied by the terrain and clearance histories because path distance alone does not establish terrain-following performance.

C.3 Ideal-plant comparison

The same five outer-guidance implementations were applied to the ideal kinematic plant. Tables C.1 and C.2 separate RMS distance from settling time to preserve readability. The comparison still includes method-specific adapters and progression conventions; it therefore supports plant-level attribution only qualitatively.

Removing the REMUS plant reduces the reported RMS of the thesis law on the four curved or survey references from 0.095–0.431 m to 0.002–0.011 m. This difference is consistent with dynamic command realisation being an important source of the remaining error; it is not a separate causal proof, and the smallest kinematic distances are below the declared path-projection resolution. The path-following extension also gives small geometric distances on the ideal plant, but its autonomous progression and coverage remain separate issues; it is therefore retained as supplementary evidence rather than as a formal contribution of the thesis.




Figure C.1: Alignment storage for the four analysed internal configurations and the exploratory progression variant on the kinematic helix experiment.




Figure C.2: Five-method geometric path distance on the inclined line.

Table C.1: Five-method RMS geometric path distance [m] on the ideal kinematic plant. **Bold** = best in column; *italic* = best of the thesis-lineage implementations when the overall best is a non-thesis-lineage method.

Method	Helix	Line	Lawnmower	Lemniscate3D	Oceanscan
M-B continuous-PP	0.005	<i>0.372</i>	0.011	0.002	0.011
MED switching predecessor	3.764	0.226	2.523	1.787	2.518
M-B continuous-PF	0.038	0.624	0.013	0.006	0.013
Coates 3-D ALOS	–	–	0.242	0.035	0.025
Degorre 3-D VRP	0.071	0.025	0.040	0.063	0.039

Note. An em-dash denotes a non-finite reported RMS value under the common geometric-distance estimator.




Figure C.3: Five-method geometric path distance on the lawnmower.

Table C.2: Five-method settling time T_s [s] on the ideal kinematic plant. **Bold** = best (smallest) finite value in column; *italic* = best of the thesis-lineage implementations when the overall best is a non-thesis-lineage method.

Method	Helix	Line	Lawnmower	Lemniscate3D	Oceanscan
M-B continuous-PP	<i>19.5</i>	40.7	<i>143.2</i>	<i>124.3</i>	<i>144.9</i>
MED switching predecessor	∞	37.0	139.9	125.2	141.5
M-B continuous-PF	21.1	45.2	167.3	134.8	168.3
Coates 3-D ALOS	∞	∞	167.9	156.3	168.6
Degorre 3-D VRP	7.8	19.9	717.3	94.2	739.1




Figure C.4: Five-method geometric path distance on the spatial Geronio lemniscate.

Table C.3: M-B continuous-PP across plant levels.

Path	Ideal kinematic plant		REMUS 100/PI plant	
	$d_{\mathcal{P}}^{\text{RMS}}$ [m]	T_s [s]	$d_{\mathcal{P}}^{\text{RMS}}$ [m]	T_s [s]
Helix	0.005	19.5	0.431	27.8
Line	0.372	40.7	0.970	48.5
Lawnmower	0.011	143.2	0.095	144.8
Lemniscate3D	0.002	124.3	0.118	125.3
Oceanscan	0.011	144.9	0.099	146.9

Note. The table isolates the selected thesis formulation. It shows the degradation introduced by dynamic command realisation, but it does not by itself prove a cascade-stability property for the REMUS 100/PI implementation.




Figure C.5: Five-method geometric path distance on the terrain-following survey.




Figure C.6: Altitude, seabed cross-section, and vehicle-depth histories for the four analysed internal configurations and the exploratory progression variant on the oceanscan reference.

C.4 Extended dynamic response checks

The following figures complement the actuator and surge plots in the main chapter by exposing attitude, body velocity, angular rates, and cumulative command activity under the five-method helix comparison.




Figure C.7: Attitude histories on the helix.




Figure C.8: Body-frame velocity histories on the helix.




Figure C.9: Body-frame angular-rate histories on the helix.




Figure C.10: Cumulative angular effort and surge-command roughness on the helix.

Appendix D

Simulation Setup Details

This appendix collects the material required to reproduce the runs of Chapter 5: the parameterisation of each reference path (Section D.1), the guidance and inner-loop parameters (Section D.2), the initial-condition composition (Section D.3), the runtime feasibility protocol (Section D.5), the REMUS 100 plant configuration (Section D.6), the ocean-current model (Section D.7), the numerical integration settings (Section D.8), the paired setup of the observer frame-consistency experiment (Section D.9), and the external comparison protocol (Section D.10). The values below define the numerical experiments reported in this thesis.

D.1 Reference-path parameterisations

The five reference geometries listed in Table 5.1 are generated by the following analytic expressions.

Helix (descending).

$$\mathbf{r}_H(t) = \begin{bmatrix} R_c \cos(\omega t) \\ R_c \sin(\omega t) \\ v_z t \end{bmatrix}, \quad R_c = 50 \text{ m}, \quad \omega = 0.025 \text{ rad s}^{-1}, \quad v_z = 0.12 \text{ m s}^{-1}. \quad (\text{D.1})$$

The reference speed is constant, $\|\dot{\mathbf{r}}_H\| = \sqrt{(R_c\omega)^2 + v_z^2} \approx 1.256 \text{ m s}^{-1}$, and the centripetal acceleration is $\|\ddot{\mathbf{r}}_H\| = R_c\omega^2 = 0.03125 \text{ m s}^{-2}$. The horizontal period is 251.3 s. The ideal own-method run uses $T = 300 \text{ s}$, the dynamic own-method run uses $T = 600 \text{ s}$, and the cross-method comparison uses $T = 400 \text{ s}$.

Straight inclined line.

$$\mathbf{r}_L(t) = t \begin{bmatrix} v_N \\ v_E \\ v_D \end{bmatrix}, \quad (v_N, v_E, v_D) = (1.0 \text{ m s}^{-1}, 0.5 \text{ m s}^{-1}, 0.3333 \text{ m s}^{-1}), \quad (\text{D.2})$$

so that $\|\dot{\mathbf{r}}_L\| \approx 1.166 \text{ m s}^{-1}$ and $\|\ddot{\mathbf{r}}_L\| = 0$. Nominal horizon $T = 300 \text{ s}$ (path length $\approx 350 \text{ m}$).

Lawnmower (5-leg boustrophedon). Five straight legs of length $L = 2400 \text{ m}$ at constant depth $z_0 = 20 \text{ m}$, joined by four semicircular U-turns of radius $R = 200 \text{ m}$, traversed at constant path speed $V = 2.0 \text{ m s}^{-1}$. The path is C^1 (tangent-continuous) but the curvature

steps at each leg/turn junction. Nominal horizon $T = (n_{\text{legs}}L + (n_{\text{legs}} - 1)\pi R)/V \approx 7256$ s, capped at 8000 s in the comparison runs.

Three-dimensional lemniscate (Gerono).

$$\mathbf{r}_{\text{Lem}}(\tau) = \begin{bmatrix} A \cos \tau \\ A \sin \tau \cos \tau \\ z_0 + A_z \sin \tau \end{bmatrix}, \quad A = 500 \text{ m}, \quad A_z = 30 \text{ m}, \quad z_0 = 60 \text{ m}, \quad (\text{D.3})$$

where $\tau(s)$ is precomputed from a persistent arc-length table so that $\|\dot{\mathbf{r}}_{\text{Lem}}\| = V = 1.5 \text{ m s}^{-1}$ (constant speed by construction). Full-period horizon $T \approx 2035$ s (closed loop).

Terrain-following (oceanscan). The horizontal pattern coincides with the lawnmower; the depth follows the smoothed seabed at a fixed altitude:

$$z_{\text{ref}}(t) = z_{\text{bed}}(x_{\text{ref}}(t), y_{\text{ref}}(t)) - h_{\text{alt}}, \quad h_{\text{alt}} = 25 \text{ m}. \quad (\text{D.4})$$

The seabed model is a three-term multi-octave sum of oriented sinusoids (Table D.1) with mean depth $z_c = 150$ m; the maximum slope is $\approx 18^\circ$, keeping the required depth rate $|\dot{z}| \leq V \tan(18^\circ) \approx 0.6 \text{ m s}^{-1}$ within the pitch authority of the REMUS. Analytic first and second partial derivatives of $z_{\text{bed}}(x, y)$ are used to construct $\ddot{\mathbf{r}}$ on the oceanscan reference. Nominal horizon $T \approx 7256$ s.

Table D.1: Seabed model $z_{\text{bed}}(x, y) = z_c + \sum_i A_i \sin(k_{x,i}x + k_{y,i}y + \phi_i)$ used for the oceanscan reference.

i	A_i [m]	$k_{x,i}$ [1/m]	$k_{y,i}$ [1/m]	ϕ_i [rad]	λ_i [m]
1	65.0	0.0016	-0.0011	0.00	3240
2	40.0	-0.0026	0.0020	1.70	1915
3	10.0	0.0045	0.0028	3.10	1190

D.2 Guidance and inner-loop parameters

The blended-alignment guidance and observer gains are collected in Table D.2, the switching-predecessor parameters in Table D.3, the inner-loop PI cascade in Table D.4, and the frozen external-baseline gains in Table D.5.

The exploratory progression implementation retains the M-B continuous direction, observer, and surge parameters of Table D.2, but replaces the prescribed reference schedule by a state-dependent particle update. It is included only as an exploratory implementation; its regularity, progression-domain, and integration assumptions are not part of the proved formulation and are listed in Section 6.2.

Table D.2: Blended-alignment guidance law and observer gains.

