


COLAV

Anti-collision for a Coordinated Fleet

Rules-governed, benchmark-pinned collision avoidance for a mothership, her unmanned screen and her aircraft, under a single operator.



Pierre-Louis Constant

Director / Lead Engineer, Oceanomatics Pty Ltd

Software baseline: RTGroundControl,
branch solved-evasion-and-audit
Companion to OCX-WP-2026-01 (μDP)

Abstract

RTGroundControl is Oceanomatics' ground station and autonomy stack for a multimodal task group, surface and air, controlled by a single operator from one console. This paper describes its anti-collision system and the tests it has been put through, on a simulated group of one mothership, four unmanned surface vessels and two fixed-wing unmanned aircraft. A hard admissibility filter has the last word over every command; above it a decision core chooses whole manoeuvres, alter, hold and rejoin, scored against every contact and the water over a receding horizon. The Collision Regulations strike out what they forbid and order what remains, and every manoeuvre is reported with the rule that required it and what it would otherwise have hit. Contacts are predicted by one shared model, encounters carry memory, the ground is a clearance field, the passage is run on paper ahead of the ship, and the fleet is one system in which the guide holds her course and her boats conform. The water is the IHO's S-100 product from the official UKHO and SHOM Channel trial datasets, adopted both to implement the new standard and to hold the chart to S-52, against which it has been audited requirement by requirement. Verification rests on three deterministic baselines regenerated by the test suite, a parameter sweep, several hundred simulated passages and a corpus of corrected radar exam papers; the paper also records the three occasions on which a green benchmark coexisted with a live fault.

Contents

1	Introduction	3
1.1	What the paper covers	3
1.2	Provenance	4
2	The design	4
2.1	Overview	4
2.2	Layers	5
2.3	Three principles	5
2.4	What the literature calls it	6
3	The picture	6
3.1	The guard set	6
3.2	The separation floor	6
3.3	Prediction: one shared model	7
3.4	Own ship: one turning law	7
3.5	The ground as a clearance field	8
3.6	The relative plot	8
4	The Rules	8
4.1	Classification	8
4.2	Memory: the encounter register	9
4.3	Obligation	12
4.4	What is forbidden, and in extremis	12
4.5	Citation	13
4.6	Restricted visibility: Rule 19, and what it would change	13
5	The decision: the manoeuvre core	16
5.1	A candidate is a manoeuvre	16
5.2	Hard, forbidden, priced	18
5.3	Tiers	19
5.4	Cost of the search	19
5.5	What it replaced	21
6	The plan: route anticipation	21
6.1	The route run on paper	22
6.2	The resolver	23
6.3	Routing and the marks	24
7	The fleet	25
7.1	The doctrine	25
7.2	The escorts' defence	25
7.3	The boats are not the ship	27
7.4	Formation shapes	28
8	The chart: S-100 data and ECDIS conformance	29
8.1	The dataset	29
8.2	What each product does in the model	30
8.3	Conformance with S-52	31
8.4	Why it matters for collision avoidance	32

9	What the operator sees	32
9.1	The helm report	32
9.2	Helm by consent	33
9.3	Intent on the chart	33
10	Verification	33
10.1	From arguments to experiments	33
10.2	Open water: the Imazu set	34
10.3	Confined water	35
10.4	The fleet's own encounters	35
10.5	The head-on doubt band	35
10.6	Passages: the campaigns	36
10.7	The formation sweep	38
10.8	The radar plotting corpus	39
11	What the benchmarks did not catch	39
11.1	Rules of coverage	40
12	Where the model runs	41
13	Roadmap	42
14	Conclusion	42
A	Settings reference	45
B	The Imazu set as run	48
C	Scoring definitions	49
D	The passage exercise	51

1 Introduction

RTGroundControl is Oceanomatics' ground station and autonomy stack for a multimodal task group under a single operator: one console, one person, every unit on the water and in the air. The group tested in this paper is one mothership, four unmanned surface vessels (USVs) and two fixed-wing unmanned aircraft (UAVs). The mothership, AURORA in the simulations, is a 152 m hull that needs six metres of water and turns at about $1.3^\circ/\text{s}$ at twelve knots. The boats, ALPHA to DELTA, draw half a metre, are ordered at up to sixteen knots in the trials reported here, and take their stations from the ship. The aircraft, UAV ALPHA and UAV BRAVO, fly between 50 and 250 knots and loiter at 400 m radius; they share the operator, the picture and the fleet plan with the boats, and Section 3 explains why they share no separation floor with them. The group transits working water: Cockburn Sound off Fremantle, the Singapore Strait and the Solent are the three operating areas built so far, each with live or replayed AIS traffic, real bathymetry and real charted marks.

Each unit must keep clear of traffic and of the other units as the International Regulations for Preventing Collisions at Sea prescribe [1], must not run aground, should not leave the planned track for vessels that would have passed clear, and must do so in a way that one operator, responsible for all seven units, can follow from the screen and the log. This paper describes the system that meets those requirements and the tests it has been put through.

The companion paper [3] describes μDP , the propulsion-agnostic manual and autonomous control stack that steers a hull to a commanded course and speed. This paper sits one layer above it: it decides what that course and speed should be when something is in the way. The interface between the two is a command, a heading and a speed. The anti-collision layer does not command a thruster, and the control layer does not read a contact.

1.1 What the paper covers

1. the design: a hard admissibility filter with the last word over every command, beneath a decision core that chooses whole manoeuvres (Sections 2 and 5);
2. the picture the decision is made against: contacts, separation floors, one shared predictor, and the ground as a clearance field (Section 3);
3. the rules: how encounters are classified, remembered and obliged, and how the rules act by striking out and ordering rather than by steering (Section 4);
4. the plan: the passage projected ahead of the ship so that a meeting minutes away is answered by a small reduction of speed rather than a late alteration (Section 6);
5. the fleet: what a mothership and her screen owe each other, and why the boats are not the ship (Section 7);
6. the chart: the S-100 trial data the two Channel areas are built on, why it was adopted, and the audit of the presentation against S-52 (Section 8);
7. what the operator sees (Section 9);
8. the tests: three deterministic baselines, a parameter sweep, the passage campaigns and the radar plotting corpus (Section 10);
9. three faults that reached live runs while every benchmark passed, and the coverage rules derived from them (Section 11);
10. deployment on the vehicles, and the roadmap (Sections 12 and 13).

Intended readers are engineers and mariners who know the Rules and are comfortable with the vocabulary of the relative plot [2], not necessarily with the collision avoidance literature. Where the literature has a name for something built here, the name is given; the surveys [12, 14] cover the field.

1.2 Provenance

The model was rebuilt in August 2026 after an external audit. The audit judged the primitives sound (the admissibility filter, the shared predictor, the timing solver and the reporting) and the decision architecture dated: the decision was memoryless, instantaneous, solved against one contact at a time, and evaluated against no benchmark. Those four findings organise this paper. Each section states what replaced the finding, and Section 10 states how the replacement was measured. The audit’s first recommendation, that results be made measurable before the model was changed, was followed; every figure quoted in this paper is regenerated by the test suite from a deterministic scenario, or comes from a passage run whose ledger is committed with the source.

Status of this revision

This is Revision A of a document about a system that is still being developed. The results in Section 10 are those of the software baseline on the cover; they are good enough to publish and not good enough to stop. Section 11 lists faults the benchmarks did not catch, Section 13 lists the work that follows from the evidence, and three figures in Section 5 are illustrative rather than measured. Each of these is expected to change the text, and further revisions will be issued as they do. Numbers quoted here should be read with their date.

Summary

A hard filter refuses any command that projects inside a separation floor or onto foul water. Above it, a decision core chooses a whole manoeuvre (alter, hold, rejoin) scored against every contact and the ground. The Rules remove the options they forbid and order the rest. Every decision is reported with the rule that required it and the contact it avoided. Every change to the model is measured against a published baseline.

2 The design

2.1 Overview

Every step, for every unit, the system builds a list of everything that could be hit, assigns each a *separation floor* it may not close inside, asks the Rules which one governs each encounter, and then refuses to issue any command that would breach a floor or run onto foul water, substituting the nearest lawful command that would not. The filter decides what the vehicle is commanded to do; the Rules act by ordering the choices the filter is allowed to make. A manoeuvre therefore has two independent justifications, reported separately: the rule that required it, and the contact or ground it would otherwise have closed with.

2.2 Layers




Figure 1: The decision chain for one unit at one step, as a stack. The guard set is built at the top; the Rules classify, remember and oblige; the filter is the one place a command can be refused; beneath it are the four primitives every layer shares.

Two things sit outside the chain. **CollisionGuard** displaces a commanded *station* away from traffic before the station is ever steered to; the filter guarantees the *command*. The two are different jobs, and the guard cannot make the guarantee, which is why the filter exists. **RoutePlot** asks of the plan, minutes ahead, the question the filter asks of the command (Section 6).

2.3 Three principles

Hard and soft stay apart. Separation floors and foul water are constraints. No cost buys them. The Rules strike out what they forbid, a port alteration under Rules 14, 15 and 17(c), asked of every vessel at risk and not only of the one that forced the decision. Everything else is priced: the size of the alteration, the side, the speed change, changing one's mind, shaving the bank, leaving the corridor, and how far from the plan the manoeuvre ends. A scheme in which a prohibition is a large penalty will eventually meet a picture in which the penalty is cheaper than the alternatives; the open-water baseline recorded seven such cases in twenty-two before the separation was made (Section 10.2).

One implementation of each judgement. The serious faults found in the system's history share a pattern, two implementations of one judgement drifting apart: two straight-line predictors, one in the classifier and one in the filter, free to disagree about the same vessel; the filter and the policy classifying the same encounter differently, so the layer deciding the manoeuvre saw no encounter at all; six local tests of "is she one of ours" that answered differently. Each was replaced by one

object that the others call. That principle also decides where the model runs on the vehicles (Section 12).

The decision is a manoeuvre, not a heading. A candidate is a whole trajectory: alter by a given angle at a given speed, hold until every contact is past and opening, then steer back onto the planned track. It is integrated through the vehicle’s own turn rate and acceleration and scored at every step. Four faults that the audit listed separately (no recovery leg, a commit deadline measured against a fixed ninety-degree reference, an evasion solved against one contact, and an alteration that could be reversed on the next frame) had a single cause: the model decided a heading every frame and had no representation of a manoeuvre. Section 5 gives the core; Section 5.5 records what it replaced.

2.4 What the literature calls it

In the terms of the literature the pipeline is as follows. The separation floor is a circular ship domain [13]. The admissibility veto is a command-level safety filter, the shape a control barrier function takes when applied to a turning hull [16]. The manoeuvre core is simulation-based model-predictive selection over a small library of whole behaviours, the architecture Johansen, Perez and Cristofaro [4] introduced and Eriksen [6] and Kufoalor [7] took to sea, with this project’s floors, timing oracle and rule precedence dropped in. The velocity-obstacle partition of Kuwata et al. [8] is what the pass-ahead and pass-astern discipline formalises. The scoring is Woerner’s [9]. The audit did not find the combination in the literature: a hard veto beneath the selector, the rule citation derived from the decision itself, and a formation in which the guide holds her course and her screen conforms.

3 The picture

This section describes the inputs to the decision: what a contact is, how far from it the vehicle must stay, where it will be, and how the ground enters the same arithmetic.

3.1 The guard set

Built fresh each step for each unit. It contains every AIS contact with a position, and every own unit that is not under command in the same task group. A boat on her cradle is part of the ship and not a contact; a UAV at altitude and a boat in the water cannot touch, and that severance is asserted rather than assumed. Since September 2026 an own unit under command (a *Subordinate* or a *Consort*) is not in the set at all: the fleet is one coordinated system, its mutual positioning is settled by the formation plan, and reactive avoidance between its members produced the helm-held cascades and in-extremis deadlocks described in Section 7. A unit that has gone off the plan, link lost or under manual control, is a contact again, because “we command her” has stopped being true.

A contact carries position, heading, speed, hull radius, rate of turn where AIS reports it (messages 1 to 3, rejected if it is the not-available sentinel or above 120°/min), vertical extent and its relationship to own ship.

3.2 The separation floor

Each contact has a floor, the distance own ship may not close inside, and a margin of 30 m on top of it. Against traffic the floor is the base traffic clearance, 370.4 m (a fifth of a mile), plus the larger of the two hull radii. Against a detached own unit it is the two hulls plus 20 m of water, or the escort’s clearance ring, whichever is larger. Two asymmetries are deliberate.

