Worked Synthesis Case Study: A Certified Hybrid Robot in the Sweedler Framework

Abstract

This chapter is a worked synthesis of the constructions developed across the tutorial series. A battery-limited autonomous robot must approach and dock at a charging station while using noisy sensing, discrete control, continuous dynamics, hybrid events, and a finite energy reserve. The purpose is not to propose a new robotics model. It is to test whether the Sweedler programme can organise a genuinely heterogeneous cyber–physical system without flattening its probabilistic, discrete, smooth, and resource-sensitive parts into one semantics. The example combines stateful runners, controlled-flow residuals, barrier and progress certificates, amortised potentials, assume–guarantee wiring, behavioural abstraction, and a schematic cross-ambient architecture. Unlike a purely schematic example, the repaired finite policy below is accompanied by explicit parameters and witnesses for which the local hypotheses are jointly satisfiable. The final repair makes probabilistic sensing load-bearing: robust safety and the mission deadline impose opposite-sided requirements on the estimator's clearance contract. The final indexed architecture of heterogeneous Sweedler worlds is deliberately presented as programme-level structure rather than as an already established theorem.

1 Purpose and relation to the tutorial series

The tutorial series developed a number of constructions separately. Part III treated tight interpretations, wiring, and tensor structure; Part IV introduced predicate-valued certification; Part V developed reachability, observability, and quantitative transport of predicates; Part VI introduced graded and amortised structure; Part VII examined concurrency and deadlock; Part IX developed behaviour and minimal realisation; Part X moved to smooth systems and Lyapunov-style reasoning; Part XA introduced probabilistic ambient worlds; and Part XI studied feedback and trace. The present chapter asks whether these ingredients can coexist in one end-to-end system.

The case study is deliberately small. A one-dimensional autonomous robot approaches a charging dock mounted at a fixed offset from a wall whose absolute location is unknown. It receives noisy range information, runs a finite-state controller, evolves continuously between clocked events, must retain enough battery energy to finish the mission, and must dock before a fixed deadline. The target is not merely to construct the closed-loop dynamics, but to obtain global guarantees from local certificates carried by separate components.

The guiding separation is

$$\text{interface theory} \neq \text{stateful implementation} \neq \text{behavioural abstraction} \neq \text{certificate witness}.$$

A Sweedler runner expresses stateful implementation as a translation from a source theory to a residual theory. Behaviour forgets implementation details. Certification records quantitative obligations that may themselves disappear under behavioural quotient while leaving their conclusions intact.

The local verification ingredients used below have standard precedents. Barrier certificates and their stochastic variants are developed by Prajna–Jadbabaie–Pappas [1], while control barrier functions provide a control-theoretic formulation of forward invariance [2]. Zeno behaviour in hybrid automata is classical [3]; contract-based design provides the nearest engineering precedent for the assume–guarantee layer [4]; and differential dynamic logic provides an important alternative framework for compositional verification of hybrid programs [5]. On the categorical side, Lerman–Schmidt's networks of hybrid open systems are a particularly close precedent for compositional open hybrid dynamics [6].

2 A repaired concrete robot with genuinely uncertain geometry

The physical motion state is

$$x_M = (r, v),$$

where r is position in a fixed laboratory frame and $v \geq 0$ is forward velocity. Battery energy $e \geq 0$ is deliberately kept in a separate subsystem. The wall position is an unknown static parameter

$$Z \in [10, 12].$$

For a Bayesian semantics one may additionally equip this interval with a non-degenerate prior μ_0; the concrete end-to-end theorem below is stated per installed wall position, using the sensor kernel indexed by fixed Z. The charging dock is physically mounted a fixed distance

$$d_* = 1$$

in front of the wall, so its absolute position is $r_*(Z) = Z - d_*$. The controller therefore cannot infer either wall clearance or absolute dock position from the clock alone. It starts at

$$r_0 = 0, \qquad v_0 = v_C = 1,$$

so the robot is already in cruise rather than passing through a separate acceleration phase.