Symbol	Value	Description
<i>Blended-alignment guidance</i>		
λ	5.0 m	Blend length scale
k	1.0	Alignment gain (kinematic sims)
k_{dyn}	0.3	Alignment gain (dynamic cascade)
<i>Luenberger observer, char. polynomial $s^2 + k_C s + k_D$</i>		
k_C	1.0	Observer position feedback
k_D	0.25	$= k_C^2/4$ (critically damped)
obs. velocity aiding	enabled	DVL water-track aiding
full-rotation current	enabled	Frame-consistent map (D.11)
<i>Hysteretic surge policy, two-sphere</i>		
ρ_{in}	2.0 m	Inner hysteresis sphere
ρ_{out}	5.0 m	Outer hysteresis sphere
γ_+	1.5	Acceleration factor ($s \geq 0$)
γ_-	0.5	Deceleration factor ($s < 0$)
γ_{hold}	1.0	Cruise inside inner sphere
<i>Smooth surge policy, $\gamma(s) = 1 + \eta \tanh(s/\alpha_0)$</i>		
η	0.4	Amplitude, $\gamma \in [0.6, 1.4]$
α_0	3.0 m	Characteristic transition scale

Table D.3: Simulation parameters of the switching MED predecessor. The second column gives the values used in the present transfer study; these are not presented as a reproduction with the conference-paper tuning.

Parameter	Simulation value	Role
k_1, k_2	1.5, 1.5	Approach- and tangent-alignment gains
k_C, k_D	1.0, 0.5	MED observer gains
α, β	1.5, 0.7	Surge increase and overshoot-recovery factors
$\rho_{\text{in}}, \rho_{\text{out}}$	1.5 m, 6.0 m	Task-switch hysteresis band
$u_{\text{default}}, u_{\text{min}}$	$1.0 \text{ m s}^{-1}, 0.1 \text{ m s}^{-1}$	Fallback and surge floor
ω_{max}	5.0 rad s^{-1}	Outer angular-command cap

Table D.4: Inner-loop PI cascade gains and actuator limits.

Symbol	Value	Loop
$K_{p,q}, K_{i,q}$	6.0, 1.5	Pitch-rate $\rightarrow$ stern-plane δ_s
$K_{p,r}, K_{i,r}$	6.0, 1.5	Yaw-rate $\rightarrow$ rudder δ_r
$K_{p,n}, K_{i,n}$	200, 40	Surge $\rightarrow$ propeller n
n_{ff}	610 RPM/(m/s)	Propeller feed-forward
Rate clamp $\dot{\theta}_{\text{max}}$	15° s^{-1} (0.262 rad/s)	Applied to $q_{\text{cmd}}, r_{\text{cmd}}$
Fin limit $ \delta_r , \delta_s $	20° (0.349 rad)	Plant fin saturation
Propeller limit	[0, 1525 rpm]	Plant propeller saturation
Anti-windup	integrator rollback on saturation	Applied in the PI interface

Table D.5: Published outer-guidance gains retained for the two external baselines.

Symbol	Value	Description
<i>Coates & Fossen (2026) spherical 3-D ALOS</i>		
Δ_h	20 m	Horizontal look-ahead
Δ_v	20 m	Vertical look-ahead
k_h	0.001	Horizontal adaptive gain
k_v	0.001	Vertical adaptive gain
M_θ	20° (0.349 rad)	Projection bound on crab estimates
R_{switch}	5 m	Waypoint arrival radius
u_{ref}	$\ \dot{\mathbf{r}}(t)\ $	Overridden by the guidance harness
Attitude PI (feeds the shared command interface)	$K_{p,\psi} = K_{p,\theta} = 2.5;$ $K_{i,\psi} = K_{i,\theta} = 0.2$	Converts (ψ_d, θ_d) to $(q_{\text{cmd}}, r_{\text{cmd}})$
<i>Degorre et al. (2024) 3-D VRP kinematic guidance</i>		
ϵ	[0.3, 0, 0] m	Tracking-point offset (VRP) ahead of vehicle
K_p	0.4	PI proportional gain
K_i	0.05	PI integral gain
u_{ref}	$\ \dot{\mathbf{r}}(t)\ $	Overridden by the guidance harness

D.3 Initial conditions

Initial conditions are composed by

$$\mathbf{q}(0) = \mathbf{r}(0) - P_{\text{behind}} \mathbf{b}(0) + P_{\text{lateral}} \mathbf{n}_{\text{lat}} + [0, 0, -P_{\text{up}}]^\top, \quad (\text{D.5})$$

where $\mathbf{b}(0) = \dot{\mathbf{r}}(0)/\|\dot{\mathbf{r}}(0)\|$ is the path tangent at $t = 0$ and $\mathbf{n}_{\text{lat}} = [-\mathbf{b}_y(0), \mathbf{b}_x(0), 0]^\top / \|\cdot\|$ is the horizontal normal. Table D.6 lists the per-path values used in the comparative runs.

Table D.6: Per-path initial-condition composition for the comparative runs.

Path	P_{behind} [m]	P_{lateral} [m]	P_{up} [m]	$\ \mathbf{e}_0\ $ [m]
Helix	10	-24	12	≈ 29
Line	5	-50	10	≈ 51
Lemniscate3D	30	200	-100	≈ 224
Lawnmower	30	24	12	≈ 40
Oceanscan	30	24	12	≈ 40

The nominal stress-test values used by the long own-method runs are $P_{\text{behind}} = 100$ m, $P_{\text{lateral}} = 100$ m, $P_{\text{up}} = 50$ m, which give $\|\mathbf{e}_0\| \approx 150$ m and define the large-offset cases. The dynamic helix own-method run is the exception: it uses the helix comparison offset [10, -24, 12] m so that its five internal traces start from the same $\|\mathbf{e}_0\| \approx 29.03$ m condition as the helix cross-method benchmark.

D.4 Performance-metric definitions

The primary geometric error used in the benchmark is the shortest distance from the vehicle to the path,

$$d_{\mathcal{P}}(\mathbf{q}) := \inf_{\sigma \in \mathcal{I}} \|\mathbf{q} - \mathbf{r}_{\mathcal{P}}(\sigma)\|. \quad (\text{D.6})$$

The reported implementation approximates this infimum on an extended sampled reference, using a search half-window $\min(0.25T, 120 \text{ s})$ about the scheduled reference index. This reduces the phase-lag bias of a narrow local search but is not an exact continuous-curve projection. On the survey paths the reference-point spacing is approximately 0.2 m; values below that scale are therefore reported as implementation outputs but are not used to infer a sub-decimetres ranking. Steady-window statistics are evaluated on the interval $\mathcal{T}_s = [T_r, T]$ with $T_r = 0.15T$, unless otherwise stated:

$$d_{\mathcal{P}}^{\text{RMS}} = \left(\frac{1}{|\mathcal{T}_s|} \int_{\mathcal{T}_s} d_{\mathcal{P}}(\mathbf{q}(t))^2 dt \right)^{1/2}, \quad d_{\mathcal{P}}^{\text{max}} = \max_{t \in \mathcal{T}_s} d_{\mathcal{P}}(\mathbf{q}(t)), \quad (\text{D.7})$$

and the terminal along-track error is $s_e^{\text{fin}} := (\mathbf{r}(T) - \mathbf{q}(T))^{\top} \mathbf{b}(T)$. The settling time T_s is the first instant after which $d_{\mathcal{P}}(t) \leq \max(0.10 d_{\mathcal{P}}(0), 0.05 \text{ m})$ for the remainder of $[0, T]$; if no such instant exists, $T_s = \infty$ is reported. Integrated efforts are

$$J_{\omega} = \int_{\mathcal{T}_s} \|\omega_{\text{cmd}}(t)\|^2 dt, \quad J_{\Delta u} = \int_{\mathcal{T}_s} \dot{u}_{\text{cmd}}(t)^2 dt, \quad (\text{D.8})$$

the actuator-saturation duty is $\text{sat}\% = |\mathcal{T}_s|^{-1} \int_{\mathcal{T}_s} \chi_{\text{sat}}(t) dt$ with $\chi_{\text{sat}} = 1$ whenever any actuator or rate clamp is active, the fallback duty is $f_{\text{fb}} = T^{-1} \int_0^T \chi_{\text{fb}}(t) dt$ (evaluated on the full horizon, see Section D.5), the current observer quantity in the benchmark tables is $\|\tilde{\mathbf{v}}_c(T_r)\|$ (it does not exist for the two external baselines), and the fitted exponential convergence rate α_{conv} is estimated on the convergence window $[0, T_r]$ by a least-squares fit of $\log d_{\mathcal{P}}(t)$ against t . The time-scale approximation measure $J_{\varepsilon}(T) = \int_0^T \|\varepsilon(t)\| dt$ is reported for the own-method runs only.

D.5 Runtime feasibility and fallback protocol

The smooth surge law of Chapter 3 requires the positive root of a quadratic whose discriminant is

$$\Delta_v := (\mathbf{h}^{\top} \hat{\mathbf{v}}_c)^2 - \|\hat{\mathbf{v}}_c\|^2 + \gamma^2 \|\dot{\mathbf{r}}\|^2. \quad (\text{D.9})$$

When $\Delta_v < 0$ no real positive surge exists and the outer loop enters a fallback branch: the last feasible surge is clipped to $[u_{\min}, u_{\max}]$, the angular request is set to zero, and the indicator $\chi_{\text{fb}}(t) \in \{0, 1\}$ is asserted. The fallback duty $f_{\text{fb}} = T^{-1} \int_0^T \chi_{\text{fb}}(t) dt$ is reported for every run. Nominal claims exclude fallback samples. The alignment-margin feasibility check $m_v = (\mathbf{h}^{\top} \bar{\mathbf{v}}_1) \cdot \|\mathbf{v}_1\|$ of Chapter 3 is monitored separately and is *not* the same quantity as the surge discriminant Δ_v introduced here.