Guard circles nest; they do not add. A guard circle drawn about a vessel already contains her hull, so two vessels each carrying one need the larger of the two between their centres, not the sum. The physical minimum stays a sum, because two hulls less than half their lengths apart really are inside one another.

The guide does not give way to boats she commands. The clearance ring is a station-keeping distance and belongs to the escort. Applied to the leader it became a veto on her own passage: in one run ninety of ninety-four helm engagements were caused by her own screen. She holds her course; the boats conform.

Until September 2026 the leader's floor was also inflated by the *formation reach*, the extent of the screen projected toward the contact, so that she manoeuvred the whole group around traffic as one body. Measured in the Solent, that gave her a 1743 m floor against a ferry 2340 m away, and she left her track for ships two kilometres off that would have passed clear. Since the boats give way for themselves and the group is coordinated by the plan, the reach duplicated their cover and was removed: the leader avoids traffic as one ship.

3.3 Prediction: one shared model

ContactMotion is the single place a contact's future position is computed. It is read by the classifier deciding whether a rule applies, by the memory deciding whether an encounter is still open, by the filter deciding whether a command is admissible, and by the route plot. Before it existed, the classifier and the filter each had a straight-line copy and were free to disagree about the same vessel.

A steady contact ($|\omega| < 0.5^\circ/\text{min}$) runs a straight line at present course and speed. A turning contact runs a circular arc, integrated in closed form:

$$h(t) = h_0 + \omega t, \quad E(t) = E_0 + \frac{v}{\omega} (\cos h_0 - \cos h(t)), \quad N(t) = N_0 + \frac{v}{\omega} (\sin h(t) - \sin h_0). \quad (1)$$

After **TurnPersistenceSeconds** (60 s) she is assumed steady on the heading the turn reached. The persistence window is the substantive judgement. A rate of turn is an observation of what a vessel is doing now, not of her intent over the next six minutes, and integrating it across the full look-ahead would carry her round a complete circle and invent encounters as readily as straight-line projection does. Close aboard, a turning contact is held to the worse of her arc and her steadied course, since a projection that assumes she keeps turning away is the optimistic one.

3.4 Own ship: one turning law

Own ship is not a point. Every candidate is run through the vehicle's own turn rate, acceleration and speed limits, and every deadline is derived from them. One law serves the track pilot's wheel-over, the collision commit deadlines and the route plot:

$$R(v) = \max(R_{\min}, v^2/(gG)), \quad \dot{\psi}(v) = \min(\dot{\psi}_{\max}, v/R(v)), \quad (2)$$

with the declared minimum radius $R_{\min}$ at and below cruise and the lateral acceleration bound G widening it as the square of the speed above. For the mothership $R_{\min} = 270$ m and $G = 0.02$, so she turns at about $1.3^\circ/\text{s}$ at twelve knots and the lateral term takes over above about fourteen. For a USV $R_{\min} = 40$ m and $G = 0.20$. These figures matter because the system's early behaviour had been tuned against a ship that turned at $8^\circ/\text{s}$, a rate no 152 m hull can make, so every deadline and reserve had been sized on time she did not have. Correcting the dynamics was the first step of the rebuild, so that the model was tuned once.

3.5 The ground as a clearance field

Terrain was a boolean consulted in one place. It is now a **ClearanceField**: one exact distance transform over the depth grid, built once per hull class, that answers “how far is the nearest foul water” at any point in constant time. It is read by the filter (a candidate that runs inside **TerrainMarginMeters** is struck out like a floor breach), by the manoeuvre core (a rising cost through a caution band of 250 m, weighted 45, and an uncapped cost for leaving the route corridor), by the route planner (a berth of the width asked for, at any radius, instead of the square Chebyshev dilation it replaced), and by the corridor alarm, which is derived from the same erosion rather than set beside it.

Each hull class has its own field. The ship’s is cut at six metres of water; the boats’ at half a metre. The distinction was enforced late, and until it was, the boats were held to the ship’s water (Section 7.3).

3.6 The relative plot

Beneath everything sits **RelativeMotion**: the relative-motion triangle of Pub. 1310 [2], giving direction and speed of relative motion, CPA and TCPA, the manoeuvring authority of each vessel, and the Pub. 1310 categories of contact. It uses the vocabulary of the bridge rather than of robotics, and it is the layer the radar plotting corpus of Section 10.8 validates directly.

4 The Rules

The Rules do not steer. They classify an encounter, name the obligation, remove the options they forbid and order the rest. The steering is done by the filter and the core.

4.1 Classification

EncounterClassifier deems risk of collision to exist when all of the following hold: range inside **EncounterRangeMeters** (4000 m); TCPA positive; and CPA inside the larger of **RiskClosestApproachMeters** (400 m) and this contact’s separation floor. The last clause repaired a fault: the Rules recognised risk at 400 m while the filter refused to close inside 446, so there was a band in which the ship manoeuvred hard for a vessel the Rules had just declared no risk, with no obligation and no prohibition on port. An encounter that can cause a manoeuvre must be one the Rules have classified.

Both vessels must be making way (**MakingWaySpeed**, half a knot) before any steering rule divides responsibilities. A stopped vessel is water to be kept clear of, which the floor does with no help from the Rules; classifying her as a crossing stand-on by the sign of a bearing produced the permanent, seven-minutes-early Rule 17(b) lapse recorded in Section 11.

With risk established, Rule 13 is tested first because it overrides the crossing and head-on rules:




Figure 2: Risk of collision. The contact is closing on a relative track whose closest point of approach falls inside the larger of the risk distance and her floor; both vessels are making way; she is inside the encounter range. Only then is a rule asked which one governs.

Test	Result
own ship abaft her beam by more than OvertakingSectorDegrees (67.5°)	Overtaking: we keep clear
she is abaft our beam by the same	Being overtaken: we stand on
headings reciprocal within 25° and she within 12° of dead ahead, or within the 30° doubt band of Rule 14(c)	Head-on: both alter to starboard
she is on our starboard bow	Crossing, give way
she is on our port bow	Crossing, stand on
otherwise	None: the separation floor alone

Table 1: Classification order. The role geometry is read from the *intended course*, the current route leg, rather than from wherever the bow points while working up; a crosser converging on the passage is a crossing, not an overtaker of a ship lying at anchor across her own route.

The doubt band is Rule 14(c) made arithmetic: a near-reciprocal contact fifteen degrees on the port bow used to fall to crossing stand-on, and by symmetry she classified us the same, so both held until their clocks lapsed and both then acted in a compressed window. Where there is doubt whether a head-on situation exists, both assume it does. The width of the band was chosen rather than measured, and Section 10.5 records the attempt to measure it.

4.2 Memory: the encounter register

Classification used to be re-derived from scratch every step with no persistence, so a single frame of adverse geometry could commit the helm, a classification flipped as a bearing drifted across a threshold, and Rule 13(d), “once overtaking, always overtaking”, was documented but could not be true of stateless code. **EncounterMemory** is the register: one entry per (own, contact) pair, created when risk is first suspected, carrying the classification, the chosen passing side and the manoeuvre in progress, and surviving until the contact is finally past and clear.




Figure 3: The classification sectors about own ship, read against her intended course. Rule 13 is tested first: a vessel in the stern sector is overtaking and keeps clear. The head-on core is $\pm 12^\circ$ with reciprocal headings within 25° ; the Rule 14(c) doubt band widens it to $\pm 30^\circ$. Elsewhere, starboard bow is crossing give-way and port bow is crossing stand-on.





Figure 4: Two gates on classification. (a) A vessel not making way is not a rules encounter. (b) At a cold start the role geometry is read from the route leg, not the bow.



Rule 13(d): a pinned *overtaking* stays one however the bearing later changes; the classifier is not asked again while the entry is open.

Figure 5: The encounter register, one entry per (own, contact) pair. Classification, role and passing side are decided once and carried; the classifier is not re-asked while the entry is open.




Figure 6: Rule 13(d) needs memory. A vessel that entered from the stern sector is the overtaking vessel until she is past and clear, whatever her bearing later becomes.

Two details are recorded because each was a live fault. A pinned role must still lapse when the geometry flips the roles mid-encounter: on 12 August a stand-on pin, when fresh geometry put own ship in the give-way chair, returned “stand on, never due” and disabled Rule 17(b) entirely; the ship stood on from thirty-four seconds to hull contact. A role-flipped pin now arms its own lapse inside a 120-second imminence. The distance at which an encounter is held open is recorded as a finding rather than set as a preference: **RiskReleaseFactor** is implemented and left at 1.0, because widening it to 1.10 or 1.25 grounded the ship in the confined-water set (Section 10.3). Four hundred metres is already most of the water a 152 m ship has in a channel.

4.3 Obligation

A give-way vessel sends her desired command to the filter, which substitutes if it must. A stand-on vessel holds course and speed, and the filter is called with the horizon cut to **StandOnHorizonSeconds** (90 s) rather than the full 360; a stand-on vessel altering at long range is the thing Rule 17 exists to prevent.

Rule 17’s lapse is solved, not clocked. Rule 17(a)(ii), *may* act, is read when the give-way vessel’s own best alteration can no longer make the distance and she is not observed to be turning: a vessel visibly swinging is taking action, however late she left it, and manoeuvring across that removes the predictability she is steering by. Rule 17(b), *must* act, is read when no alteration of ours alone still makes the distance. Her assumed rate of turn is half a degree per second at two hundred metres of hull, scaled by AIS length; the earlier constant of one and a half degrees credited a bulk carrier with three times the rate she has, and we stood on well past the point at which she could still have resolved the encounter.




Figure 7: The stand-on vessel’s timeline. Both lapses of Rule 17 are solved from the geometry by the timing oracle: 17(a)(ii) when the give-way vessel’s own best alteration can no longer make the distance and she is not seen to be turning; 17(b) when no alteration of ours alone can.

4.4 What is forbidden, and in extremis

A port alteration is struck out under Rule 14 (she is altering to starboard at the same moment), under Rules 15 and 16 (it crosses ahead), and under Rule 17(c). It is struck out of every decision path and against every vessel at risk, not only the one that forced the decision: a picture containing a head-on and an overtaking used to permit port whenever the overtaking was the urgent one. Rule 8(b), which widens a small alteration to 20° so another bridge can see it, asks which side it is widening before it does so, and widens only if the wider alteration is admissible in its own right.

Nothing is forbidden in extremis; that is Rule 2. In extremis is read against *collision*, the sum of the two hull radii, and not against the passing distance. Reading it against the passing distance switched the whole rules layer off whenever the comfort margin had been lost, routinely and often

deliberately by slowing down, and two of the twenty-two benchmark encounters ran that way from end to end.




Figure 8: Where Rule 2 begins. A candidate that ends between the passing distance and the collision distance is a lawful degradation; only a candidate that cannot stay outside hull-to-hull lifts the prohibitions.

4.5 Citation

RuleCitation derives the wording of every report from the encounter and the manoeuvre, so it cannot drift from the decision. Where nothing governs it says “no rule (separation floor only)” rather than inventing an authority, which is what most of the unexplained manoeuvres in the early logs proved to be. Rule 8(e) is cited when way is taken off, 8(b) when an alteration was widened, 17(b) when stand-on lapsed, and 2 in extremis.

Not modelled

Sound signals, restricted visibility (Rule 19; its design is set out in Section 4.6), vessels not under command, constrained by draught, fishing or sailing, and traffic separation schemes (Rule 10) are not modelled. Every contact is treated as a power-driven vessel underway, in sight. Navigation marks enter the picture as dangers (Section 6.3), not as Rule 9 obligations.

4.6 Restricted visibility: Rule 19, and what it would change

Everything above assumes the vessels are in sight of one another. That is the condition of Section II of the Steering and Sailing Rules; Rules 11 to 18 apply only then. In fog, heavy rain or snow the governing rule is Rule 19, which is a different scheme from the crossing and overtaking rules rather than a variant of them. It is not implemented. What follows is the design: what the rule says, what it does to each layer of Sections 3 to 5 and 7, and what would test it.

What the rule says. A vessel in or near an area of restricted visibility proceeds at a safe speed adapted to the conditions, with her engines ready for immediate manoeuvre (19(b)). A vessel which detects by radar alone the presence of another vessel determines whether a close-quarters situation is developing or a risk of collision exists, and if so takes avoiding action in ample time; when that action is an alteration of course she avoids, so far as possible, an alteration to port for a vessel forward of the beam other than one being overtaken, and an alteration toward a vessel abeam or abaft the beam (19(d)). Except where it has been determined that no risk of collision exists, a vessel which hears the fog signal of another apparently forward of her beam, or which cannot avoid a close-quarters situation with a vessel forward of her beam, reduces her speed to the minimum at which she can be kept on her course, takes all her way off if necessary, and navigates with extreme caution until the danger is over (19(e)).