Write

$$d = Z - r$$

for true wall clearance. We require

$$d \geq d_{\text{safe}}, \qquad d_{\text{safe}} = 0.5,$$

and impose the mission deadline

$$T_{\max} = 15.$$

The clock period, creep speed, and braking deceleration are

$$\Delta = 0.25, \qquad v_K = 0.25, \qquad b = 1.$$

The final-brake trigger clearance is

$$d_B = d_* + \frac{v_K^2}{2b} = 1.03125.$$

Because the trigger is sampled at clock times, we do not require exact landing. Put

$$\varrho = v_K \Delta = 0.0625.$$

The docking region is

$$\mathcal{D}_\varrho = \{(d, v) : d_* - \varrho < d \le d_*,\ v = 0\}.$$

The landing argument is therefore robust to one creep-sample overshoot. For arithmetic transparency we do retain the commensurability

$$v_C - v_K = 3b\Delta,$$

so the first braking phase reaches v_K exactly at a sampling instant. A non-commensurate saturated brake would require a shortened final braking subsegment, or an additional clock-dependent margin; nothing below claims otherwise.

The controller state is

$$Q = \{C, B, K, D\},$$

standing for cruise, brake, creep, and final docking brake. The probabilistic range estimator exports a lower clearance bound $\underline{d}$ and hence a flag

$$\chi = \begin{cases} \mathsf{clear}, & \underline{d} > 2, \\ \mathsf{brake}, & \underline{d} \le 2. \end{cases}$$

The speed class is

$$s \in \{\mathsf{fast}, \mathsf{slow}, \mathsf{stopped}\},$$

with fast meaning $v > v_K$, slow meaning $0 < v \le v_K$, and $\mathsf{stopped}$ meaning $v = 0$. Finally, a separate short-range dock beacon supplies

$$\kappa \in \{\mathsf{approach}, \mathsf{dockZone}\},$$

where $\kappa = \mathsf{dockZone}$ exactly when the robot is at a sampling instant with $d \le d_B$. The beacon is idealised as exact in the worked instance. Only the range channel is noisy; an additional beacon error budget could be incorporated in exactly the same contract layer. The controller observation alphabet is therefore

$$O = \{\mathsf{clear}, \mathsf{brake}\} \times \{\mathsf{fast}, \mathsf{slow}, \mathsf{stopped}\} \times \{\mathsf{approach}, \mathsf{dockZone}\}.$$

The policy is

q	condition	flow mode	q'
C	$\chi = \mathsf{clear}$	Cruise	C
C	$\chi = \mathsf{brake}$	Brake	B
B	$s = \mathsf{fast}$	Brake	B
B	$s = \mathsf{slow}$	Creep	K
K	$\kappa = \mathsf{approach}$	Creep	K
K	$\kappa = \mathsf{dockZone}$	DockBrake	D
D	$s \neq \mathsf{stopped}$	DockBrake	D
D	$s = \mathsf{stopped}$	terminal	D.

All nonterminal flow segments use the closed clock guard $\theta = \Delta$, with reset $\theta := 0$. The policy uses only quantities actually exposed by its observation interface; in particular it never branches on the hidden wall position Z.

Mission and controller interfaces. Take

$$T_M = \mathrm{Free}(y + y^{O_M}), \qquad O_M = \{\mathsf{enroute}, \mathsf{docked}, \mathsf{failed}\},$$

whose command advance requests one hybrid control cycle. Each nonterminal advance advances physical time by exactly Δ, so mission traces are timed by the count of advance steps: after n such steps the elapsed time is $n\Delta$. Thus the property $T_{\mathrm{dock}} \leq T_{\max}$ is externally expressible even though the mission theory has no separate clock observation. Let

$$F = \{\mathsf{ok}, \mathsf{docked}, \mathsf{collision}, \mathsf{energyFail}\}$$

and

$$\widehat{T}_C = \mathrm{Free}(y + y^{Q \times F}).$$

There is a tight interpretation

$$j : T_M \longrightarrow \widehat{T}_C$$

sending advance to the controller tick and postcomposing the controller/status observation with

$$(q, f) \longmapsto \begin{cases} \mathsf{docked}, & f = \mathsf{docked}, \\ \mathsf{failed}, & f \in \{\mathsf{collision}, \mathsf{energyFail}\}, \\ \mathsf{enroute}, & f = \mathsf{ok}. \end{cases}$$

The substantive content is not this tight map but the state-Q runner below, whose residual theory contains estimator queries, controlled flow, and status queries to the motion and battery subsystems.