D.6 REMUS 100 plant configuration

The dynamic runs use the six-degree-of-freedom REMUS 100 model of Section 4.1. The parameter set is reproduced in Tables D.7 to D.9; the reader is referred to Fossen (2021, §8.4.2) and Allen et al. (2000) for the physical justification of each choice.

Table D.7: REMUS 100 geometry, mass, and environment.

Symbol	Value	Description
L_{auv}	1.6 m	Hull length
D_{auv}	0.19 m	Hull diameter
a	0.8 m	Spheroid semi-major axis
b	0.095 m	Spheroid semi-minor axis
S	0.2128 m^2	Planform area, $0.7 L_{\text{auv}} D_{\text{auv}}$
$\mathbf{r}_{bg}$	$[0, 0, 0.02]^\top \text{ m}$	CG w.r.t. body-frame origin CO
$\mathbf{r}_{bb}$	$[0, 0, 0]^\top \text{ m}$	CB w.r.t. CO
r_{44}	0.3	Roll added-inertia ratio, $A_{44} = r_{44} I_x$
ρ	1026 kg m^{-3}	Water density
μ	63.4468°	Latitude (Trondheim)
g_μ	$\approx 9.815 \text{ m s}^{-2}$	Local gravity

Table D.8: Hydrodynamic damping, added-mass, and lift/drag parameters.

Symbol	Value	Description
T_1, T_2	20 s, 20 s	Surge/sway time constants
T_6	1 s	Yaw time constant
ζ_4, ζ_5	0.3, 0.8	Roll and pitch damping ratios
Damping decay	$D_{11}, D_{22} \propto e^{-3U_r}$	Surge/sway damping vs relative speed
Cross-flow drag	strip-theory integration	Applied to v_r
Body lift/drag	spheroidal lift/drag model	$C_D^0 = 0.42\pi b^2/S$, $C_D = 0.42$
Added-mass matrix	prolate-spheroid coefficients	Standard spheroidal formulation
Cross-coupling	Selected C_A entries zeroed	For stability without quadratic rot. damping

The generalised propulsion force vector is composed from $X_r = -\frac{1}{2}\rho U_{rh}^2 A_r C_{L,\delta_r} \delta_r^2$, $X_s = -\frac{1}{2}\rho U_{rv}^2 A_s C_{L,\delta_s} \delta_s^2$, $Y_r = -\frac{1}{2}\rho U_{rh}^2 A_r C_{L,\delta_r} \delta_r$, $Z_s = -\frac{1}{2}\rho U_{rv}^2 A_s C_{L,\delta_s} \delta_s$, and $X_{\text{prop}}, K_{\text{prop}}$ from the KT/KQ linear-in- J_a polynomials. The propeller-induced roll moment is scaled down by a factor of 10 to match the reference-model calibration.

D.7 Ocean-current model

Two current representations were considered. The yaw-only representation is

$$u_c = \|\mathbf{v}_c\| \cos(\beta_{v_c} - \psi), \quad v_c = \|\mathbf{v}_c\| \sin(\beta_{v_c} - \psi), \quad w_c = v_{c,z}, \quad (\text{D.10})$$

Table D.9: Propeller and fin parameters.

Symbol	Value	Description
D_{prop}	0.14 m	Propeller diameter (three blades, $A_E/A_0 = 0.718$)
t_{prop}	0.1	Thrust deduction number
V_a	$0.944 U$	Advance speed
$J_{a,\text{max}}$	0.6632	At n_{max} , $U = 2.5 \text{ m s}^{-1}$
$K_{T,0}, K_{T,\text{max}}$	0.4566, 0.1798	Thrust coefficients
$K_{Q,0}, K_{Q,\text{max}}$	0.0700, 0.0312	Torque coefficients
S_{fin}	$6.65 \times 10^{-3} \text{ m}^2$	Single-fin area
$A_r = A_s$	$1.33 \times 10^{-2} \text{ m}^2$	Rudder and stern-plane area
C_{L,δ_r}	0.5	Rudder lift coefficient
C_{L,δ_s}	0.7	Stern-plane lift coefficient
$x_r = x_s$	$-a = -0.8 \text{ m}$	Fin longitudinal position

and treats the horizontal current as yaw-projected only; under large pitch excursions (terrain following) the observer sees a rotating apparent current and the current-estimate error persists. The frame-consistent full-rotation representation used throughout Chapter 5 is

$$\mathbf{v}_c^{\text{body}} = \mathbf{R}_{zyx}(\phi, \theta, \psi)^\top \mathbf{v}_c, \quad (\text{D.11})$$

which produces a genuinely constant NED current and is required for $\|\tilde{\mathbf{v}}_c\| \rightarrow 0$ under prolonged pitch excursions. The NED current vector used in the reported runs is

$$\mathbf{v}_c = [0.35, 0.22, 0.06]^\top \text{ m s}^{-1}, \quad \|\mathbf{v}_c\| \approx 0.42 \text{ m s}^{-1}, \quad (\text{D.12})$$

oriented at $\approx 32^\circ$ east of north with a slight downward component. A parametric sweep of $\|\mathbf{v}_c\|$ over $\{0, 0.1, 0.2, 0.35, 0.5\}$ m/s is left as future work; see Section 5.5.

D.8 Numerical settings

All numerical simulations were carried out in MATLAB 2025. The ideal plant runs use the reduced kinematic equations directly, whereas the dynamic runs use the REMUS 100 model implemented through the Fossen MSS Toolbox. Both plant levels are integrated with a fixed-step fourth-order Runge–Kutta scheme. No pseudo-random number generator is invoked in the simulation setup; every run is deterministic for a given parameter set. Table D.10 summarises the numerical settings.

D.9 Observer frame-consistency experiment

Case D5 isolates the implementation issue analysed in Section 4.2. Two paired closed-loop runs are executed with the yaw-only and the full-rotation body-current maps of Section D.7. Both share guidance law, gains, intended NED current, initial state, integrator, and — since the simulation is deterministic — trajectories are exactly comparable within each map. The

Table D.10: Numerical integration settings.

Symbol	Value	Description
h (default)	0.05 s	Integration and control step
h (survey paths)	0.10 s	Used on lawnmower / oceanscan
T (default)	300 s	Overridden by path/user
Integrator	fixed-step RK4	rk4 (Fossen MSS Toolbox)
ε_e	10^{-4}	Threshold treated as $\ \mathbf{e}\ \approx 0$
ε_{den}	10^{-3}	Denominator floor in ω -inversion
Randomisation	none	Deterministic runs
Software environment	MATLAB 2025	Fossen MSS Toolbox used for the REMUS 100 dynamic plant

instantaneous plant states diverge between the two maps because the current representation itself differs. The comparison is conducted first on the helix and then on the oceanscan reference, because the frame-conversion error in (4.20) is attitude-dependent.

Table D.11: Controlled conditions for the observer frame-consistency experiment (case D5).

Item	Helix test	Oceanscan test
Reference geometry	(D.1)	(D.4)
Intended NED current $\mathbf{v}_c$	(D.12)	(D.12)
Initial pose and velocity	identical paired runs	identical paired runs
Observer gains and initial estimate	identical paired runs	identical paired runs
Current representation	yaw-only vs full rotation	yaw-only vs full rotation
Recorded attitude statistic	peak/RMS roll and pitch	peak/RMS roll and pitch
Primary observer output	componentwise $\tilde{\mathbf{v}}_c$ and RMS norm	componentwise $\tilde{\mathbf{v}}_c$ and RMS norm
Coupled performance output	depth/path error, $\mathbf{d}_{\text{al}}$, and δ	depth/path error, $\mathbf{d}_{\text{al}}$, and δ

The full-rotation result is compared with the homogeneous observer prediction; the yaw-only result is interpreted through the forced system (4.21). Numerical values are reported only from paired runs with identical numerical settings; this prevents a frame-definition change from being confounded with a gain or scenario change. A separate observer-replay test feeds one recorded attitude and water-relative-velocity history to both current maps and observers. This open-loop replay holds vehicle motion fixed and therefore isolates the frame-induced observer forcing from the subsequent closed-loop plant response.

D.10 External comparison protocol

The common conditions, method-specific adapters, metrics, and interpretation rules used for the comparative study are summarised below.

D.10.1 Comparative evaluation protocol

Compared implementations

The comparative study contains five implementations:

1. the M-B continuous law developed in this thesis;
2. the switching MED predecessor [4];
3. an exploratory state-dependent path-progression extension;
4. spherical 3-D ALOS [5]; and
5. 3-D virtual-reference-point guidance [6].

The first two provide a predecessor–successor comparison. The third is an exploratory extension and is not covered by the formal analysis of Chapter 3. The last two represent recent LOS and non-LOS three-dimensional path-following designs.

Controlled transfer conditions

Every implementation uses the same REMUS 100 plant, reference path, initial condition, constant NED current $\mathbf{v}_c = [0.35, 0.22, 0.06]^T$ m/s, sampling interval, numerical solver, actuator limits, and anti-windup rule. The external methods retain their published outer-guidance gains where available. Method-specific conversion is nevertheless unavoidable: ALOS uses an attitude-to-rate adapter, while VRP and the thesis-lineage methods deliver different outer-loop command variables. The experiment is therefore a controlled transfer comparison of the complete outer-law–adapter implementations under a shared plant, not a pure guidance-law isolation and not a reproduction of each publication’s native closed loop.

No equal-budget retuning of all five methods is performed. The MED implementation uses the gains declared for the present simulation, while the published external baselines are transferred without path-specific retuning. Every table entry is one deterministic run; no Monte Carlo or parametric robustness envelope is implied.