Figure 9: Rule 19(d), one panel per sector, every manoeuvre made going ahead. There is no give-way and no stand-on vessel: each vessel acts on her own plot, in ample time, and where a close-quarters situation forward of the beam cannot be avoided she comes down to steerage way (19(e)).

Table 2 takes the model of this paper layer by layer and says what the rule does to each.

Implications.

- **The Section II machinery is replaced, not attenuated.** Rule 17 and its lapses, the give-way and stand-on roles, the head-on doubt band and the crossing tests have no meaning when the vessels are not in sight. A visibility state must select a rule table; it must not be a penalty added to the existing one.
- **Rule 19(d)(ii) is a prohibition the model already prices.** TurnTowardPenalty charges a turn toward a closing contact inside 1200 m. In restricted visibility, for a contact abeam or abaft the beam, it becomes hard, exactly as port did under Rules 14, 15 and 17(c) in Section 10.2; the lesson of that section, that a prohibition priced as a penalty is eventually bought, applies unchanged.
- **The overtaken exception survives.** 19(d)(i) allows port for a vessel being overtaken, which is the same geometry Rule 13(d) pins today. The register’s “entered from the stern sector” fact is what decides it, so the memory of Section 4.2 is needed in fog as much as in sight.
- **Speed is the rule’s own remedy.** 19(e) is the one place in the Rules that orders way taken off. The core already has speed as a candidate and the escorts already defend by speed first; what is missing is a *steerage way* fraction and a tier that prefers it when the forward sector cannot be cleared.

layer	in sight (today)	restricted visibility (Rule 19)
trigger	always	a <i>visibility state</i> , set by the operator or a sensor, per area; “in or near” an area of restricted visibility
classification	head-on, crossing give-way, crossing stand-on, overtaking, being overtaken	<i>aspect</i> only: forward of the beam, abeam or abaft; overtaken or not
obligation	one vessel gives way, the other stands on (Rule 17) until her lapses fall due	<i>both act</i> ; no stand-on, no lapse, no “may act” and “must act”; the register keeps no role pin, only the aspect and the passing side
threshold	risk of collision: CPA inside the floor	<i>close-quarters situation developing</i> : a larger distance, and the plot’s trend, not only its CPA
forbidden	port under Rules 14, 15, 17(c)	port for a contact forward of the beam unless she is being overtaken; any turn toward a contact abeam or abaft; both hard, not priced
preferred	starboard, then slowing (8(e))	starboard forward of the beam; <i>speed</i> is the rule’s own answer once a close-quarters situation forward of the beam cannot be avoided
speed set	1.0, 0.85, 0.6, 0.3, 0.0, 1.5	adds <i>steerage way</i> , the least speed that keeps her on her course, and makes it the 19(e) candidate; no increase
timing	the timing oracle against her assumed manoeuvre	the same oracle, but ample time measured against a longer passing distance and without crediting her with an obligation to act
the plan	the resolver eases passage speed for a meeting minutes ahead	becomes the 19(b) safe-speed ladder: the ease is the first answer, not the fallback
escorts	speed first, helm when no speed answers	unchanged in principle; the helm, when granted, obeys the two sector prohibitions
own units	no encounter under command	unchanged; the fleet knows its own units by link, and Rule 19 is about vessels that cannot see one another
reporting	the rule cited	“Rule 19(d)(i)”, “19(d)(ii)”, “19(e): way taken off”

Table 2: Rule 19 against each layer of the model. Nothing structural changes: the filter, the shared predictor, the register and the core are the same objects with a different rule table in front of them.

- **Close quarters is farther than a floor.** Risk of collision at 400 m is the in-sight threshold; a close-quarters situation in fog is judged at miles, and on the plot's trend. The floors of Section 3.2 would carry a visibility multiplier, and the watch of Section 5 would open earlier.
- **AIS is not sight.** A contact known from AIS is still a vessel not in sight; Rule 19 applies. The picture is built the same way; the rule table is what differs.
- **Sound signals are out of reach.** 19(e)'s fog signal apparently forward of the beam is an aural cue the platform has no sensor for. It would be an operator input.

What would test it. Two benchmarks already exist. The Imazu set of Section 10.2 run under the Rule 19 table, where every vessel acts and none stands on, gives a restricted-visibility baseline. Its expected signature, earlier and smaller alterations, more speed answers, and no port for a forward contact except the overtaken cases 3, 7, 15, 17, 20 and 22, is a prediction of this design and has not been measured. The second half of every ENMM paper of Section 10.8 is a Rule 19 manoeuvre with its answer worked by hand: the course to adopt at a stated time to pass at a stated distance from an echo, whether that passes ahead or astern of her, the time of return to the original course, and the echo's CPA and TCPA after the manoeuvre. Each is an alter-hold-rejoin candidate of Section 5 with its answer given, and the harness of Section 10.8 already holds the plot it starts from.

5 The decision: the manoeuvre core

5.1 A candidate is a manoeuvre

ManoeuvreCore chooses among candidates that are whole trajectories. A candidate is an alteration and a speed fraction:

$$\delta \in \{0, \pm 20, \pm 30, \pm 45, \pm 60, \pm 90\}^\circ, \quad s \in \{1.0, 0.85, 0.6, 0.3, 0.0, 1.5\}, \quad (3)$$

sixty-six in all. Each is run as alter, hold, rejoin: turn onto the offset course at the vehicle's own rate through Eq. (2), hold it at the fraction of the ordered speed until every contact is *past and clear* (range at or beyond floor plus margin, and opening), then steer back onto the corridor leg line with the track pilot's own course-to-steer law and a look-ahead of $\max(200 \text{ m}, 120 \text{ s} \times v)$. The trajectory is run out to a horizon of 900 s at ten-second steps and scored at every step against every contact, advanced by the shared predictor, and against the clearance field.





Figure 10: The sixty-six candidates for the crossing of Figure 12, coloured by fate. The colouring is illustrative, not the output of a benchmark run; a later revision will draw it from a recorded case. The port half is struck out by Rule 15; holding course and the smallest starboard alterations breach the floor; hard starboard at speed runs onto foul water. Among what remains the cheapest is taken.



At every step the position is checked against every contact (advanced by the shared predictor) and against the clearance field. A breach anywhere, including on the rejoin, strikes the candidate out.

Figure 11: One candidate in time. Alter, hold, rejoin, then on track to the fifteen-minute horizon; the hold ends on a condition, not on a clock.




Figure 12: Three of sixty-six candidates. Floors and foul water strike a candidate out; the Rules strike out port; what remains is priced.

5.2 Hard, forbidden, priced

A candidate that projects inside any floor plus **MarginMeters**, or whose clearance falls below the terrain margin at any step, including on the rejoin, is inadmissible and no cost brings it back. A port-helm candidate is removed when the Rules forbid port against any vessel at risk (Section 4). What remains is priced:

$$\begin{aligned}
 J = & |\delta| + |\Delta v| c_{\text{speed}} \max(1, r) \left[4 \text{ if increasing} \right] + p_{\text{port}} \left[\delta < -1^\circ \right] \\
 & + (c_{\text{commit}} + p_{\text{ahead}}) \left[\text{reverses the committed side} \right] + \frac{\text{drift}}{100 \text{ m}} + c_{\text{caution}} d_{\text{caution}} \\
 & + c_{\text{corridor}} \frac{\text{excess cross-track}}{\text{half-width}} + c_{\text{ahead}}(\text{held leg}), \tag{4}
 \end{aligned}$$

with $c_{\text{speed}} = 3$, the reluctance $r = 6$ for a ship and 1 for a boat, $p_{\text{port}} = 15$ where no rule governs, $c_{\text{commit}} = 55^\circ$, $p_{\text{ahead}} = 600$, $c_{\text{caution}} = 45$ per unit of depth into the 250 m caution band, and $c_{\text{corridor}} = 40$, uncapped. The unit is degrees of helm; every other term is expressed in it so that helm and speed answers are comparable. Until course and speed were priced in one unit the search never selected a speed change.

Two weights need explanation. A speed increase is priced at four times a reduction. The candidate is needed: against a contact whose plot runs across the scope the helm does not answer and slowing is not enough (Pub. 1310, rule 5). At twice the price it was selected routinely and the benchmark lost four arrivals to the resulting detours. The crossing-ahead penalty is six hundred, raised from sixty, so that any candidate passing astern is cheaper than any candidate crossing ahead; the reversal charge carries the crossing penalty on top, so that escaping a bow cannot pay for a change of side.




Figure 13: The cost of four candidates, term by term. The values are illustrative, chosen to show the relative size of the terms for the crossing of Figure 12, and will be replaced by a recorded case in a later revision. The port candidate would be struck out under Rule 15; it is shown priced as it would be where no rule governs, with the reversal charge it would carry if the helm had already been committed to starboard.

5.3 Tiers

Solve answers in four tiers, taking the first that applies:

1. **Watch.** The held course is inadmissible but afloat, no breach falls inside the act window (`ActWithinSeconds`, 300s; 90 if a contact is turning), and clearance is at least `WatchClearanceMeters` (500m). Report *ample time*, hold, and show the operator the proposed manoeuvre and its deadline (Section 9).
2. **The cheapest admissible candidate.** Holding course, if admissible, is ample time; anything else is action due.
3. **Degrade.** Nothing makes the required passing distance: take the widest *lawful* opening, run in full. Standing on wins unless the widest opens more than a meaningful gain, 75m, and that bar applies only while holding course still makes the worst floor; failing it, any metre of gain is taken.
4. **Rule 2.** Even the widest lawful opening ends inside hull-to-hull: the port restriction lifts and the widest candidate is taken, flagged *no safe option*.

The first version of the degradation rule dropped to the collision distance as soon as the passing distance became unachievable; it failed seven benchmark cases before the graded version replaced it. The meaningful-gain bar, when first released without the “only while holding makes the floor” clause, stood a guide on into her own subordinate at ample time and stood eight of twenty-two Imazu cases into 812 to 885m passes for four days while the suite stayed green (Section 11).

5.4 Cost of the search

Candidates are ordered by what is knowable about them before they are run: the alteration, the speed change, the side, and whether they reverse something already begun. That partial price is a true lower bound, since every remaining term can only add, so the search stops as soon as the best complete answer so far is no dearer than the next candidate’s floor. It returns the same answer as running all sixty-six and in practice runs a few; each run also stops at its first refusal. Without the bound the core exceeded the time budget of the formation diagnostic.




Figure 14: The four tiers of Solve, taken in order.




Figure 15: Branch and bound over the candidates, with illustrative costs. The floor cost is known before a candidate is run and can only grow once it is; the search stops at the first candidate whose floor already exceeds the best complete cost.

5.5 What it replaced

The previous decision core is kept behind `UseManoeuvreCore` so that the two can be measured against each other on every baseline. It solved a tangent to one contact's guard circle (the two edges of the velocity obstacle), tested the tangents against the rest, and fell to a grid of headings in ten-degree steps crossed with four speeds when no tangent cleared the set. Its urgency was read against a fixed reference alteration of ninety degrees, so it held until nothing short of hard over would still make the distance and then altered fifty or sixty degrees; lowering the reference to 35° was tried and rejected by a regression test, because committing earlier meant committing smaller and a smaller alteration in a crossing no longer passed astern. The dilemma was an artefact of the candidate set: the missing candidate was the early, larger-than-minimal alteration constrained to pass astern, which is what Rule 8(a) together with Rule 15 describes. With whole manoeuvres in the set the reference disappears, and the passing side is a property of the candidate rather than a consequence of lateness. The differences are tabulated in Section 10.2: the same safety figures, a seventh of the reversals, and every case arriving.




Figure 16: Why the commit reference could not be lowered in the old core, and why the new one needs none. With only tangent candidates, acting early (a) means a gentle turn that crosses ahead, so the core waited until only a bold turn (b) remained. The missing candidate (c) is early, deliberately larger than minimal, and constrained to pass astern.

`EvasionTiming`, the solved last-point-of-alteration oracle, survives from the old core and is what the Rule 17 lapse and the watch deadline are read from.

6 The plan: route anticipation

Every projection in the pipeline extended the *current* command: the filter walks a candidate onto its heading and assumes it thereafter, and the classifier reads the geometry of the bow. None projected the *route*. A passage whose next leg turned to starboard across a slow crosser's water produced, over three runs in August 2026, a bow crossing, a corridor departure and a collision. Each of those manoeuvres was lawful at the moment it was made, because by then only late and large alterations remained. The requirement was to anticipate in general, not to repair the case.