3 The finite controller as a runner

Let $T_E = \mathrm{Free}(y^O)$ expose the estimator, speed, and dock-beacon observation, let T_{flow} expose the controlled flow modes, and let T_F expose the combined motion/battery status flag F. Schematically put

$$R_C = T_E + T_{\mathrm{flow}} + T_F.$$

The external controller theory $\widehat{T}_C$ has a state-Q runner

$$\rho_C : \widehat{T}_C \rightsquigarrow R_C.$$

For a pure return,

$$[\![\eta x]\!]_q = \eta(x, q).$$

On a control tick, the runner first queries the observation interface and then emits the flow mode selected by the policy table:

$$[\![\mathsf{tick}(p)]\!]_q = \mathsf{estimate}\big(o \mapsto \mathsf{flow}_{\alpha(q,o)}\big([\![p]\!]_{\delta_C(q,o)}\big)\big)\,.$$

A status query is implemented by querying T_F and returning the pair (q, f). Thus collision and energy-failure information comes from the plant and battery residuals rather than being invented as controller state.

By the Sweedler adjunction, the same stateful implementation is a coalgebra

$$\xi : Q \longrightarrow \{\widehat{T}_C, R_C\}(Q).$$

The measuring-comonad law is ordinary state-threading coherence: after one tick changes the hidden mode, the next source operation is interpreted in the new mode. Composing $j : T_M \to \widehat{T}_C$ tightly with ρ_C loosely yields the mission-level implementation once the estimator, flow, status, and battery residuals have been wired.

4 Controlled flow and hybrid time

Let $A_H = \{\mathsf{Cruise}, \mathsf{Brake}, \mathsf{Creep}, \mathsf{DockBrake}\}$ and let $\mathcal{H}_{A_H}$ be the monoid of finite-duration piecewise-constant mode signals. The unary theory

$$T_{\text{flow}} = \mathcal{H}_{A_H} y$$

has operations flow_σ satisfying concatenation. A smooth or hybrid plant realises this operational theory by flow maps.

For the concrete instance, in clearance coordinates $d = Z - r$ the mode dynamics are

mode	$\dot{d}$	$\dot{v}$
Cruise	$-v_C$	0
Brake	$-v \quad -b \quad (v > v_K),$	$0 \quad (v = v_K)$
Creep	$-v_K$	0
DockBrake	$-v \quad -b \quad (v > 0),$	$0 \quad (v = 0).$

The invariant domains are $v = v_C$ in cruise, $v \in [v_K, v_C]$ in brake, $v = v_K$ in creep, and $v \in [0, v_K]$ in final brake. In particular, neither braking mode can reverse the robot.

The clock variable satisfies

$$\dot{\theta} = 1, \qquad G_{\text{sample}} = \{\theta = \Delta\}, \qquad \theta := 0 \text{ on reset.}$$

The guard is closed, so the sampling hitting time is attained. This is the structural origin of the dwell-time lower bound used below. In a smooth-topos/SDG reading, infinitesimal vector fields supply local dynamics, while integration to finite flow maps is an additional hypothesis in general; for the elementary piecewise-linear flows used here it is explicit.

5 Probabilistic sensing, sharpness, and the correct Markov state

The wall position is fixed during one installation but unknown to the controller. At the kth sampling time the range sensor returns

$$Y_k = d_k + \eta_k, \qquad d_k = Z - r_k.$$

The estimator exports a lower clearance value $\underline{d}_k$. The certification interface is deliberately two-sided. Put

$$A_k^- = \{d_k \geq \underline{d}_k\}, \qquad A_k^+ = \{d_k \leq \underline{d}_k + \delta\}, \qquad \delta = 0.5.$$

For each fixed installed $Z \in [10, 12]$, write $\Pr_Z$ for probability over the sensor noise and require

$$\Pr_Z((A_k^-)^c) \leq \varepsilon_k^-, \qquad \Pr_Z((A_k^+)^c) \leq \varepsilon_k^+. \tag{E}$$

The first inequality is *calibration*: the lower bound must not overstate true clearance. The second is *sharpness*: the lower bound must not be more than δ below true clearance. These obligations have opposite operational variance. Safety uses calibration, whereas the mission deadline uses sharpness.