Reference geometries

The five paths are the descending helix, inclined straight line, C^1 lawnmower, spatial Geronon lemniscate, and terrain-following oceanscan of Section D.1. The lawnmower and oceanscan contain tangent-continuous leg–turn junctions with curvature jumps. The spatial lemniscate has a crossing in horizontal projection but not a three-dimensional self-intersection for the stated vertical parameterisation.

Metrics and reporting policy

The common primary metrics are RMS and peak geometric path distance, settling time, integrated angular effort, surge-command roughness, and saturation duty. Observer error is reported only for implementations that contain the thesis observer. Signed scheduled-reference

error, internal particle-relative error, and terminal progress are architecture-dependent quantities and are not ranked in a common column.

The geometric distance is evaluated using the projection procedure of Section D.4. Values at or below the reference discretisation scale are not used to claim sub-decimetre superiority. If a method does not enter and remain in the prescribed tube within the simulation horizon, the settling time is reported as $T_s = \infty$. This means that settling was not observed under the stated transfer protocol; it is not a claim of general instability of the published method. Conclusions are drawn per path and per metric, and all comparative claims retain this scope.

Appendix E

Conditional Cascade-Stability Analysis

This appendix collects supplementary calculations for a possible stability analysis of the interconnection between the kinematic guidance loop and the REMUS 100/PI implementation. The thesis does not prove closed-loop stability of the full six-degree-of-freedom cascade; the main analytical result remains the kinematic convergence analysis of Chapter 3. The estimates below are conditional: they hold only under the inner-loop hypotheses stated in each section. They are included to clarify what would be required to extend the kinematic result to the dynamic implementation studied numerically in Chapter 5.

The three routes below proceed from the weakest and most general hypothesis to the strongest and most restrictive.

E.1 Interconnection model and standing hypotheses

Throughout this section the full-rotation current map of (4.22) is in force, so by Theorem 4.1 the observer error is a homogeneous exponentially stable mode and $\hat{\mathbf{v}}_c \rightarrow \mathbf{v}_c$; its transient enters the analysis only as an exponentially decaying, absolutely integrable signal, exactly as in the kinematic composite argument of Section 3.3.8. For the purposes of the analytical routes of this appendix, the interconnection between the outer position error $\mathbf{e}$ and the inner-loop velocity-tracking error δ is *modelled as* a series interconnection

$$\dot{\xi} = f_i(\xi, \mathbf{w}_i), \quad \delta = h_i(\xi), \quad \dot{\mathbf{e}} = \dot{\mathbf{r}} - \mathbf{v}_1 - \delta.$$

In the actual closed loop this simplification is not exact: the outer error influences the guidance command $\mathbf{v}_1$, whose rate is in turn the input to the inner loop, and the inner-loop velocity error δ re-enters the outer dynamics. The series model of the next display is therefore a *singular-perturbation* approximation that is expected to be tight only when the inner loop is much faster than the outer one; each estimate derived below should be read under this hypothesis

$$\Sigma_i : \dot{\xi} = f_i(\xi, \mathbf{w}_i), \quad \delta = h_i(\xi), \quad \Sigma_o : \dot{\mathbf{e}} = \dot{\mathbf{r}} - \mathbf{v}_1 - \delta, \quad (\text{E.1})$$

where ξ collects the inner-loop state (REMUS body velocities and the PI integrators of (4.28)), $\mathbf{w}_i := \dot{\mathbf{v}}_1$ is the command rate seen by the inner loop, and δ is the water-relative velocity-tracking error of (4.34). Differentiating $V_e = \frac{1}{2} \|\mathbf{e}\|^2$ along the outer dynamics of (E.1) gives the *perturbed position identity*

$$\dot{V}_e = \|\dot{\mathbf{r}}\| \Xi - \mathbf{e}^\top \delta, \quad (\text{E.2})$$

where Ξ is the (nonpositive) ideal position-channel expression of Theorem 3.16. The interconnection is a cascade: the inner block drives the outer block through δ , and the outer block feeds back into the inner block only through the guidance command rate $\dot{\mathbf{v}}_1$, which is treated as the exogenous input to the inner loop.

Two standing hypotheses close the analysis. The first isolates the minimum information about the reference and the surge policy that is required for the outer estimate to be uniform in time; the second isolates the minimum information about the inner loop that is required for it to have an ISS estimate with respect to the command rate.

Definition E.1 (ISS and ISS-Lyapunov function). A system $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$ is *input-to-state stable* (ISS) if there exist $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$ such that $\|\mathbf{x}(t)\| \leq \beta(\|\mathbf{x}_0\|, t) + \gamma(\sup_{\tau \leq t} \|\mathbf{u}(\tau)\|)$ along every trajectory. A smooth V is an *ISS-Lyapunov function* if there exist $\underline{\alpha}, \bar{\alpha}, \alpha_3 \in \mathcal{K}_\infty$ and $\chi \in \mathcal{K}$ such that $\underline{\alpha}(\|\mathbf{x}\|) \leq V \leq \bar{\alpha}(\|\mathbf{x}\|)$ and

$$\|\mathbf{x}\| \geq \chi(\|\mathbf{u}\|) \implies \nabla V \cdot f(\mathbf{x}, \mathbf{u}) \leq -\alpha_3(\|\mathbf{x}\|). \quad (\text{E.3})$$

Existence of an ISS-Lyapunov function is equivalent to ISS [43].

Assumption E.2 (Reference and feasibility regularity). The reference satisfies Theorem 3.1, so that $v_r^{\min} \leq \|\dot{\mathbf{r}}(t)\| \leq v_r^{\max}$ and $\|\ddot{\mathbf{r}}(t)\| \leq a_r^{\max}$ for all $t \geq 0$. The surge policy is feasible in the sense of Section D.5, so that the commanded velocity $\mathbf{v}_1 = u_{\text{cmd}} \mathbf{h} + \hat{\mathbf{v}}_c$ and its rate $\dot{\mathbf{v}}_1$ are bounded: there exists $W_1 > 0$ with $\sup_t \|\dot{\mathbf{v}}_1(t)\| \leq W_1$. Let

$$\underline{\gamma} := \min_{t \geq 0} \gamma(t) = \begin{cases} \gamma_-, & \text{hysteresis policy (3.65),} \\ 1 - \eta, & \text{continuous policy (3.79),} \end{cases} \quad (\text{E.4})$$

denote the *effective speed-retention floor* of the surge policy.

Assumption E.3 (Explicit outer-decay dominance). For the continuous Modulation B law on the selected compact admissible domain, the nominal decay function of (A.54) satisfies

$$G(\mathbf{e}, \mathbf{b}) \geq \frac{\underline{\gamma} \|\mathbf{e}\|^2}{\underline{\lambda} + \|\mathbf{e}\|}. \quad (\text{E.5})$$

This explicit inequality is stronger than the positive-definiteness used in the layered kinematic proof. It does not follow from the speed floor $\gamma \geq \underline{\gamma}$ alone and has not been verified uniformly for the implemented controller. It is introduced only to make the closed-form ISS and ultimate-bound routes below mathematically traceable. A future analysis may instead construct a different class- $\mathcal{K}$ lower bound for G on the admissible domain.

Assumption E.4 (Inner-loop dissipativity). There exist a smooth $V_{\text{in}} : \mathbb{R}^{n_\xi} \rightarrow \mathbb{R}_{\geq 0}$ and constants $\underline{c}, \bar{c}, C_\delta, c_1, c_2 > 0$ such that, along Σ_i ,

$$\underline{c} \|\delta\|^2 \leq V_{\text{in}}(\xi) \leq \bar{c} \|\xi\|^2, \quad (\text{E.6})$$

$$V_{\text{in}}(\xi) \leq C_\delta \|\delta\|^2, \quad (\text{E.7})$$

$$\dot{V}_{\text{in}} \leq -c_1 \|\delta\|^2 + c_2 \|\mathbf{w}_1\|^2. \quad (\text{E.8})$$

The second inequality (E.7) is an *output-dominance* hypothesis: on the reachable trajectories of the closed loop, the storage V_{in} is required to be comparable to the squared output velocity error $\|\delta\|^2$. It is this hypothesis that makes the estimate below an ISS-Lyapunov bound rather than merely an output-error bound; its verification for the six-degree-of-freedom REMUS/PI plant is part of the same future work as the verification of the dissipation inequality itself.

Theorem E.4 is the dissipativity property that would make the inner block ISS with respect to the command rate. It *may hold locally* for selected linearised channels of the PI-controlled REMUS 100 (e.g. the exponentially stable Nomoto yaw linearisation), but it has *not been verified* for the complete six-degree-of-freedom saturated closed loop of Section 4.1 with its cross-coupling, fin and propeller nonlinearities, PI integrators, anti-windup, and rate clamps. Stability of an individual linearised single-input channel does not, on its own, imply dissipativity of the full multivariable saturated plant. Verifying the assumption for the actual implementation is the analytical work required to turn this appendix into a closed-loop stability proof and is left as future work.

E.2 Outer-loop input-to-state stability

The first cascade route strengthens the pointwise sign condition $\dot{V}_e = \|\dot{\mathbf{r}}\| \Xi \leq 0$ of the ideal channel into a local small-input estimate for the outer error $\mathbf{e}$ driven by the inner-loop velocity error δ , provided the explicit dominance assumption below is available.