6.1 The route run on paper

RoutePlot runs the remaining passage on paper with the hull's real capability: the same wheel-over law the track pilot steers by, the turning circle of Eq. (2), the ease into sharp corners made good by the wheel-over line, the eased speed held through the swing, and the braking curve into the destination. The plot is not a commitment; the guard and the Rules may override it. It is the statement of intent against which every anticipation question is asked.




Figure 17: The route run on paper. Wheel-over before each corner at the radius the turning law gives for the speed made good; the ease into a sharp corner made good by the wheel-over line and held through the swing; the braking curve into the destination.

RoutePlot.Conflicts scans the plot against every rules-apply contact, advanced by the same ContactMotion every other layer uses, a turning contact held to the worse of arc and her steadied course. It reports each meeting inside the floor: when, where, how close, and on which side of her beam own ship would pass. That is the crossing-ahead question, answered for the plan, minutes before the filter could ask it.





Figure 18: The plot scanned against a contact. Each meeting inside the floor is reported with its time, place, closest approach and the side of her beam own ship would pass, which is the crossing-ahead question answered for the plan.

6.2 The resolver

The leader, every few seconds, works through three steps:


1. Nothing in the plan's way: carry on, and release any held ease once the full-speed plot has read clear for a dwell.
2. A meeting minutes ahead: the smallest ease of passage speed (85, 70 or 50 %) whose re-run plot clears everything. This generalises the operator's instruction that a small reduction of speed, made while there is ample time, resolves most meetings; the ease is announced once, with the meeting it resolves.
3. No ease clears it: reported once. The helm answers through the Rules when the meeting comes within the filter's horizon, and it arrives correctly classified, because the classifier reads the intended course.



The plot is judged against the *seaway* clearance for the hull, not the live floors. The live floor is the filter's standard; against reach-inflated floors no route through a working sound ever read clear, and the resolver thrashed.

Figure 19: The resolver's ladder.

The plan is judged against the *seaway* clearance for the hull, not against the live floors. The first run of the campaign judged it against live floors inflated to a kilometre with the screen deployed, against which no route through a working sound ever reads clear; the resolver thrashed between fifty-per-cent eases and "no ease clears" and every passage timed out. A long-horizon question has to be asked against a long-horizon standard; the live floor is the filter's standard, not the plan's.



Judged against the live floor, the route through the sound never reads clear: every lane vessel is a meeting. Judged against seaway clearance it is a passage with two crossings, each resolvable by a small ease.

Figure 20: Two standards for one question. The plan is judged against the seaway clearance for the hull; the filter enforces the live floors on the command.

6.3 Routing and the marks

The route planner is not a contribution of this work. Passages and station routes are found by Theta*, a textbook any-angle grid search [18], used as published. This project's contribution is the grid it searches: the clearance field of Section 3.5, eroded by the berth of the hull class asking, with charted dangers, wrecks and obstructions merged into it as foul discs, and each class of hull given its own field so that a boat is not routed round a shoal she clears by four metres. In the Solent's hundreds of danger marks the navigability test was an $O(\text{marks})$ scan per cell and a route took twenty seconds on the main thread; a uniform-grid spatial index over the marks brought it to milliseconds.

Navigation marks are drawn, not yet interpreted

The marks are read from real S-101 cells, classified as IALA buoyage (lateral with its hand, cardinal with its quadrant, isolated danger, safe water, special) and drawn from the IHO portrayal catalogue [20]. Their *navigational meaning* is not yet used by the planner: which hand to leave a lateral mark on, the quadrant of safe water a cardinal mark names and the sector it closes, a channel as the water between two lines of lateral marks, the conventional direction of buoyage, and the narrow-channel obligations of Rule 9. Today the planner reads depth and foul discs only. The passage ledger counts, after the fact, how often a unit passed a lateral mark on the wrong hand, judged against the direction of the leg because the conventional direction is not in the schema; that count is a measurement of the gap, not a control. The interpretation is on the roadmap (Section 13). Where an operating area has no vector cell the buoyage layer draws nothing and reports that it has nothing to draw; hand-typed positions would put invented navigation data on an operator's screen.




Figure 21: Routing over the clearance field. The field is eroded by the berth of the hull class and charted dangers are merged into it as foul discs; the standard any-angle search then returns a few straight legs through the water that remains. The lateral pair and the cardinal mark are drawn as the chart draws them, and the planner does not yet read them.

7 The fleet

7.1 The doctrine

One decision, taken by the guide, with the stations carried along: that is the arrangement of an escort screen, and it is the arrangement here. The guide holds her course and speed for her screen; her boats conform to her. The boats keep their stations, and when traffic makes a station untenable they break station and give way for themselves, then regain. Between units under command there is no encounter: no floor, no helm, no hold. Their mutual positioning is the formation plan's business, and the plan is coordinated over a link that reactive avoidance does not have.

Three faults followed from departures from this arrangement, each found in a live run. The guide answered her own screen (ninety of ninety-four helm engagements). Deploying the whole screen in one frame put four boats inside each other's floors, so that every motion candidate of every boat breached a neighbour and each boat's least-bad answer was to stop, in extremis against her own squadron, for the rest of the run; the remedy is sequential launch, the next boat going when the last is clear and under way. In the Solent the screen stopped to avoid itself ("DELTA starboard against USV CHARLIE", "BRAVO port against AURORA"), which is what removed own units from each other's guard sets altogether.

7.2 The escorts' defence

An escort under command keeps her station by course and speed and defends herself against traffic by speed first: `HoldOffTraffic` asks, per contact and with terrain excluded, whether any speed in $[0, v]$ clears her over a ninety-second horizon, and holds the largest that does. When no speed answers, the helm is granted and the boat gives way under the Rules, passing astern, then presses on to her station. The hold-up ledger names the contact for every hold, and the log reads:

- 1 USV DELTA has been held 60 s by MMSI235026142, 592 m from her station; waiting for the traffic to clear.
- 2 USV BRAVO holding speed for MMSI235026142: 16 kn ordered, 8 kn held.
- 3 USV DELTA has been held long enough: taking her helm to pass astern of MMSI235026142 and press on to her station.

The per-contact question must not be answered by the coastline. In September 2026 the governor asked it through the terrain-inclusive projection, so a boat within her terrain horizon of a bank found no admissible speed against every contact in turn, and the reduction was charged to




Figure 22: The doctrine. The guide holds her course for her screen; a boat whose station a merchant makes untenable breaks station, gives way astern of her and regains. Between units under command there is no floor.




Figure 23: The launch raft, and its cure.

whichever contact held the slot: units slowed “for” ships fifteen kilometres away whose ninety-second CPA was also fifteen kilometres, while their terrain clearance sat at 75 to 150 m. Foul ground is now handled separately, reported as *easing for foul ground* with no contact named, and does not fall through to a Rule 2 helm engagement, because easing for a bank is not an encounter.




Figure 24: The escort speed governor. (a) The per-contact question and the held speed. (b) Why the coastline must never answer it.

7.3 The boats are not the ship

Stated by the operator across three operating areas

The mothership is a 152 m hull needing six metres of water. The USVs draw half a metre and are small, fast and nimble. Their depth, foul-ground and anti-collision rules must not be the same.

The boats had been held to the ship’s parameters throughout, which was the cause of most of the operator’s reports of units stopping without visible reason. A USV station in five metres of water read as on the ground because the depth veto used the ship’s six against a half-metre boat; half the Solent read as foul under them; a boat detoured around shoals she cleared by four metres; and a USV held for distant or opening traffic because she inherited the ship’s 370 m standoff. Since September 2026:

- each unit has her own draught (0.5 m for a USV, 6 m for the ship) and her own clearance field, and grounding is scored on her own draught;
- a USV plans her station route on her own field with a 15 m berth, and eases for foul ground with an 8 m terrain margin rather than the ship’s 40;
- the leader’s traffic floor is her own, base clearance plus the larger hull, with no formation reach.

Two of the operator’s standing instructions remain to be enforced in code and are on the roadmap: a USV handles targets within about a kilometre, reacts late, and never holds for an opening contact with a safe CPA; and any vessel can go astern, which station keeping honours but the manoeuvre core’s speed set (no negative fraction) does not.




Figure 25: The boats are not the ship. (a) Per-class water. (b) The leader avoids traffic as one ship.

7.4 Formation shapes

The formation layer assigns stations by optimal assignment with hysteresis, relocates through transitions between shapes (line ahead, line abreast, echelon, vee, wedge, diamond, circular screen, trail, picket line), and was itself swept before the anti-collision rebuild: ninety-six transitions per run, sixteen runs, with ring incursions, bow crossings and leash exceedances scored (Section 10.7). The general lesson from that sweep was that a clause written for the leader's problem must be scoped to the leader; applied to the whole fleet it removes an arrangement the fleet depends on.




Figure 26: The ten formation shapes cycled by the stress harness; the leader at the centre, four boats on station.

8 The chart: S-100 data and ECDIS conformance

Everything in Sections 3 and 6 that concerns the water, the clearance field, the per-class contours, the corridor, the marks merged into routing, is built from a chart. The chart dataset behind the Solent and Saint-Malo operating areas was adopted for two reasons, which are different kinds of claim and are stated separately.

Two purposes, one dataset

Implementation and test of the S-100 series. The IHO’s S-100 framework went operational in January 2026 with its Phase 1 products (S-101 ENC, S-102 bathymetric surface, S-104 water levels, S-111 surface currents, S-124, S-128, S-129), with the IMO transition running to 1 January 2029. The platform ingests the four data products directly from the official UKHO and SHOM trial datasets, and every one of them does a job in the collision model, not on the display alone.

ECDIS conformance. The chart presentation was audited requirement by requirement against S-52 Edition 6.1.1 and the Presentation Library 4.0.3 Part I Addendum, and the audit’s delta is published with named departures. An operator who has stood a watch on a type-approved ECDIS should find the chart familiar and should be told, on the chart itself, where it departs from the standard.

8.1 The dataset

The UKHO and SHOM joint sea-trials programme of 2026 is, to the authors’ knowledge, the most complete S-100 test campaign published, and it runs down the western Channel: Southampton and Portsmouth, across to Jersey, down to Saint-Malo. That corridor is where the data is, and it is why the two S-100 operating areas are the Solent and Saint-Malo rather than Dover. Both agencies license the data for evaluation and exploration only, not for navigation and not commercially. A simulator trialling the standard falls within that use, and the licence is one reason the chart states, in its legend, that it is not a type-approved ECDIS.

product	Solent (UKHO)	Saint-Malo (SHOM)
S-101 ENC	19 cells, overview to berthing scale (Cowes, Southampton docks); 136 depth and land areas	11 cells including the Channel Islands approaches; 1022 depth and land areas; the harbour cell alone carries 5055 features
S-102 bathymetry	12 cells; 161959 grid cells shoaled against the open base; high-resolution patch 15003×5276 at the survey’s own 0.00004°	Saint-Malo approaches, v3.0.0; 7963 cells shoaled against the open base; per-node uncertainty carried
S-104 water levels	3 cells, 40 % coverage, 11 to 17 May 2026	Norman-Breton Gulf and Cherbourg, 1 June to 31 August 2026
S-111 surface currents	cells present, values uniformly zero in the trial window (checked in the raw HDF5)	Norman-Breton Gulf, 1 June to 31 August 2026; 5 kn streams round the Minquiers

Table 3: The S-100 Phase 1 products behind the two Channel operating areas, as ingested. Gaps in the trial data are reported as found.

8.2 What each product does in the model





Figure 27: The S-100 products and what consumes them. Each product reaches both the collision model and the display through one reader, so the chart the operator sees and the water the filter enforces are the same data.

S-101 is read from the ISO 8211 container by a reader driven by the cell's own data descriptive record, decoding fields and subfields against the format control strings rather than against assumptions about the specification. Feature records resolve their spatial pointers through the point, curve, composite-curve and surface tables into real geometry, ring by ring, hole by hole. Depth and land areas are triangulated and drawn as the vector chart, coloured by their own depth ranges; depth contours and the coastline are drawn as the line features the cell carries, not as polygon boundaries. Navigational marks are classified as real IALA buoyage and lights from their S-101 attributes and drawn from the complete S-101 portrayal catalogue, all 725 symbols extracted, rather than from borrowed artwork. Charted dangers are merged into the clearance field. The marks' navigational meaning, which hand to leave a lateral mark on, the quadrant a cardinal mark closes, is parsed and drawn but not yet read by the planner (Section 6.3). An operator can click a mark and read its attributes.

S-102 is converted into the packed depth grid that the terrain veto is built on, merged pessimistically with the open base so that the shoalest source wins, and where the survey carries per-node uncertainty the veto margin inside survey coverage is the surveyor's declared doubt. The same grid paints the chart fill at the survey's native resolution where no vector cell reaches; where neither exists the chart shows nothing rather than an interpolated guess.

The order of loading matters: a harbour-scale cell must win outright over the general-scale cell beneath it, so each cell of a portfolio is lifted a fraction above the coarser one. Reading and triangulating on the main thread froze the application for twenty-two minutes on one cell whose depth area merges fifty-nine rings; on a background thread with a complexity cap it takes twelve seconds. Three tooling faults surfaced only when both agencies' cells sat in one directory, because the two agencies store the same HDF5 attributes in different shapes and name their files by different conventions, and one of them produced a 17 GB grid for a 150 MB survey before cells outside the area were excluded.

The audit graded the presentation against the two documents an ECDIS is graded against, with the requirements extracted into two checklists and each given an identifier. Three deviations are declared up front and kept: the chart is a fadeable wash over imagery and terrain, which S-52 knows nothing of; own ship, routes, intent and guard rings use the platform's tactical symbology rather than the S-52 own-ship figure; and the chart fill is painted only from genuine survey product, never from the open, interpolated base.



October 2026

What is partial or absent is named rather than hidden: sounding digits as placed labels rather than glyph compositions; no automatic isolated-danger overlay derived from depth against the safety contour; no IMO display-base, standard and other layering, the platform's own per-layer panel standing in; sector lights, scale minimum and overscale not yet drawn; the mariner's chartwork and pick reports; and the conditional-symbology engine of the Presentation Library, whose ISO 8211 foundation exists and whose lookup tables and procedures are the long arc still open. Radar overlay, the physical chart area and the alarm-panel formalisms of a type-approved unit are not applicable to a ground station and are listed as such.

What conformance means here, and what it does not

RTGroundControl is a ground-control station that draws an ECDIS-style chart. It is not a type-approved ECDIS, it is not seeking IEC 61174 approval, and it has not been through IHO conformance testing; the legend on the chart says so. The claim this paper makes is narrower and checkable: the presentation was measured against every requirement of S-52 6.1.1 and the Presentation Library addendum, the result is published with each requirement's state, and the data under it is the IHO's own Phase 1 product, ingested by readers written against the standard and checked against the official trial cells.

8.4 Why it matters for collision avoidance

The chart and the collision model share their data. The safety contour the operator sets is the contour the terrain veto enforces; the depth areas the vector chart draws are the areas the clearance field is cut from; the danger a mark names is a foul disc in routing; the tide that S-104 reports is the water under the keel in the grounding score. A conformant chart therefore gives the operator what the collision model cannot give on its own: a picture, in the symbology they were trained on, of the water the system is refusing to enter. The S-100 pipeline gives the model survey-resolution depth with declared uncertainty, which is the quantity a margin should be sized on. Cockburn Sound, the first operating area, still runs on the open base with no licensed product, and each result in Section 10 states which area it was run in.

9 What the operator sees

The system is designed for one operator. That person is responsible for the mothership, the four boats and the two aircraft at once, so the anti-collision layer may not demand attention in proportion to the number of units: it reports once per event, names the rule and the contact, asks for consent only where there is time to give it, and draws its intent on the same chart the operator is already watching. Everything below follows from that constraint.

9.1 The helm report

When the filter takes or releases the helm the event is reported once, edge-triggered:

```
1 AURORA: filter has the helm -- starboard against STIRLING TENDER, 60 deg off ordered course.
2 Rule 15/16 (crossing; we give way). On the ordered course she passes 22 m in 198 s;
3 this manoeuvre makes it 664 m.
```

Each of the four parts was missing at some point and each absence caused a misdiagnosis: the rule, derived so it cannot drift from the decision; the contact the Rules named, or failing that the *binding* contact, the one with least room; what the ordered course was heading for, which is why the helm went over; and what the manoeuvre buys. When only the third was reported, "CPA 22 m" was read as the filter accepting 22 m, when 22 m was the distance being avoided.

Events reach the operator through **StationEvents**: green for routine, amber for warning, red for alarm. The helm being taken is always red, because the operator is being told that something

other than their order is steering the ship. An ease held for minutes is a standing condition with its own row, not a thirty-second event, and every held deviation, rally, detour, aground, station on the ground, holds a row, newest on top.

9.2 Helm by consent

For the group leader, while action is not yet forced, the core's chosen manoeuvre is presented as a *proposal*. It is shown at the bottom centre of the display in an orange frame, with its options priced:

```
1 AURORA PROPOSES starboard 30 deg, 2 kn off for SEA URCHIN -- makes 640 m;
2 standing on makes 210 m (held off by CAPE LEEUWIN). Approve (Y) or veto (N -- she stands on).
```

The window is between five and twenty-five seconds and never later than sixty seconds before TCPA. An approval executes the manoeuvre; a veto stands for three minutes and the ship stands on, with what blocks standing on named. At execution the proposal is revalidated: if the subject or the manoeuvre has materially changed it is withdrawn and asked again, and only an explicit approval rides through the change. The proposed track is drawn on the chart out to its release and the regain to the leg line, beside the green path a veto keeps, so that the operator compares two projected tracks. Once action is due the proposal becomes a demand with a countdown.

9.3 Intent on the chart

The audit found the intent display drawing a manoeuvre the model had not decided: the applied command held for a display constant of thirty seconds, then a fabricated pursuit back to the desired heading, re-anchored to the ship every frame so the deviation could never appear, ending before the closest point of approach, and drawn only after the helm had already gone over. Since the core decides a whole manoeuvre, the drawn line is the chosen trajectory: the avoidance leg to the predicted closest approach, marked and labelled with the distance and the contact, the rejoin the core actually planned, and a bar showing the separation gained. During the watch the planned manoeuvre is dashed with its deadline counting down. Escort speed holds have a mark of their own, so a boat visibly slowing for a merchant is distinguishable from a boat malfunctioning.

10 Verification

10.1 From arguments to experiments

Before August 2026 questions about the model were argued from first principles and settled by observing runs against live Fremantle AIS, which differed from run to run. The audit's first recommendation was to make results measurable before touching the model, and it was followed: a deterministic encounter harness in the Core assembly, no Unity and no live traffic, scripted contacts at a fixed timestep, own fleet under the real pipeline. Every run is scored per encounter in the manner of Woerner [9]: the closest approach attained against the distance required; whether the helm went to port while a rule forbade it and Rule 2 did not permit it; time to first action; the number of starboard and port alterations and of reversals; steps in extremis; extra distance; and whether the waypoint was reached in the time allowed (Section C).

The suite has 280 tests. The three baselines below are regenerated by it and committed to the repository. A change to the model is measured against them, and both safety counts, no collisions and every case making its distance, are asserted by the tests rather than only written to the baseline document (Section 11 records why that distinction matters). The previous decision core, called the *incumbent* in the tables, runs the same sets beside the manoeuvre core.

10.2 Open water: the Imazu set

The twenty-two encounters of Imazu [11], in the form tabulated by Sawada et al. [10] (their Appendix 1, Table 4; the thesis itself, in Japanese, was not consulted and the set is cited through that table): own ship at twelve knots from six miles out, every target set to arrive at the origin together with a thirty-minute time to collision, targets holding course and speed throughout, cases 1 to 4 with one target, 5 to 12 with two, 13 to 22 with three. Required passing distance half a mile (926 m). Section B lists the geometries.

	incumbent, first table	incumbent, after fixes	manoeuvre core
made the required distance	22 of 22	22 of 22	22 of 22
collisions	0	0	0
worst closest approach	944 m	944 m	944 m
port under a rule forbidding it	7 of 22	0 of 22	0 of 22
reversals, all cases	66	7	10
arrived within the time allowed	19 of 22	20 of 22	21 of 22

Table 4: The Imazu set. The first column is the incumbent core as first published in August 2026, the second the same core after the compliance corrections of Section 4, the third the manoeuvre core.

case	ships	CPA (m)	first action (s)	stbd	port	reversals	extra (m)	arrived
1	1	957	1422	30	0	0	0	yes
2	1	949	1390	45	0	0	0	yes
3	1	969	990	0	90	0	0	yes
4	1	1502	1298	60	0	0	5354	no
5	2	950	1390	30	0	0	0	yes
6	2	946	612	45	0	0	0	yes
7	2	983	990	90	45	8	10	yes
8	2	950	1390	30	0	0	0	yes
9	2	948	1200	0	0	0	0	yes
10	2	944	906	20	0	0	0	yes
11	2	948	1200	0	0	0	0	yes
12	2	946	612	45	0	0	0	yes
13	3	948	612	30	0	0	0	yes
14	3	951	612	20	0	0	0	yes
15	3	962	990	45	0	0	0	yes
16	3	949	1320	0	0	0	0	yes
17	3	989	612	90	0	0	0	yes
18	3	947	906	20	0	0	0	yes
19	3	946	906	20	0	0	0	yes
20	3	962	906	90	90	2	0	yes
21	3	944	906	20	0	0	0	yes
22	3	962	990	45	0	0	0	yes

Table 5: The manoeuvre core on the Imazu set, per case. Port entries are lawful (Rule 13 expresses no side, or Rule 2). Case 4 arrives late because the core detours; case 7 reverses eight times and is the open legibility item. Every case made the required half mile.

Four faults were present when compliance stood at seven violations in twenty-two, and only one of them was in the prohibition itself. The measure charged deviation from the ordered course rather than the helm, so a vessel well off her track with the waypoint astern recorded a 178° “port alteration” while steering steadily to starboard. Rule 2 was read against the passing distance. The prohibition lived in one path and was asked of one ship. Rule 8(b) widened onto the forbidden

side. The incumbent's reversals fell from sixty-six to seven on the same fixes; the core's ten are concentrated in one case.

10.3 Confined water

Five encounters on a synthetic coastline, own ship the mothership running a four-mile leg, required passing distance 400 m against traffic, the corridor the strip a route eroded by 150 m guarantees with 40 m held in hand.

- 1 **head-on, bank to starboard.** With the lawful side foul, does she take the ground or the wrong side of Rule 14?
- 2 **crossing, starboard water runs out.** Is the give-way alteration made inside the water that exists?
- 3 **rejoin across a shoal.** Is the return from a lawful avoidance checked at all?
- 4 **own boat closing, bank close aboard.** With the helm withheld, is the throttle enough?
- 5 **narrows, meeting end on.** When there is no way out, is the fallback still afloat, and is it reported?

Table 6: The confined-water cases.

	incumbent	manoeuvre core
stayed afloat	5 of 5	5 of 5
kept the distance owed to traffic	4 of 5	4 of 5
least water anywhere	117 m	118 m
seconds outside the corridor	4348	3100
steps with no admissible command	67	0
arrived	4 of 5	5 of 5

Table 7: Confined water, both cores.

Case 1 remains open: head-on with the bank to starboard, the lawful side is the foul side, she passes at 340 to 354 m against 400 owed, and both cores behave the same way. This set is also what settled the question of risk hysteresis: widening the release distance to 1.10 or 1.25 took the set from five afloat to four and from four arriving to two.

10.4 The fleet's own encounters

Three encounters containing what the Imazu set does not, own units. Strangers hold course and speed; own units hold their bearing and range on own ship, which swings their heading through the compass and is why a straight-line projection of them is wrong. Required 400 m against traffic; against one of ours, the physical minimum of 101 m and nothing more.

10.5 The head-on doubt band

Rule 14(c)'s doubt band was chosen at 30° rather than measured, so it was swept: the Imazu set at each width from 10 to 35° under the manoeuvre core. The result is the same at every width: 22 of 22 safe, worst CPA 944 m, ten reversals, 21 of 22 arriving. The set places each target either reciprocal or well round on the bow, so no case lies in the region where the band decides which rule governs. A benchmark that cannot tell two settings apart is not evidence that they are the same. The figure stays at thirty, the wider reading of the Rule, and the symmetry condition, that a head-on classification is reciprocal at every width, is asserted rather than tabulated, because two bridges working to different rules are each behaving correctly by their own.

	encounter	CPA traffic	CPA own	shaped by a subordinate	arrived
1	guide, screen and a crossing merchant	423 m	520 m	0 steps	yes
2	escort screening, merchant on her starboard bow	423 m	284 m	0	yes
3	three ships, no pairwise answer	432 m	—	0	yes

Table 8: The fleet set, identical under both cores. The middle column is the one the set exists for: a manoeuvre taken because of one of our own boats is a fault whatever it did for the separation, and it stayed a fault through three earlier fixes, each of which closed one channel by which an own unit reached her decision and left the others open.