Theorem E.5 (Outer-loop local ISS estimate, conditional). *Under the continuous policy (3.79), ideal alignment $\bar{\mathbf{v}}_1 \equiv \bar{\mathbf{a}}$, the corrected current map (4.22), and Theorems E.2 and E.3, the outer error system $\dot{\mathbf{e}} = \dot{\mathbf{r}} - \mathbf{v}_1 - \delta$ admits the following local ISS-Lyapunov implication. For every $\theta \in (0, 1)$,*

$$\|\mathbf{e}\| \geq \frac{(\lambda + \|\mathbf{e}\|) \|\delta\|}{\theta \underline{\gamma} \|\dot{\mathbf{r}}\|} \implies \dot{V}_e \leq -(1 - \theta) \|\dot{\mathbf{r}}\| \underline{\gamma} \frac{\|\mathbf{e}\|^2}{\lambda + \|\mathbf{e}\|}. \quad (\text{E.9})$$

Let $d := \sup_{\tau \geq 0} \|\delta(\tau)\|$. If $d < \theta \underline{\gamma} v_r^{\min}$, the implicit threshold in (E.9) can be solved explicitly, giving the local gain

$$\limsup_{t \rightarrow \infty} \|\mathbf{e}(t)\| \leq \gamma_o(d), \quad \gamma_o(d) = \frac{\lambda d}{\theta \underline{\gamma} v_r^{\min} - d}. \quad (\text{E.10})$$

For $d \rightarrow 0$, $\gamma_o(d) = \lambda d / (\theta \underline{\gamma} v_r^{\min}) + o(d)$. No gain is claimed when the denominator is not positive.

Proof. Differentiating $V_e = \frac{1}{2} \|\mathbf{e}\|^2$ along $\dot{\mathbf{e}} = \dot{\mathbf{r}} - \mathbf{v}_1 - \delta$ and using the nominal identity $\dot{V}_e|_{\delta=0} = \|\dot{\mathbf{r}}\| \Xi$ of (E.2),

$$\dot{V}_e = \|\dot{\mathbf{r}}\| \Xi - \mathbf{e}^\top \delta \leq -\|\dot{\mathbf{r}}\| \psi(s, \|\mathbf{e}\|) + \|\mathbf{e}\| \|\delta\|, \quad (\text{E.11})$$

where the first inequality uses the two-branch bound $\Xi \leq -\psi(s, \|\mathbf{e}\|)$ of Theorem 3.16 and Cauchy–Schwarz on the perturbation term. Applying the additional dominance condition

(E.5) then gives

$$\dot{V}_e \leq -\|\dot{\mathbf{r}}\| \underline{\gamma} \frac{\|\mathbf{e}\|^2}{\lambda + \|\mathbf{e}\|} + \|\mathbf{e}\| \|\delta\|. \quad (\text{E.12})$$

Split the negative-definite term as $\theta + (1 - \theta)$ with $\theta \in (0, 1)$. The perturbation is dominated by the θ -fraction whenever

$$\|\mathbf{e}\| \|\delta\| \leq \theta \|\dot{\mathbf{r}}\| \underline{\gamma} \frac{\|\mathbf{e}\|^2}{\lambda + \|\mathbf{e}\|} \iff \|\mathbf{e}\| \geq \frac{(\lambda + \|\mathbf{e}\|) \|\delta\|}{\theta \underline{\gamma} \|\dot{\mathbf{r}}\|}, \quad (\text{E.13})$$

which is (E.9); on that set $\dot{V}_e \leq -(1 - \theta) \|\dot{\mathbf{r}}\| \underline{\gamma} \|\mathbf{e}\|^2 / (\lambda + \|\mathbf{e}\|) < 0$. To obtain an explicit threshold, set $d := \sup_{\tau \geq 0} \|\delta(\tau)\|$ and use $\|\dot{\mathbf{r}}\| \geq v_r^{\min}$ in (E.13). Rearrangement gives

$$\|\mathbf{e}\| (\theta \underline{\gamma} v_r^{\min} - d) \geq \lambda d. \quad (\text{E.14})$$

When $d < \theta \underline{\gamma} v_r^{\min}$, this is precisely $\|\mathbf{e}\| \geq \gamma_o(d)$ with (E.10). Standard local comparison arguments give the stated limit superior. If the denominator is non-positive, this calculation supplies no finite threshold and no conclusion is asserted. $\square$

Remark E.6 (Role of the surge policy in the outer gain). The gain (E.10) is inversely proportional to the product $\underline{\gamma} v_r^{\min}$: the larger the guaranteed deceleration floor $\underline{\gamma}$ and the reference-speed floor $v_r^{\min}$, the smaller the ultimate outer error for a given inner-loop velocity error. This is the quantitative content of the surge-law design $\|\mathbf{v}_1\| \geq \|\dot{\mathbf{r}}\| > \|\hat{\mathbf{v}}_c\|$ (cf. (3.46)), now expressed as an ISS gain.

Remark E.7 (Local and small-input scope of the outer gain). Two restrictions matter. First, the assumed outer-decay inequality is posed only on a selected compact admissible domain; the gain is therefore local in the state. Second, the explicit denominator in (E.10) is positive only for $d < \theta \underline{\gamma} v_r^{\min}$. Thus the result is also local in the input magnitude. Equation (E.9) should not be read as a data-verified invariant ball. The separate result in Theorem E.17 obtains an explicit ultimate ball only after assuming an eventual uniform bound on δ .

E.3 Inner-loop input-to-state stability

The second building block states the dissipativity property that would be sufficient to make the PI-controlled inner loop ISS with respect to the command rate. It is a conditional route, not a verified property of the REMUS/PI model.

Remark E.8 (Scope downgrade of the inner-loop result). The statement of Theorem E.9 below is best read as a *velocity-error bound conditional on Theorem E.4 and on an additional storage lower bound* $V_{\text{in}} \leq C_\delta \|\delta\|^2$ that has not been established for the six-degree-of-freedom REMUS/PI plant of Section 4.3. Without that additional bound the estimate bounds the tracking error δ but does not establish boundedness of the full inner state ξ .

Theorem E.9 (Inner-loop output-ISS velocity-error bound, conditional). *Under Theorems E.2 and E.4, V_{in} is an ISS-Lyapunov function with respect to the output velocity error δ for the*

inner loop Σ_i driven by the command rate $\mathbf{w}_i = \dot{\mathbf{v}}_1$; in particular

$$\|\delta(t)\|^2 \leq \frac{C_\delta}{\underline{c}} e^{-(c_1/C_\delta)t} \|\delta_0\|^2 + \frac{c_2 C_\delta}{\underline{c} c_1} \sup_{\tau \leq t} \|\dot{\mathbf{v}}_1(\tau)\|^2. \quad (\text{E.15})$$

Consequently the input–output map from $\dot{\mathbf{v}}_1$ to δ admits an ISS estimate with $\mathcal{KL}$ term exponential in rate $c_1/(2C_\delta)$ and linear class- $\mathcal{K}$ gain

$$\gamma_i(s) = \sqrt{\frac{c_2 C_\delta}{\underline{c} c_1}} s. \quad (\text{E.16})$$

The estimate is an output-error bound and does not, on its own, establish boundedness of the complete internal state ξ ; see Theorem E.10.

Remark E.10 (Scope of the inner-loop bound). The lower storage bound $\underline{c} \|\delta\|^2 \leq V_{\text{in}}$ used above is a bound on the *output* velocity error δ of the inner loop, not on the full internal state ξ . Consequently the estimate (E.15) gives an ISS-type output bound from the command rate $\dot{\mathbf{v}}_1$ to the velocity error δ , and does not by itself establish boundedness of the plant velocities, the PI integrator states, and the attitude states that also enter ξ . An ISS-in-state result for the complete inner loop would require the additional storage bound $\underline{c}_\xi \|\xi\|^2 \leq V_{\text{in}}(\xi)$ for some $\underline{c}_\xi > 0$; this bound is neither assumed nor verified for the six-degree-of-freedom REMUS 100 with the PI cascade of Section 4.3. Where the term “inner-loop ISS” is used elsewhere in this appendix, it refers to the velocity-error version (E.15).

Proof. From (E.8), along Σ_i ,

$$\dot{V}_{\text{in}} \leq -c_1 \|\delta\|^2 + c_2 \|\mathbf{w}_i\|^2 \leq -\frac{c_1}{C_\delta} V_{\text{in}} + c_2 \|\mathbf{w}_i\|^2, \quad (\text{E.17})$$

using the output-dominance bound $V_{\text{in}} \leq C_\delta \|\delta\|^2$ from (E.7), which is the step that requires the strengthened assumption of Theorem E.4 and fixes the decay rate at c_1/C_δ . Applying the comparison lemma [43, Lem. 3.4] to (E.17) with that same rate,

$$V_{\text{in}}(t) \leq e^{-(c_1/C_\delta)t} V_{\text{in}}(0) + c_2 \int_0^t e^{-(c_1/C_\delta)(t-\tau)} \|\mathbf{w}_i(\tau)\|^2 d\tau \leq e^{-(c_1/C_\delta)t} V_{\text{in}}(0) + \frac{c_2 C_\delta}{c_1} \sup_{\tau \leq t} \|\mathbf{w}_i\|^2. \quad (\text{E.18})$$

Applying the lower storage bound $\underline{c} \|\delta\|^2 \leq V_{\text{in}}$ on the left and the output-dominance bound $V_{\text{in}}(0) \leq C_\delta \|\delta_0\|^2$ on the right gives (E.15). Its square root is bounded by a $\mathcal{KL}$ term plus the linear class- $\mathcal{K}$ function (E.16) of $\sup_\tau \|\dot{\mathbf{v}}_1\|$, which is the stated output-ISS estimate. By Theorem E.2, $\sup_t \|\dot{\mathbf{v}}_1\| \leq W_1 < \infty$, so δ is bounded. $\square$

Remark E.11 (When Theorem E.4 holds). For the Nomoto yaw channel $T_{\text{yaw}} \dot{r} + r = K_{\text{yaw}} \delta_r$ implicit in the PI-loop design of (4.28), the natural inner-loop storage $V_{\text{in}}^{\text{yaw}} = \frac{1}{2} e_r^2$ with $e_r = r_{\text{cmd}} - r_b$ gives $\dot{V}_{\text{in}}^{\text{yaw}} = -T_{\text{yaw}}^{-1} e_r^2 + e_r \dot{r}_{\text{cmd}}$, i.e. (E.8) with $c_1 = T_{\text{yaw}}^{-1}$ and $c_2 = T_{\text{yaw}}/4$ after Young’s inequality. Analogous quadratic storage functions apply to the pitch loop and to the surge loop of (4.28); the estimate (E.15) degrades (the effective c_1 drops) precisely while the rate clamp (4.24) or the fin and propeller saturation of (4.29) is active, which is the analytical origin of the ultimate bound in Theorem E.17 below.