10.6 Passages: the campaigns

The deterministic sets say what the model does to a geometry. The passage exercise says what the fleet does to a day: boats deployed, synthetic traffic alive, a fixed seed per run so runs compare, time accelerated four to six times, and per run a ledger of arrivals, timeouts, off-track fraction, groundings scored per unit's own draught, station keeping, helm engagements and releases, own-unit manoeuvres, escort holds, machinery movements, no-safe-option verdicts and the closest approach to traffic (Section D). Between runs the method was fixed: read the ledger and the log, correct the largest fault, run again.

Cockburn Sound, 11 to 12 August 2026. Five builds, four twenty-destination runs at four times real time. The faults corrected were event storms: a helm engaged and released on the same condition every frame. Two remedies generalised: engage and release on different conditions (latches, dwells and cooldowns), and credit a projection only as far as its source justifies. The fourth run, complete and instrumented: thirty helm engagements in two and a half simulated hours, all released; minimum approach to traffic 301 m; 91.9 % of samples on station with a mean error of 35 m; fifteen machinery movements. Three of twenty passages arrived in time, because the allowance demanded thirteen knots of a twelve-knot ship.

Cockburn Sound, 12 to 14 August 2026: route anticipation. Sixteen runs and about fifteen root causes, of which Section 6 describes the anti-collision half. The navigation half followed the same method. The leader did not wait for her screen, so traffic holds left a boat twenty-seven kilometres astern; the off-plan allowance then inflated the floors until a vessel passing 1692 m clear was classified as an encounter, and the resulting manoeuvre put the ship in four metres of water. Four changes followed: the ship rallies when her worst station error exceeds 700 m; the allowance is capped at 400 m; a formation forms up before it proceeds; boats are launched one at a time. Run 15 recorded the first arrivals of the campaign, five of eight, with a mean station error of 87 m, 76.6 % of samples on station, a minimum approach of 838 m, no collisions and no grounding. Seed 7, whose destinations run along the traffic lanes, exposed the limit of the escort doctrine: 534 escort hold-and-resume pairs, each correct on its own, and no passage completed in the time allowed. No collision occurred. That result is the first item of the roadmap.

A note on method

Each root cause in the campaign was found with an instrument added for the purpose, not by reasoning from the symptom. A once-a-minute station snapshot found three boats stationary, one of which had not left the anchorage. A two-hertz track log identified both causes within minutes, after several hours of incorrect inference. Twice a plausible correction addressed the wrong mechanism because it was designed before the instrument existed. The harness itself also penalised the fleet for its own geometry three times: a straight-line time allowance against island-rounding routes, an arrival radius smaller than the station-holding envelope, and synthetic traffic replayed at four times its reported speed, which falsified every CPA in the area. A metric that deteriorates is checked before the model is.

Solent, 4 September 2026: screen trial. Twenty passages over the marked reference route, seed 11, the formation shape changed every passage so that the reshape is exercised over open water and shallows, with the per-class USV rules of Section 7.3 in force.

passages	20
arrived / timed out	16 / 4
off track	21.2 %
aground samples (per own draught)	0
shoalest water under any unit	3.5 m
helm engagements / releases	27 / 26
escort traffic holds	23
no safe option	11
closest approach to traffic	70 m
closest approach between own units	0.8 m (sim 11)

Table 9: The seed-11 screen trial ledger. Station figures are omitted: the reference passage is long and the mean station error (9 km) reflects boats lawfully held by traffic over a working strait rather than station keeping, which the campaign runs measured at tens of metres.




Figure 29: Seed 11, twenty simulations. Own units under command no longer avoid each other, so the inter-unit separation is what the formation plan alone delivers; sims 11 and 13 (line ahead, echelon right) closed to under ten metres and are the reason station keeping, not avoidance, is where the inter-unit guarantee now has to be made. Water is per the unit's own draught; nothing touched.

10.7 The formation sweep

Before the anti-collision rebuild, the formation layer was swept on its own: ninety-six shape transitions per run with no traffic, sixteen runs, scoring ring incursions, bow crossings, leash exceedances, unsettled transitions, settling time and error, closest pair and closest approach to the leader. Over the sweep the transitions that settled went from 82 to 92 of 96, ring incursions to nil, mean settled error from 38 to about 25 m, and the closest approach to the leader held above 190 m except in one deliberately tight run. Two runs of identical code differed by one ring incursion, which is why the passage ledgers carry a seed and why no single run is quoted as a result. Table 10 gives the per-run figures.

run	failures	ring	bow	leash	unsettled	settle (s)	error (m)	closest pair (m)	to leader (m)
1	19	1	11	11	14	107	37.8	50	191
2	19	0	8	12	14	105	37.6	50	210
3	22	0	8	15	12	110	37.8	50	222
4	25	1	9	15	12	109	37.9	51	195
5	20	0	10	10	4	109	26.2	56	228
6	20	0	8	12	4	106	25.2	53	229
7	19	0	7	12	5	106	26.2	50	215
8	20	0	8	13	62	81	30.7	50	227
9	22	0	9	14	6	109	26.2	50	226
10	16	0	5	11	3	96	22.2	56	215
11	20	0	8	12	2	101	25.2	56	210
12	20	0	8	12	3	102	24.9	52	227
13	18	0	7	12	2	109	24.9	52	227
14	14	6	0	8	24	100	23.9	50	119
15	18	0	8	10	15	103	25.4	55	223
16	12	0	4	8	18	96	25.5	50	227

Table 10: The formation transition sweep, per run: failures, ring incursions, bow crossings and leash exceedances out of 96 transitions; transitions unsettled at the end of the run; mean settling time and settled station error; closest approach between any two units and to the leader. Run 14 was a deliberately tight configuration. Source: docs/formation-runs-no-traffic.csv.




Figure 30: The formation transition sweep, sixteen runs of ninety-six transitions each, no traffic.

10.8 The radar plotting corpus

The relative-motion layer is validated against the corrected radar plotting exam papers of the École Nationale de la Marine Marchande, Marseille [19]. Each paper gives own ship's course and speed and two instants, plots four echoes on a relative-motion, north-up scope, and its corrigé gives, worked by hand, the CPA and TCPA of each echo, the target's surface course and speed, and the set and drift of the current from a fixed echo. Two checks fall out of every paper: the fixed object's apparent motion must be the reciprocal of the current, and every echo must agree on it.

Scope: the plot, not the manoeuvre

Each paper has a second half, a manoeuvre in restricted visibility. That is Rule 19, which this stack does not implement (Section 4), and it is not checked. What is checked is the plot: the arithmetic a watchkeeper does before any rule is applied.

The check runs against the Core sources as they are. The blank plot pages are vector PDF, so each echo's two fixes are read exactly from the drawing (the filled symbol and the digit beside it) rather than by eye; the corrigé values are transcribed from the handwritten answer pages. A small harness compiles the Core's relative-motion code from the source tree, derives each target's course and speed from the two fixes, solves the plot with the Core's own `RelativeMotion.Solve`, and compares (`Tools/radar-corpus`, results in `docs/trials/`).

quantity	agree	tolerance, and the misses
closest point of approach	18 of 18	0.2 NM; largest difference 0.13 NM
time to CPA	17 of 18	2 min; the miss is 16 min against 18 on a 23 kn target
target course	18 of 18	4°; largest difference 2°
target speed	16 of 18	0.6 kn; misses of 0.6 and 0.9 kn on 22 and 32 kn targets
current, set and drift	3 of 6	5° / 0.3 kn; within 3° and 0.4 kn where the current is 2 to 4 kn, 10 to 19° off where it is about 1 kn

Table 11: Six papers (April 2022 A, B and C; May 2023; January 2018; May 2019), eighteen target echoes and six fixed echoes, 30 September 2026. A current of one knot is the residual of two thirteen-knot vectors on a hand plot, and the disagreement there is within the plot's precision.

One definitional difference was found. Once TCPA is past, the Core reports the present range as the closest approach, deliberately: the historical figure would understate the separation a decision is being made at. The exam wants the geometric closest approach, and the harness recovers it from the Core's own direction and speed of relative motion. Both are reported; the geometric one is what the table scores.

The remaining papers of the corpus are the manoeuvre exercises and the ARPA and theory papers; the six above are the ones whose first question is the plot alone.

11 What the benchmarks did not catch

A published baseline of 22 of 22 safe and no unlawful port alterations coexisted, three times in one week, with a fault in a live run. Each is recorded here because the tables of Section 10 certify only the geometries they contain, and a reader should know what they did not. In what follows a vessel's *chair* is the role she occupies in an encounter: give-way or stand-on, overtaking or overtaken.

11 August: the overtaken chair

With every baseline green, the fleet turned to starboard across the bow of an overtaking vessel, in an encounter the classifier had correctly named Rule 13 the whole time. The manoeuvre core had shipped carrying the port prohibition and none of the other manoeuvre disciplines the older decision paths had (crossing-ahead cost, turn-toward threat), and no case in any suite put own ship in the overtaken vessel's chair and forced her to act: the Imazu set contains overtaking only with own ship as the overtaker. Pinned by `AnOvertakenVesselNeverCutsAcrossTheOvertaker`.

12 August: the role flip

The register's Rule 13(d) pin, when fresh geometry flipped own ship into the give-way chair, returned "stand on, never due", discarding the lapse flags as solved for the other chair and disabling Rule 17(b) entirely. The ship stood on from TCPA 34 seconds to hull contact, on a demonstration day. No benchmark exercised a transition between chairs mid-encounter. Fixed, and replayed verbatim in `APinnedStandOnStillLapsesWhenTheRolesFlipCloseAboard`.

17 August: a baseline that was not asserted

The meaningful-gain invariant, "a manoeuvre must buy something", stood a guide on into her own subordinate at ample time and stood eight of twenty-two Imazu cases into 812 to 885 m passes, for four days with the suite green throughout. Two mechanisms: the invariant was accepted on a suite run against a different tree from the one committed; and the baseline test asserted only zero collisions, while the made-the-distance count was written to a document that no test checked.

The same fault class also produced the incident of 11 August described in Section 4: a stopped vessel classified as a crossing stand-on, a derived Rule 17(b) that correctly concluded her action alone could never avoid anything, and 5,328 helm engagements and releases in one session as the live loop re-derived engagement every frame. The making-way gate and the engagement latch were the fixes.

11.1 Rules of coverage

1. **A benchmark certifies only what it exercises.** "22 of 22, no violations" means compliant across these twenty-two geometries. The published tables carry that sentence, and so does this paper.
2. **Classification coverage is not action coverage.** Every rule has two halves, naming the situation and doing the right thing when action falls due. The classification tests all passed while the ship cut the overtaker's bow.
3. **Role coverage is a matrix.** Each rule needs cases with own ship in every chair: overtaking and overtaken, give-way and stand-on, acting early and forced to act late, and transitions between chairs mid-encounter. A published benchmark inherits its author's question.
4. **A new decision path resets coverage to zero.** The rules the core dropped were all enforced somewhere, in paths the benchmark no longer routed through.
5. **A number written to a document rather than asserted by a test is not coverage,** and the suite must be run on the tree being committed.

gap	case to write	status
Rule 13, overtaken chair	escape goes away from the overtaker or takes way off, never across her bow	pinned
Rule 17, stand-on chair	act neither early nor into the give-way vessel's chosen side	lapse pinned; side discipline open
Rule 14 doubt band	meetings fine on the bow, swept through the band, both vessels acting on their own classification	open
Rule 2	the departure chosen is the least that avoids collision	open
every rule $\times$ every chair	a systematic audit of the three suites against the matrix	open

Table 12: The coverage worklist. The structural companion is one rule-discipline provider that every decision path must price candidates through, so that the next new path cannot compile without the Rules rather than cannot be caught by a benchmark that never asked.

12 Where the model runs

Each vehicle carries an ArduPilot stack, ArduRover for the boats, and the question of how much firmware to modify was settled by the project's own history. The serious faults found in August 2026 were two implementations of one judgement drifting apart; the Core assembly is the single tested implementation of the Rules and is to remain the only one. It already runs without the Unity runtime, as the benchmark harness demonstrates, and the simulator already speaks MAVLink v2 over UDP into the same receive path as a real vehicle, so simulated vehicles can be replaced by real ones under the same ground station.

level	what it is	verdict
0	stock firmware, parameters only: missions, geofences, Follow mode, lost-link failsafes, AIS ingest, fence-aware path planning	the baseline everybody maintains for us
1	Lua on the autopilot	a deliberately thin layer for the last-ditch vetoes that must survive companion failure
2	forked firmware with the Rules in C++	never: a fork to rebase across every release, a second implementation of the benchmarked judgement, and GPLv3 obligations on delivery
3	a companion computer running the unmodified Core assembly per vehicle, speaking MAVLink to the autopilot	the model as built, benchmark-pinned, under the project's own terms

Table 13: Firmware levels. The recommendation is 0 plus 3 with a thin 1.

The physics interface is built: the simulator's hull models serve ArduPilot's SITL JSON backend, a unit can be handed to her own firmware end to end by full name, and the guided chain has been proven. What steers the boats in the water is the μ DP stack of the companion paper; what tells it where to go is this one.

13 Roadmap

In the order the evidence of Sections 10 and 11 suggests.

1. **The escort doctrine under dense traffic.** Speed-only defence performed correctly under sparse traffic and stopped the screen along a working sound's lanes (seed 7). The candidate answer, a bounded escort helm against traffic once a speed hold has stood unresolved past a dwell, changes an arrangement the fleet depends on, and will be designed against the benchmarks.
2. **The USV standing instructions in code.** Targets within about a kilometre; react late; never hold for an opening contact with a safe CPA; a station change passes astern of the mothership, never ahead. The manoeuvre core's speed set gains a negative fraction, since every vessel can go astern and station keeping already does.
3. **Inter-unit separation by the plan.** With reactive avoidance removed between own units, two of the seed-11 simulations closed to under ten metres. The guarantee has to be made where the coordination lives, in station assignment and transition planning, and scored there.
4. **The coverage matrix** of Section 11, and the single rule-discipline provider.
5. **Reading the marks.** The planner today reads depth and foul discs; the marks it draws carry meaning it does not use. In order: the hand of a lateral mark given the conventional direction of buoyage (which must first enter the schema); the sector a cardinal mark closes around the danger it names; a channel as the water between two lines of lateral marks, with Rule 9's keep-to-starboard obligation; then isolated-danger, safe-water and special marks. Each step turns a mark from a symbol into a constraint on the clearance field, the same way a charted danger already is, and each is scored by the wrong-hand count the ledger already keeps (Section 6.3).
6. **Confined-water case 1**, and a risk distance that is a function of sea room rather than a constant.
7. **Domains.** Asymmetric ship domains displaced toward bow and starboard [13], which would make pass-astern a property of the geometry and retire the crossing-ahead penalty; and margins that grow with prediction time, the cheap robustness the audit asked for against a deterministic picture.
8. **Rule 19, Rule 10 and vessel categories:** restricted visibility, traffic separation schemes, and vessels not under command, constrained by draught, fishing and sailing.
9. **The S-52 conditional-symbolology engine** (the Presentation Library's lookup tables and procedures over the S-101 reader), the isolated-danger overlay, sector lights and scale minimum: the items the audit lists as absent that a watchkeeper would notice first. Then the IHO S-158 validation checks, held locally, run against the ingested cells.
10. **Rule 19.** The design is in Section 4.6: a visibility state that selects a rule table, aspect in place of roles, both vessels acting, the two sector prohibitions hard, steerage way as a candidate, floors and watch opened out. The Imazu set under that table and the radar corpus's manoeuvre half are its benchmarks.
11. **On the water.** The companion-computer deployment of Section 12, with the passage exercise re-run against real vehicles under the same ground station.

14 Conclusion

The audit of August 2026 found primitives worth keeping and an architecture that was memoryless, instantaneous, pairwise and unmeasured. Each of the four was replaced: encounters have a register; the decision is a whole manoeuvre run through the vehicle's own dynamics; it is scored against every contact and the ground at once; and every change is an experiment against baselines the test suite regenerates. The hard filter, the shared predictor and the rule citation survived the rebuild unchanged.

Twenty-two of twenty-two Imazu encounters make their half mile with no unlawful port alteration; five confined-water cases stay afloat and four make their distance; the fleet's own encounters shape no manoeuvre of the guide's. On passages, a fleet that recorded five thousand

helm engagements in one session records thirty in a night, and one that lost its screen now arrives. The paper also records the three faults the tables did not measure, and the coverage rules derived from them.

The water the model refuses to enter is charted water. The two Channel areas run on the IHO's S-100 Phase 1 products from the official trial datasets, read by readers written against the standard, and the chart that shows them was audited requirement by requirement against S-52 and its Presentation Library. The limits of that statement are stated with it: an evaluation licence, no type approval, and a list of what is still absent.

What remains is doctrine as much as code: what an escort may do with her helm under dense traffic, and where the guarantee between own units is made once they stop avoiding each other. Both will be designed against the benchmarks.

Acknowledgements

We thank Professor Thomas Bräunl of the University of Western Australia for his ongoing support of our work. The radar plotting corpus is the work of the instructors of the École Nationale de la Marine Marchande, Marseille, made freely available by them. Chart data for the Solent testbed is licensed from the UKHO and SHOM under their terms.

References

- [1] International Maritime Organization, *Convention on the International Regulations for Preventing Collisions at Sea, 1972 (COLREGs)*, as amended. Consolidated in U.S. Coast Guard, *Navigation Rules and Regulations Handbook*, 2024.
- [2] National Geospatial-Intelligence Agency, *Radar Navigation and Maneuvering Board Manual*, Pub. 1310, 7th ed., 2001.
- [3] P.-L. Constant, μDP : *Portable Marine Control*, Oceanomatics white paper OCX-WP-2026-01, Rev. C, September 2026.
- [4] T. A. Johansen, T. Perez and A. Cristofaro, "Ship collision avoidance and COLREGS compliance using simulation-based control behavior selection with predictive hazard assessment," *IEEE Trans. Intelligent Transportation Systems*, vol. 17, no. 12, pp. 3407–3422, 2016.
- [5] B.-O. H. Eriksen, G. Bitar, M. Breivik and A. M. Lekkas, "Hybrid collision avoidance for ASVs compliant with COLREGs rules 8 and 13–17," *Frontiers in Robotics and AI*, vol. 7, 2020 (arXiv:1907.00198).
- [6] B.-O. H. Eriksen, M. Breivik, E. F. Wiltail, A. L. Flåten and E. F. Brekke, "The branching-course model predictive control algorithm for maritime collision avoidance," *J. Field Robotics*, vol. 36, no. 7, pp. 1222–1249, 2019.
- [7] D. K. M. Kufoalor, E. F. Wiltail, I. B. Hagen, E. F. Brekke and T. A. Johansen, "Autonomous COLREGs-compliant decision making using maritime radar tracking and model predictive control," *Proc. European Control Conference*, 2019.
- [8] Y. Kuwata, M. T. Wolf, D. Zarzhitsky and T. L. Huntsberger, "Safe maritime autonomous navigation with COLREGS, using velocity obstacles," *IEEE J. Oceanic Engineering*, vol. 39, no. 1, pp. 110–119, 2014.
- [9] K. Woerner, M. R. Benjamin, M. Novitzky and J. J. Leonard, "Quantifying protocol evaluation for autonomous collision avoidance," *Autonomous Robots*, vol. 43, pp. 967–991, 2019.
- [10] R. Sawada, K. Sato and T. Majima, "Automatic ship collision avoidance using deep reinforcement learning with LSTM in continuous action spaces," *J. Marine Science and Technology*, vol. 26, pp. 509–524, 2021. (Source of the 22-case Imazu set as used here.)

- [11] H. Imazu, *Research on collision avoidance manoeuvre*, PhD thesis, Tokyo University of Mercantile Marine, 1987.
- [12] Y. Huang, L. Chen, P. Chen, R. R. Negenborn and P. H. A. J. M. van Gelder, “Ship collision avoidance methods: state-of-the-art,” *Safety Science*, vol. 121, pp. 451–473, 2020.
- [13] M. Wielgosz and Z. Pietrzykowski, “The ship domain in navigational safety assessment,” *PLOS ONE*, vol. 17, no. 4, 2022.
- [14] A. Vagale, R. Ouicheikh, R. T. Bye, O. L. Osen and T. I. Fossen, “Path planning and collision avoidance for autonomous surface vehicles I: a review,” and “II: a comparative study of algorithms,” *J. Marine Science and Technology*, vol. 26, 2021.
- [15] M. Akdağ, T. I. Fossen and T. A. Johansen, “Collaborative collision avoidance for autonomous ships using informed scenario-based model predictive control,” *IFAC-PapersOnLine*, vol. 55, no. 31 (CAMS), 2022.
- [16] Lee, Park and Kim, “Turning-circle control barrier functions for maritime collision avoidance,” arXiv:2504.19247, 2025.
- [17] *Review of performance testing and evaluation of collision avoidance for intelligent vessels*, *J. Marine Science and Engineering*, 2025.
- [18] A. Nash, K. Daniel, S. Koenig and A. Felner, “Theta*: any-angle path planning on grids,” *Proc. AAAI*, 2007.
- [19] École Nationale de la Marine Marchande, Marseille, radar plotting exam papers and tutorials (corrected), http://pemlb.free.fr/NAV_radar.htm, accessed August 2026.
- [20] International Hydrographic Organization, *S-52 Specifications for Chart Content and Display Aspects of ECDIS*, ed. 6.1.1; *S-100 Universal Hydrographic Data Model*, ed. 5.2; *S-101 ENC Product Specification*; *S-102 Bathymetric Surface*.
- [21] ArduPilot Development Team, *Rover documentation; AIS; object avoidance; SITL JSON interface*, <https://ardupilot.org>, accessed September 2026.

A Settings reference

Defaults as of the software baseline on the cover, grouped by the object that owns them. The last column says what the setting decides and, where a value was learned rather than chosen, where in the paper the lesson is. Settings marked *old path* bind only the tangent-and-grid core kept behind `UseManoeuvreCore` and are retained so that the comparison between the two cores stays meaningful.

Table 14: Settings by owner.

setting	default	decides
<i>Separation floors (built per contact in FleetSimulator)</i>		
<code>TrafficClearanceMeters</code>	370.4 m	base separation from traffic, a fifth of a mile; the floor is this plus the larger hull radius (Section 3.2)
<code>HullMarginMeters</code>	20 m	water between two hulls before they are touching; the floor against a detached own unit is both hulls plus this
formation reach	removed	the screen’s extent projected toward a contact, once added to the leader’s floor; retired 4 Sep 2026 (Section 7.3)
<i>SafetySettings (Core/Control/SafetyFilter.cs)</i>		
<code>UseManoeuvreCore</code>	true	the manoeuvre core, or the tangent-and-grid incumbent (Section 5.5)
<code>MarginMeters</code>	30 m	margin above every floor
<code>HorizonSeconds</code>	360 s	look-ahead against traffic in the filter
<code>CooperativeHorizonSeconds</code>	90 s	look-ahead against a detached own unit; fixed after two failed attempts to derive it per frame
<code>StepSeconds / SteppedHorizonSeconds</code>	1 s / 30 s	the stepped walk before the closed-form solve
<code>VerticalGuardMeters</code>	60 m	air and surface are severed beyond this
<code>TerrainMarginMeters / TerrainHorizonSeconds / TerrainStepSeconds</code>	40 m / 150 s / 6 s	foul-water margin, horizon and step for the ship (Section 3.5)
<code>CautionClearanceMeters / CautionWaterPenalty</code>	250 m / 45	the caution band and its rising cost
<code>CorridorExcursionPenalty</code>	40	cost per half-width of corridor excursion, uncapped
<code>SpeedChangeCost</code>	3.0	degrees of helm per m/s of speed change (Eq. (4))
<code>PortAlterationPenalty</code>	15	starboard preference where no rule governs
<code>ManoeuvreCommitmentDegrees</code>	55	cost of reversing an alteration already begun (Rule 8(b) in arithmetic)
<code>CrossingAheadPenalty / CrossingAheadRangeMeters</code>	600 / 1200 m	any candidate passing astern beats any crossing ahead; raised from 60 on 12 Aug 2026 (Section 6)
<code>CrossingAheadHysteresisMeters / CrossingAheadTimeSeconds</code>	150 m / 90 s	when a crossing-ahead reading may change
<code>ProhibitCrossingAhead</code>	false	a hard version of the above, kept for experiments
<code>ActWithinSeconds / TurningActWithinSeconds</code>	300 s / 90 s	the watch window before action is due; shorter against a turning contact (Section 5)
<code>WatchClearanceMeters</code>	500 m	clearance below which the watch is not offered
<code>ReleaseHysteresis</code>	0.10	the helm is released only once the picture has cleared by this fraction