E.4 Cascade input-to-state stability (Route 1)

Composing Theorem E.5 and Theorem E.9 gives a conditional local estimate for the input–output chain from command rate to outer position error. It does not establish full-state ISS of the REMUS/PI interconnection.

Theorem E.12 (Conditional cascade input–output estimate). *Under Theorems E.2 to E.4, the interconnection (E.1) of the output-ISS inner loop (Theorem E.9) and the local ISS outer loop (Theorem E.5) yields a local ISS estimate from the command rate $\dot{\mathbf{v}}_1$ to the outer error $\mathbf{e}$; on the local regime of Theorem E.7, define*

$$d_i := \gamma_i(\sup_{\tau} \|\dot{\mathbf{v}}_1(\tau)\|). \quad (\text{E.19})$$

If $d_i < \theta \underline{\gamma} v_r^{\min}$, then

$$\limsup_{t \rightarrow \infty} \|\mathbf{e}(t)\| \leq \gamma_o(d_i) = \frac{\lambda d_i}{\theta \underline{\gamma} v_r^{\min} - d_i}, \quad (\text{E.20})$$

where the factor C_δ enters through the output-ISS gain γ_i of (E.16). If $\dot{\mathbf{v}}_1(t) \rightarrow \mathbf{0}$ then $\delta(t) \rightarrow \mathbf{0}$ by Theorem E.9 and $\mathbf{e}(t) \rightarrow \mathbf{0}$ on the local regime by outer ISS.

Proof. The two blocks form a local input–output cascade: Σ_i has the conditional output-ISS estimate from $\dot{\mathbf{v}}_1$ to δ (Theorem E.9), and Σ_o has the conditional local estimate from δ to $\mathbf{e}$ (Theorem E.5). By the cascade theorem for ISS systems [41, 43, Lem. 4.7], their local estimates compose while both trajectories remain in the assumed admissible domains. Let b_i, γ_i and b_o, γ_o denote the corresponding $\mathcal{KL}$ and gain functions. Then

$$\|\mathbf{e}(t)\| \leq b_o(\|\mathbf{e}_0\|, t) + \gamma_o\left(b_i(\|\xi_0\|, t/2) + \gamma_i(\sup_{\tau} \|\dot{\mathbf{v}}_1\|)\right), \quad (\text{E.21})$$

which is a local $\mathcal{KL}$ estimate plus a class- $\mathcal{K}$ function of $\sup_{\tau} \|\dot{\mathbf{v}}_1\|$. Taking $\limsup_{t \rightarrow \infty}$ eliminates the $\mathcal{KL}$ terms and leaves $\limsup \|\mathbf{e}\| \leq \gamma_o(\gamma_i(\sup \|\dot{\mathbf{v}}_1\|))$; substituting the linear γ_i of Theorem E.9 and the exact local γ_o of Theorem E.5 gives (E.20), provided its denominator is positive. If $\dot{\mathbf{v}}_1(t) \rightarrow \mathbf{0}$ then $\gamma_i(\sup_{\tau \geq T} \|\dot{\mathbf{v}}_1\|) \rightarrow 0$ as $T \rightarrow \infty$, so $\delta \rightarrow \mathbf{0}$ and hence $\mathbf{e} \rightarrow \mathbf{0}$ by outer-loop ISS. $\square$

Remark E.13 (Physical reading of the conditional cascade gain). Estimate (E.20) makes the qualitative claim of Section 4.3 quantitative within its local small-input domain: for sufficiently small command rate, the ultimate outer error is approximately proportional to $\sup \|\dot{\mathbf{v}}_1\|$ (turn aggressiveness) and to the inner-loop gain ratio $\sqrt{c_2/(\underline{c} c_1)}$, and inversely proportional to $\underline{\gamma} v_r^{\min}$. Slower manoeuvres (smaller $\dot{\mathbf{v}}_1$), a faster inner loop (larger c_1), a higher reference speed floor, or a larger deceleration margin all shrink the ultimate error; this interpretation is valid only when the stated storage and small-input assumptions have first been verified.

E.5 Composite Lyapunov route (Route 2)

The input–output route above uses an abstract inner-loop storage inequality rather than the complete REMUS equations. If a storage function V_{in} satisfying Theorem E.4 can be constructed, a composite candidate yields the conditional estimate below. The required storage inequality and time-scale separation have not been verified for the REMUS/PI implementation.

Theorem E.14 (Conditional composite estimate). *Let Theorems E.2 to E.4 hold and consider the composite function*

$$V_{\text{cas}} := \alpha V_e + V_a + V_{\text{in}}, \quad \alpha > 0, \quad (\text{E.22})$$

where V_e , V_a , and V_{in} are the position, alignment, and inner-loop storage functions of the preceding sections. Define the near-origin outer decay rate

$$\kappa_o := \frac{\gamma v_r^{\min}}{\lambda} > 0 \quad (\text{E.23})$$

from the near-target estimate of Theorem 3.16. If the inner-loop dissipation rate c_1 and the coupling weight α satisfy the time-scale-separation inequality

$$\boxed{c_1 \kappa_o > \frac{\alpha}{4}, \quad \text{equivalently} \quad c_1 > \frac{\alpha}{4 \kappa_o}}, \quad (\text{E.24})$$

then there exist $\varepsilon^* > 0$ and $c_3 > 0$ such that, in a neighbourhood of $(\mathbf{e}, V_a, \delta) = (\mathbf{0}, 0, \mathbf{0})$ and modulo the exponentially decaying observer transient,

$$\dot{V}_{\text{cas}} \leq -c_3 (\|\mathbf{e}\|^2 + V_a + \|\delta\|^2) + c_2 \|\dot{\mathbf{v}}_1\|^2. \quad (\text{E.25})$$

If $\dot{\mathbf{v}}_1 \equiv \mathbf{0}$, the origin of the conditional autonomous error model is locally asymptotically stable. If instead $\dot{\mathbf{v}}_1 \in L^2[0, \infty)$ and $\dot{\mathbf{v}}_1(t) \rightarrow \mathbf{0}$, convergence of the error variables follows only with the additional boundedness and uniform-continuity conditions stated in the proof. For merely bounded $\dot{\mathbf{v}}_1$, the inequality is a local practical estimate rather than a convergence theorem. A persistently curved moving reference need not satisfy the vanishing-command-rate condition.

Remark E.15 (Formal-candidate status of V_{cas}). The composite storage $V_{\text{cas}} = \alpha V_e + V_a + V_{\text{in}}$ sums quantities of different physical units (metres-squared, dimensionless, and mixed inner-state units). It is used here as a *formal Lyapunov candidate* for the purpose of a conditional analytical route; a dimensionally consistent version would replace each term by its scaled counterpart, e.g. $\mathbf{e}^\top \mathbf{S}_e \mathbf{e}$ with the scaling matrices $\mathbf{S}_e$, $\mathbf{S}_\xi$, $\mathbf{S}_\delta$ of Theorem E.2. The numerical values of these matrices are not identified in the present thesis; their identification is part of the same future work as the verification of Theorem E.4 for the REMUS/PI plant.

Proof. Differentiating (E.22). Near the origin, the local exponential estimate implied by Theorem 3.16 together with the perturbed derivative (E.2) gives $\dot{V}_e \leq -\kappa_o \|\mathbf{e}\|^2 - \mathbf{e}^\top \delta$. For

the alignment term, on the sublevel set $V_a \leq V_{a,\max} < 1$ the logistic decay of Theorem 3.7 yields $\dot{V}_a = -2k V_a(1 - V_a) \leq -k_a V_a$ with $k_a := 2k(1 - V_{a,\max}) > 0$. For the inner-loop term, Theorem E.4 gives $\dot{V}_{\text{in}} \leq -c_1 \|\delta\|^2 + c_2 \|\dot{v}_1\|^2$. Collecting,

$$\dot{V}_{\text{cas}} \leq -\alpha \kappa_o \|\mathbf{e}\|^2 - k_a V_a - c_1 \|\delta\|^2 - \alpha \mathbf{e}^\top \delta + c_2 \|\dot{v}_1\|^2. \quad (\text{E.26})$$

The sign-indefinite cross term is bounded via Young's inequality: for any $\varepsilon > 0$,

$$-\alpha \mathbf{e}^\top \delta \leq \frac{\alpha \varepsilon}{2} \|\mathbf{e}\|^2 + \frac{\alpha}{2\varepsilon} \|\delta\|^2. \quad (\text{E.27})$$

Substituting and grouping the quadratic terms,

$$\dot{V}_{\text{cas}} \leq -\left(\alpha \kappa_o - \frac{\alpha \varepsilon}{2}\right) \|\mathbf{e}\|^2 - k_a V_a - \left(c_1 - \frac{\alpha}{2\varepsilon}\right) \|\delta\|^2 + c_2 \|\dot{v}_1\|^2. \quad (\text{E.28})$$

Both quadratic coefficients are *strictly* positive precisely when ε lies in the open interval

$$\frac{\alpha}{2c_1} < \varepsilon < 2\kappa_o, \quad (\text{E.29})$$

the left inequality forcing $c_1 - \alpha/(2\varepsilon) > 0$ and the right forcing $\alpha\kappa_o - \alpha\varepsilon/2 > 0$. The interval is non-empty if and only if its endpoints are ordered, $\alpha/(2c_1) < 2\kappa_o$, which is exactly the time-scale separation (E.24). Under (E.24), any admissible ε from (E.29) makes both coefficients strictly positive; the midpoint $\varepsilon^\star := \frac{1}{2}(\alpha/(2c_1) + 2\kappa_o)$ gives one such choice, and with

$$c_3 := \min\left\{\alpha \kappa_o - \frac{\alpha \varepsilon^\star}{2}, k_a, c_1 - \frac{\alpha}{2\varepsilon^\star}\right\} > 0 \quad (\text{E.30})$$

the estimate (E.25) follows. The perturbation term $c_2 \|\dot{v}_1\|^2$ is bounded by Theorem E.2, but boundedness on $[0, \infty)$ does not by itself imply absolute integrability. Two mutually exclusive additional hypotheses are considered separately below.