setting	default	decides
TurnTowardPenalty / TurnTowardRangeMeters	30 / 1200 m	do not turn toward a closing contact inside this range
AnswerBeamContactWithSpeed / BeamSpeedSamples	true / 4	speed candidates above the ordered speed against a beam contact (Pub. 1310 rule 5)
MaxHeadingDeviationDegrees / HeadingSamplesPerSide / SpeedSamples	90 / 9 / 4	<i>old path</i> : the grid search, hence 10° steps
SolveEvasion	true	<i>old path</i> : the tangent solver before the grid
LookAheadMeters	15 NM	<i>old path</i> : effectively dead, the 900 s ceiling always binds
<i>ManoeuvreCore (Core/Control/ManoeuvreCore.cs; constants, not settings)</i>		
Alterations	0, ±20, ±30, ±45, ±60, ±90	the eleven helm candidates
SpeedFractions	1.0, 0.85, 0.6, 0.3, 0.0, 1.5	the six speed candidates; no negative fraction yet (Section 13)
HorizonSeconds / StepSeconds	900 s / 10 s	the trajectory horizon and step (Figure 11)
IncreaseCost	4	multiplier on a speed increase; at 2 it was bought routinely and cost four arrivals
MetersPerDegreeOfDrift	100 m	drift from the plan priced as one degree of helm per hundred metres
MeaningfulGainMeters	75 m	bar for a degradation to beat standing on; 1 m once holding fails the floor (Section 11)
<i>EncounterSettings (Core/Colregs/Encounter.cs)</i>		
EncounterRangeMeters	4000 m	how far off a rule can apply at all
RiskClosestApproachMeters	400 m, floored by the separation floor	when a rule applies; anything worth manoeuvring for is worth classifying
MakingWaySpeed	0.26 m/s	both vessels must be making way before a steering rule divides responsibilities (Section 11)
HeadOnBearingDegrees / HeadOnDoubtBearingDegrees	12 / 30	the head-on core and the Rule 14(c) doubt band; the band is unmeasured (Section 10.5)
HeadOnHeadingToleranceDegrees	25	how far from reciprocal two headings may be
OvertakingSectorDegrees	67.5	the Rule 13 sector, 22.5° abaft the beam
IntendedCourseDegrees	the current leg	the course the role geometry is read against, not the bow (Figure 4)
GiveWayAlterationDegrees	30	the alteration the give-way vessel is credited with when her deadline is solved
AssumedTurnRate / ReferenceHullRadiusMeters	0.5°/s at 100 m radius, clamped 0.2 to 3.0	her manoeuvring credit for the Rule 17 lapse, scaled by AIS length; was 1.5°/s for everything afloat
StandOnMustActSeconds	60 s	the last fixed clock, now bounded by the solved lapse
<i>EncounterMemory (Core/Colregs/EncounterMemory.cs)</i>		
GraceSeconds	30 s	a suspected encounter that lapses inside this leaves no entry (Figure 5)
RoleFlipLapseSeconds	120 s	imminence inside which a role-flipped pin arms its own lapse

setting	default	decides
RiskReleaseFactor	1.0	risk hysteresis; built, switched off, a finding (Section 10.3)
<i>ColregsSettings (Core/Colregs/ColregsPolicy.cs)</i>		
Posture	Defensive	the posture handed to the filter
StandOnHorizonSeconds	90 s	how far a stand-on vessel may look, against the filter's 360
ReadilyApparent	20	Rule 8(b) minimum, applied only where admissible and never onto the forbidden side
AlterationDegrees		
PortAlterationPenalty /	120 /	port priced under a rule (before it was struck out)
Ungoverned	15	and where no rule governs
PortAlterationPenalty		
GiveWaySpeedChangeCost /	12 / 200 / 45	the policy layer's own weights, <i>old path</i>
CrossingAheadPenalty /		
TurnTowardPenalty		
<i>ContactMotion (Core/Control/ContactMotion.cs)</i>		
TurnPersistenceSeconds	60 s	how long an observed rate of turn is believed (Section 3.3)
steady threshold	0.5°/min	below this a contact is run straight
rate-of-turn rejection	>120°/min	AIS values above this, or the not-available sentinel, are ignored
<i>Fleet, escorts and consent (Unity/FleetSimulator.cs)</i>		
EscortHoldHorizonSeconds	90 s	the escort speed governor's horizon (Figure 24)
UsvTerrainMarginMeters	8 m	the boats' foul-water margin against the ship's 40
ShipHullRadiusMeters	40 m	a hull at or above this is a ship: speed-change reluctance 6 rather than 1
rally threshold / release	700 m / 350 m	the leader eases to 40 % when her worst station error exceeds the first and resumes inside the second
off-plan allowance cap	400 m	station scatter covered by the floors; beyond it a member looks after herself
departure gate	700 m	a formation forms up before it proceeds
ProposalWindowSeconds /	25 s / 5 s	how long a helm proposal waits for consent (Section 9)
ProposalMinimumSeconds		
ProposalMarginSeconds	60 s	a proposal is never open later than this before TCPA
VetoStandsSeconds	180 s	how long a veto holds
<i>Hulls (Unity/FleetSimulator.cs, Core/Control/Kinematics.cs)</i>		
mothership	152 m, 6.0 m draught	hull radius 76 m; cruise 12 kn, max 18 kn, astern 0.35 of max; $R_{\min}$ 270 m, $\dot{\psi}_{\max}$ 3°/s, G 0.02, 0.08 m/s ²
USV	0.5 m draught	hull radius 5 m; max 65 kn in the simulator; $R_{\min}$ 40 m, $\dot{\psi}_{\max}$ 20°/s, G 0.20, 2.5 m/s ²
ChartDepthSettings safety contour	ship depth, adjustable deeper	the contour the terrain veto enforces; never allowed shallower than the hull requires (Section 8)

B The Imazu set as run

Own ship starts six miles south of the meeting point heading north at twelve knots; every target is placed at the position in Table 15 and given the speed that brings her to the meeting point in thirty minutes, holding course and speed throughout. Required passing distance half a mile. Cases 1 to 4 have one target, 5 to 12 two, 13 to 22 three. The geometries follow Imazu [11] as tabulated by Sawada et al. [10], Appendix 1, Table 4, transcribed once in the harness's **ImazuCases**; the paper's frame (X forward, Y to starboard, heading from +X) maps exactly onto north, east and compass bearing.

The set's known limits, established by running it: own ship is never the overtaken vessel (case 3, 7, 15, 17, 20 and 22 place a slower target dead ahead that own ship overtakes; no case puts a faster vessel astern); no target lies in the head-on doubt band; and no case includes one of our own units. The confined-water and fleet sets exist for the last, and the coverage matrix of Section 11 for the first two.

Table 15: The Imazu set as encoded. Target start positions in the paper's frame (X forward along own ship's course, Y to starboard, nautical miles, origin at the meeting point), target course as a compass bearing, and the target's initial bearing from own ship at her start six miles astern of the origin. Every target's speed is set so that she reaches the origin in 30 minutes.

case	ships	X (NM)	Y (NM)	course	initial bearing
1	1	6.000	0.000	180	000
2	1	0.000	6.000	270	045
3	1	-4.200	0.000	000	000
4	1	-4.243	-4.243	045	292
5	2	6.000	0.000	180	000
		0.000	6.000	270	045
6	2	-5.909	1.042	350	085
		-4.243	4.243	315	068
7	2	-4.200	0.000	000	000
		-4.243	4.243	315	068
8	2	6.000	0.000	180	000
		0.000	6.000	270	045
9	2	-5.196	3.000	330	075
		0.000	6.000	270	045
10	2	0.000	6.000	270	045
		-5.796	-1.553	015	277
11	2	0.000	-6.000	090	315
		-5.196	3.000	330	075
12	2	-4.243	4.243	315	068
		-5.909	1.042	350	085
13	3	6.000	0.000	180	000
		-5.909	-1.042	010	275
		-4.243	-4.243	045	292
14	3	-5.909	1.042	350	085
		-4.243	4.243	315	068
		0.000	6.000	270	045

case	ships	X (NM)	Y (NM)	course	initial bearing
15	3	-4.200	0.000	000	000
		-4.243	4.243	315	068
		0.000	6.000	270	045
16	3	-2.970	-2.970	045	316
		0.000	-6.000	090	315
		0.000	6.000	270	045
17	3	-4.200	0.000	000	000
		-5.909	-1.042	010	275
		-4.243	4.243	315	068
18	3	4.243	4.243	225	023
		-5.796	1.553	345	083
		-5.196	3.000	330	075
19	3	-5.796	-1.553	015	277
		-5.796	1.553	345	083
		4.243	4.243	225	023
20	3	-4.200	0.000	000	000
		-5.796	1.553	345	083
		0.000	6.000	270	045
21	3	-5.796	1.553	345	083
		-5.796	-1.553	015	277
		0.000	6.000	270	045
22	3	-4.200	0.000	000	000
		-4.243	4.243	315	068
		0.000	6.000	270	045

C Scoring definitions

Per encounter, per case, after Woerner [9]:

CPA (m) the closest approach attained to any contact over the run.

safe CPA at or beyond the required passing distance (926 m open water, 400 m confined and fleet).

collision closest approach inside the sum of the two hull radii.

first action (s) time from the start at which the issued command first departed from the ordered one.

stbd / port (°) the largest starboard and port helm applied, charged on the helm and only where the resulting course also lies to that side of the plan, so that the port wheel put on to come back after a lawful starboard avoidance is not a port alteration.


port under a rule steps on which the helm was to port while Rule 14, 15 or 17(c) forbade it against any vessel at risk and Rule 2 did not permit it.

in extremis (steps) steps on which Rule 2 was in force.

reversals changes of the side of the helm during the run.

extra (m) distance run beyond the ordered track's length.

arrived the waypoint reached inside the time allowed.



own ship (blue) starts 6 NM south of the meeting point (red) at 12 kn, heading north; each target (grey) is placed at the position tabulated in Appendix B and given the speed that brings her to the meeting point in 30 minutes. Frame 13 NM square, north up.

Figure 31: The twenty-two Imazu encounters at their start. The construction is what makes the set worth having: every ship is aimed at the origin, so each case is a genuine collision rather than a near miss that resolves itself, and the cases are comparable with every result published against them.

least water (m) confined water only: the smallest clearance-field reading under own ship at any step; *aground* when inside the hull's draught contour.

outside corridor (s) confined water only: seconds spent beyond the alarm line.

shaped by a subordinate (steps) fleet only: steps on which the issued command differed from the ordered one because of a unit that takes her station from us.

D The passage exercise

The runtime harness is a player build driven from the command line. The reference passage in the Solent runs out and back between 50.7800 N, 1.2800 W and 50.6700 N, 1.0500 W so that N runs measure one passage N times:

```
1 ./Build/Linux64/RTGroundControl.x86_64 -rtgc-area solent \
2   -rtgc-passages <seed> -rtgc-passage-count 20 -rtgc-reference \
3   -rtgc-formations -rtgc-timescale 6 -rtgc-tracklog
```

-rtgc-passages <seed> fixed seed for destinations and synthetic traffic.

-rtgc-passage-count <n> passages per run, 1 to 100.

-rtgc-reference [lat lon] the marked destination instead of a drawn one.

-rtgc-formations cycle the formation shape every passage.

-rtgc-timescale <1-6> time acceleration.

-rtgc-tracklog two-hertz track log and avoidance log, flushed on SIGTERM.

Outputs, under the player's persistent data path:

passage-ledger.csv one row per run: seed, passages, arrived, timed out, off-track %, worst cross-track, aground samples and units, shoalest water, distance run, wrong side of a mark, nearest danger, helm engagements and releases, own-unit manoeuvres, escort holds, ship machinery movements, no-safe-option verdicts, closest traffic, and station mean, worst and on-station %.

sim-runs.csv one row per simulation: the shape sailed, outcome, aground samples, shoalest water, closest own-unit pair, station worst and mean, traffic holds and foul-ground eases.

passage-holdups.csv every hold-up, tagged traffic, foul ground, station aground, helm held or press on, with the contact named for traffic holds.

tracks.csv time, unit, kind, east, north, up, heading, speed, avoiding.

avoidance.csv time, unit, event, contact, manoeuvre, alteration, speed off, encounter, CPA, TCPA, range, no-safe, held off by, standing on, lapsed, own and target state, floor.

The ledgers quoted in Section 10.6 are committed under `docs/trials/`.