Case A (zero or integrable command rate). If $\dot{v}_1 \equiv \mathbf{0}$, (E.25) is negative definite and the conditional autonomous error model is locally asymptotically stable. More generally, if $\dot{v}_1 \in L^2[0, \infty)$, the state remains inside the local domain, and the derivative of $\|\mathbf{e}\|^2 + V_a + \|\delta\|^2$ is bounded, integration of (E.25) followed by Barbalat's lemma gives $\mathbf{e} \rightarrow \mathbf{0}$, $V_a \rightarrow 0$, and $\delta \rightarrow \mathbf{0}$. These extra boundedness and uniform-continuity requirements are not properties of the guidance law alone and must be verified for the actual trajectory.

Case B (bounded command rate). For merely bounded $\dot{v}_1$, (E.25) provides only a local practical estimate. No convergence to the origin is claimed. The outer-error bound of Theorem E.17 applies only if the realised velocity error itself also satisfies the eventual uniform bound assumed there. $\square$

Remark E.16 (Verification required for the implemented inner loop). Condition (E.24) is the precise form of the requirement that the inner loop be fast relative to the guidance. The dissipation rate c_1 must exceed $\alpha/(4\kappa_o)$ with $\kappa_o = \underline{\gamma} v_r^{\min}/\lambda$. The REMUS/PI implementation contains surge, yaw, and pitch channels with different bandwidths and may also encounter rate and actuator saturation. Neither a common inner-loop storage function nor a numerical value of c_1 has been identified in this thesis. Consequently, (E.24) is a condition to be

verified in future work, not a verified property of any channel of the implemented plant. The simulation traces can reveal command–response differences, but they cannot replace the dissipativity estimate required by Theorem E.4.

E.6 Conditional ultimate bound (Route 3)

In the presence of persistent rate saturation and sideslip, suppose the realised velocity error δ is bounded uniformly by a constant $\bar{\delta} > 0$ after a finite transient. Under this explicit hypothesis, the outer error admits the conditional ultimate bound derived below. The assumption is stronger than a finite-horizon percentile or an asymptotic lim sup estimate.

Theorem E.17 (Conditional ultimate bound). *Let Theorems E.2 and E.3 hold and suppose the inner-loop error satisfies the eventual uniform bound*

$$\|\delta(t)\| \leq \bar{\delta} \quad \text{for every } t \geq T_\delta \quad (\text{E.31})$$

for some finite T_δ . Then, for every $\theta \in (0, 1)$ and $\bar{\delta} > 0$, the outer error e is ultimately bounded: there exists a finite $T \geq T_\delta$ such that

$$\|e(t)\| \leq r(\theta) := \frac{\lambda \bar{\delta}}{(1 - \theta) \underline{\gamma} v_r^{\min} - \bar{\delta}} \quad \text{for every } t \geq T, \quad (\text{E.32})$$

provided the denominator is positive, i.e. $\bar{\delta} < (1 - \theta) \underline{\gamma} v_r^{\min}$. The ultimate radius grows linearly with the inner-loop velocity-error bound $\bar{\delta}$ and vanishes as $\bar{\delta} \rightarrow 0$.

Proof. From (E.12) together with $\|\dot{r}\| \geq v_r^{\min}$,

$$\dot{V}_e \leq -v_r^{\min} \underline{\gamma} \frac{\|e\|^2}{\lambda + \|e\|} + \|e\| \|\delta\| = \|e\| \left(-v_r^{\min} \underline{\gamma} \frac{\|e\|}{\lambda + \|e\|} + \|\delta\| \right). \quad (\text{E.33})$$

Fix $\theta \in (0, 1)$ and write $A := v_r^{\min} \underline{\gamma}$. For $t \geq T_\delta$ and $\|e\| \geq r(\theta)$, monotonicity of $r \mapsto r/(\lambda + r)$ gives

$$A \frac{\|e\|}{\lambda + \|e\|} \geq A \frac{r(\theta)}{\lambda + r(\theta)} = \frac{\bar{\delta}}{1 - \theta}. \quad (\text{E.34})$$

Substitution into (E.33) yields

$$\dot{V}_e \leq -\frac{\theta \bar{\delta}}{1 - \theta} \|e\| < 0 \quad \text{whenever } \|e\| \geq r(\theta). \quad (\text{E.35})$$

The strict margin makes the ball in (E.32) forward invariant after T_δ and forces finite-time entry if the state is initially outside it. When $\bar{\delta} = 0$, the formula reduces to the nominal asymptotic-convergence case rather than a positive-radius ultimate bound. $\square$

E.7 Summary and unmodified controller

The three routes above form a hierarchy of conditional estimates for one and the same unchanged guidance law of Section 3.3:

- **Conditional cascade input–output estimate (Theorem E.12).** On the local outer regime of Theorem E.7, the interconnection admits an ISS estimate from the command rate $\dot{\mathbf{v}}_1$ to the position error $\mathbf{e}$, with explicit asymptotic gain (E.20). The gain depends on the inner-loop constants $c_1, c_2, \underline{c}, C_\delta$ of Theorem E.4: the factor C_δ measures how tightly the internal storage V_{in} is dominated by the output velocity error δ and enters the cascade gain in the same way as the ratio $c_2/(\underline{c} c_1)$. Because the inner-loop estimate is output-ISS in δ rather than full-state ISS in ξ (Theorem E.10), the cascade estimate applies to $\mathbf{e}$ and δ and does not by itself bound the plant integrators or attitude states.
- **Conditional composite estimate (Theorem E.14).** Under the explicit time-scale separation (E.24), the single composite Lyapunov function (E.22) yields a local practical estimate with respect to the command rate. Local asymptotic stability follows for the zero-input conditional error model; convergence under an L^2 command rate additionally requires the boundedness and uniform-continuity conditions stated in the proof. The claim is dimensionally formal (Theorem E.15) and holds in the non-saturated regime.
- **Conditional ultimate bound (Theorem E.17).** When rate saturation and sideslip leave an eventual uniform velocity-error bound $\bar{\delta} > 0$, the outer error is ultimately bounded by the explicit ball (E.32), whose radius is linear in $\bar{\delta}$; the constant C_δ enters the UUB radius only indirectly through $\bar{\delta}$ (as it controls how δ absorbs the internal storage) and the radius is stated in the physically interpretable variable $\bar{\delta}$. Applying the formula to data requires an independently established uniform bound; a trajectory percentile is insufficient.

All three routes additionally require the explicit outer-decay dominance assumption (E.5). Their remaining ingredients are the outer position sign condition $\Xi \leq 0$ of Theorems 3.14 and 3.16 (which supplies the nominal sign structure of Theorem E.5) and the inner-loop dissipativity of the PI cascade (Theorem E.4, exploited by Theorem E.9). No route requires modifying the kinematic guidance law of Section 3.3 or the command interface of Section 4.3; Route 2 would additionally yield local asymptotic stability when the inner loop is fast enough, and would be made fully constructive by replacing the PI inner loop with a backstepping design that enforces Theorem E.4 with a prescribed dissipation rate c_1 . Such a replacement is outside the scope of the present thesis and is recorded in Chapter 6 as the natural direction for follow-on work.

Bibliography

- [1] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*, 2nd ed. Wiley, 2021.
- [2] N. Gu, D. Wang, Z. Peng, J. Wang, and Q.-L. Han, “Advances in line-of-sight (LOS) guidance for path following of autonomous marine vehicles: An overview,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 53, no. 1, pp. 12–28, 2023.
- [3] L. He, M. Xie, and Y. Zhang, “A review of path following, trajectory tracking, and formation control for autonomous underwater vehicles (AUVs),” *Drones*, vol. 9, no. 4, p. 286, 2025.
- [4] M. Sabouri and G. Indiveri, “Kinematic 3d path following controller for an underactuated underwater robot subject to a constant unknown current,” in *2026 34th Mediterranean Conference on Control and Automation (MED)*, Ancona, Italy, 2026, pp. 914–919.
- [5] E. M. Coates and T. I. Fossen, “A spherical amplitude–phase formulation for 3-D adaptive line-of-sight (ALOS) guidance with USGES stability guarantees,” *Automatica*, vol. 192, p. 113149, 2026, available online.
- [6] L. Degorre, T. I. Fossen, E. Delaleau, and O. Chocron, “A virtual reference point kinematic guidance law for 3-D path-following of autonomous underwater vehicles,” *IEEE Access*, vol. 12, pp. 109 822–109 831, 2024.
- [7] N. Hung, F. Rego, J. Quintas, J. Cruz, M. Jacinto, D. Souto, A. Potes, L. Sebastião, and A. Pascoal, “A review of path following control strategies for autonomous robotic vehicles: Theory, simulations, and experiments,” *Journal of Field Robotics*, vol. 40, no. 3, pp. 747–779, 2023.
- [8] A. M. Eissa, S. A. Mohamed, M. I. Awad, H. E. Abd El Munim, and D. Emad, “Trajectory tracking control of autonomous underwater vehicles (AUVs): A review from classical methods to AI-based approaches,” *International Journal of Dynamics and Control*, vol. 14, no. 1, p. 32, 2026.
- [9] P. Encarnação, A. Pascoal, and M. Arcak, “Path following for marine vehicles in the presence of unknown currents,” in *Proceedings of the 6th IFAC International Symposium on Robot Control*, Vienna, Austria, 2000, pp. 469–474.
- [10] M. Breivik and T. I. Fossen, “Guidance-based path following for autonomous underwater vehicles,” in *Proceedings of the MTS/IEEE OCEANS 2005*, vol. 3, 2005, pp. 2807–2814.
- [11] L. Lapierre, D. Soetanto, and A. Pascoal, “Nonlinear path following with applications to the control of autonomous underwater vehicles,” in *Proceedings of the 42nd IEEE Conference on Decision and Control*, Maui, HI, USA, 2003, pp. 1256–1261.

- [12] A. P. Aguiar, J. P. Hespanha, and P. V. Kokotović, “Path-following for nonminimum phase systems removes performance limitations,” *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 234–239, 2005.
- [13] A. P. Aguiar and J. P. Hespanha, “Trajectory-tracking and path-following of underactuated autonomous vehicles with parametric modeling uncertainty,” *IEEE Transactions on Automatic Control*, vol. 52, no. 8, pp. 1362–1379, 2007.
- [14] A. M. Lekkas and T. I. Fossen, “Integral LOS path following for curved paths based on a monotone cubic hermite spline parametrization,” *IEEE Transactions on Control Systems Technology*, vol. 22, no. 6, pp. 2287–2301, 2014.
- [15] W. Caharija, K. Y. Pettersen, M. Bibuli, P. Calado, E. Zereik, J. Braga, J. T. Gravdahl, A. J. Sørensen, M. Milovanović, and G. Bruzzone, “Integral line-of-sight (LOS) guidance and control of underactuated marine vehicles: Theory, simulations, and experiments,” *IEEE Transactions on Control Systems Technology*, vol. 24, no. 5, pp. 1623–1642, 2016.
- [16] D. Belleter, M. A. Maghenem, C. Paliotta, and K. Y. Pettersen, “Observer based path following for underactuated marine vessels in the presence of ocean currents: A global approach,” *Automatica*, vol. 100, pp. 123–134, 2019.
- [17] T. I. Fossen and A. M. Lekkas, “Direct and indirect adaptive integral line-of-sight path-following controllers for marine craft exposed to ocean currents,” *International Journal of Adaptive Control and Signal Processing*, vol. 31, no. 4, pp. 445–463, 2017.
- [18] S. Moe, K. Y. Pettersen, T. I. Fossen, and J. T. Gravdahl, “Line-of-sight curved path following for underactuated USVs and AUVs in the horizontal plane under the influence of ocean currents,” in *Proceedings of the 24th Mediterranean Conference on Control and Automation (MED)*, Athens, Greece, 2016, pp. 38–45.
- [19] B. M. Abdurahman, A. Savvaris, and A. Tsourdos, “Switching LOS guidance with speed allocation and vertical course control for path-following of unmanned underwater vehicles under ocean current disturbances,” *Ocean Engineering*, vol. 182, pp. 412–426, 2019.
- [20] T. I. Fossen, “An adaptive line-of-sight (ALOS) guidance law for path following of aircraft and marine craft,” *IEEE Transactions on Control Systems Technology*, vol. 31, no. 6, pp. 2887–2894, 2023.
- [21] T. I. Fossen and A. P. Aguiar, “A uniform semiglobal exponential stable adaptive line-of-sight (ALOS) guidance law for 3-D path following,” *Automatica*, vol. 163, p. 111556, 2024.
- [22] X. Xu, B. He, N. Dai, T. Wang, and Y. Shen, “Tracking control study of AUV large curvature path based on artificial physics method,” *Ocean Engineering*, vol. 303, p. 117737, 2024.

- [23] R. Skjetne, T. I. Fossen, and P. V. Kokotović, “Robust output maneuvering for a class of nonlinear systems,” *Automatica*, vol. 40, no. 3, pp. 373–383, 2004.
- [24] G. Indiveri and A. A. Zizzari, “Kinematics motion control of an underactuated vehicle: A 3d solution with bounded control effort,” in *Proceedings of the IFAC Workshop on Navigation, Guidance and Control of Underwater Vehicles*, Killaloe, Ireland, 2008.
- [25] G. Indiveri, S. Cretí, and A. A. Zizzari, “A proof of concept for the guidance of 3d underactuated vehicles subject to constant unknown disturbances,” in *Proceedings of the 9th IFAC Conference on Manoeuvring and Control of Marine Craft*, Arenzano, Italy, 2012, pp. 307–312.
- [26] W. Caharija, K. Y. Pettersen, P. Calado, and J. Braga, “A comparison between the ILOS guidance and the vector field guidance,” *IFAC-PapersOnLine*, vol. 48, pp. 89–94, 2015.
- [27] W. Yao and M. Wang, “Integral vector field guidance with adaptive sliding mode control for 3d path following of underactuated AUVs,” *Ocean Engineering*, vol. 236, p. 109468, 2021.
- [28] M. Bibuli, “Virtual-target based path-following in the 3D underwater environment,” *IFAC-PapersOnLine*, vol. 58, no. 20, pp. 177–182, 2024.
- [29] H. Xu, T. I. Fossen, and C. G. Soares, “Uniformly semiglobally exponential stability of vector field guidance law and autopilot for path-following,” *European Journal of Control*, vol. 53, pp. 88–97, 2020.
- [30] L. Degorre, T. I. Fossen, O. Chocron, and E. Delaleau, “A model-based kinematic guidance method for control of underactuated autonomous underwater vehicles,” *Control Engineering Practice*, vol. 152, p. 106068, 2024.
- [31] S. K. Mallipeddi, M. Menghini, S. Simani, and P. Castaldi, “Kinematic guidance using virtual reference point (VRP) for underactuated marine vehicles with sideslip compensation,” *Control Engineering Practice*, vol. 169, p. 106769, 2026.
- [32] A. M. Lekkas and T. I. Fossen, “Trajectory tracking and ocean current estimation for marine underactuated vehicles,” in *2014 IEEE Conference on Control Applications*, Antibes, France, 2014, pp. 905–910.
- [33] X. Liang, X. Qu, Y. Hou, and Q. Ma, “Three-dimensional trajectory tracking control of an underactuated autonomous underwater vehicle (AUV) based on ocean current observer,” *International Journal of Advanced Robotic Systems*, vol. 15, no. 5, 2018.
- [34] E. Kim, S. Fan, N. Bose, and H. Nguyen, “Current estimation and path following for an autonomous underwater vehicle (AUV) by using a high-gain observer based on an AUV dynamic model,” *International Journal of Control, Automation and Systems*, vol. 19, no. 1, pp. 478–490, 2020.

- [35] L. He, Y. Zhang, S. Li, B. Li, and Z. Yuan, “Three-dimensional path following control for underactuated AUV based on ocean current observer,” *Drones*, vol. 8, no. 11, p. 672, 2024.
- [36] L. Ašanin, L. Martinović, and Ž. Zečević, “Improved line-of-sight guidance law for path following of autonomous underwater vehicle with compensation of drift and lift forces,” in *2025 29th International Conference on Information Technology (IT)*, 2025.
- [37] M. Lei, Y. Li, T. Zhang, and D. Jiang, “Disturbance observer-based anti-windup control for path-following of underactuated AUVs via singular perturbations: Theory and experiment,” *IEEE Transactions on Vehicular Technology*, vol. 73, no. 8, pp. 11 044–11 058, 2024.
- [38] P. Maurya, A. P. Aguiar, and A. Pascoal, “Marine vehicle path following using inner-outer loop control,” in *Proceedings of the 8th International IFAC Conference on Manoeuvring and Control of Marine Craft (MCMC 2009)*, ser. IFAC Proceedings Volumes, vol. 42, no. 18, 2009, pp. 38–43.
- [39] Z. Zhang, M. Lin, D. Li, and R. Lin, “Gaze-assisted prescribed performance controller for AUV trajectory tracking in time-varying currents,” *Journal of Marine Science and Engineering*, vol. 12, no. 9, p. 1643, 2024.
- [40] Y. Qu, Q. Zhang, X. Xiang, S. Yang, and L. Cai, “Cascaded guidance, state estimation, and control for depth tracking of underactuated autonomous underwater vehicles,” *Communications Engineering*, vol. 4, no. 1, pp. 1–12, 2025.
- [41] E. Panteley and A. Loria, “On global uniform asymptotic stability of nonlinear time-varying systems in cascade,” *Systems & Control Letters*, vol. 33, no. 2, pp. 131–138, 1998.
- [42] E. D. Sontag, “Input to state stability: Basic concepts and results,” in *Nonlinear and Optimal Control Theory*, ser. Lecture Notes in Mathematics, P. Nistri and G. Stefani, Eds. Berlin: Springer, 2008, vol. 1932, pp. 163–220.
- [43] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2002.
- [44] C.-T. Chen, *Linear System Theory and Design*, 3rd ed. New York: Oxford University Press, 1999.
- [45] D. E. Koditschek, “The application of total energy as a Lyapunov function for mechanical control systems,” in *Dynamics and Control of Multibody Systems*, ser. Contemporary Mathematics, J. E. Marsden, P. S. Krishnaprasad, and J. C. Simo, Eds. Providence, RI: American Mathematical Society, 1989, vol. 97, pp. 131–157.
- [46] S. P. Bhat and D. S. Bernstein, “A topological obstruction to continuous global stabilization of rotational motion and the unwinding phenomenon,” *Systems & Control Letters*, vol. 39, no. 1, pp. 63–70, 2000.