A concrete per-installation realisation is immediate for Gaussian range noise of known variance. If $\eta_k \sim N(0, \sigma^2)$ and $c = 0.25 = z_{1-\alpha/2}\sigma$, then

$$\underline{d}_k = Y_k - c$$

satisfies $\delta = 2c = 0.5$ and

$$\varepsilon_k^- = \varepsilon_k^+ = \alpha/2.$$

No independence assumption is needed for the union bounds used below. More elaborate estimators may maintain a posterior state π_k and export tighter intervals; the same interface contract applies. For notational convenience, stipulate that the range sensor is sampled at every clock instant, in every mode, although the controller consumes the range output only while in cruise. Thus the random variables Y_k and events $A_k^{\pm}$ below are defined even on executions that have already left cruise.

There are two probabilistic readings worth separating. The concrete theorem below uses $\Pr_Z$ and therefore gives a guarantee for every fixed installed wall position. If instead $\underline{d}_k$ is a Bayesian posterior lower quantile under a prior μ_0, the analogous conditional-coverage statement is a statement in the joint law of $(Z, Y_0, Y_1, \ldots)$ and the resulting end-to-end probability is Bayes-averaged over $Z \sim \mu_0$. In the QBS picture this is precisely the distinction between a probability kernel indexed by Z and the joint measure obtained after integrating that kernel against a prior.

If the estimator maintains history, the sampled process is not Markov on $X_M \times Q$. The correct sampled state then contains the estimator state:

$$S_\Delta = X_M \times Q \times \Pi \times \mathbb{R}_{\geq 0} \times F.$$

For a finite controller observation, a one-step kernel has the form

$$K_\Delta((x_M, q, \pi, e, f), -) = \sum_{o \in O} K_E(\pi, x_M; o)\, \delta^{\mathrm{Dir}}_{\Psi_\Delta(x_M, q, \pi, e, f, o)},$$

where δ^{Dir} denotes the Dirac measure and Ψ_Δ includes estimator update, controller update, controlled flow, battery update, and status update. For the memoryless Gaussian interval above, the Π factor can be omitted.

Why the threshold 2 is safe. If a cruise query returns $\chi = \mathsf{clear}$ and A_k^- holds, then

$$d_k \geq \underline{d}_k > 2.$$

During the following clock interval the clearance falls by exactly $v_C \Delta = 0.25$, so the next sampling instant still has

$$d_{k+1} > 1.75.$$

At the first query with true clearance $d_k \leq 2$, calibration forces $\underline{d}_k \leq 2$ and hence braking. Thus the safety side of (E) prevents braking from being delayed beyond the critical query.

Why sharpness is needed for the deadline. If $\chi = \mathsf{brake}$ and A_k^+ holds, then

$$d_k \leq \underline{d}_k + \delta \leq 2.5.$$

Hence, on the two-sided good event, the brake-entry clearance d_s satisfies

$$1.75 < d_s \leq 2.5. \tag{D}$$

After the first brake phase the clearance is $d_s - 0.46875$. The number n_K of creep intervals before the dock beacon can fire obeys

$$n_K \leq \left\lceil \frac{d_s - 1.5}{0.0625} \right\rceil \leq 16(d_s - 1.5) + 1.$$

Therefore, using $Z \le 12$,

$$\begin{aligned} T_{\text{dock}} &= (Z - d_s) + 0.75 + 0.25 n_K + 0.25 \\ &\le Z + 3d_s - 4.75 \\ &\le 14.75 < T_{\max} = 15. \end{aligned} \tag{T}$$

Thus safety and deadline impose opposite-sided demands on the same estimator output.

This makes sensing genuinely load-bearing for the available mode set, not merely for one chosen switching schedule. Remove the range channel but allow the controller to retain the exact dock beacon, clock, and speed class. In cruise the speed is invariant at v_C, and there is no acceleration mode, so once braking reduces the speed the lost speed cannot be recovered. Beacon silence implies only $d > d_B = 1.03125$. To spend one further cruise period and then stop safely using the braking modes requires the present clearance to exceed

$$d_{\text{safe}} + v_C \Delta + \frac{v_C^2 - v_K^2}{2b} + \frac{v_K^2}{2b} = 0.5 + 0.25 + 0.46875 + 0.03125 = 1.25.$$

Thus beacon silence alone never justifies an extra cruise step near the wall. Combining the silence information with the support bound $Z \ge 10$, a range-free controller that is safe for every installation must initiate braking no later than $r = 9$. Even granting the latest and fastest such safe switch, an installation with $Z = 12$ requires at least

$$9 + 0.75 + 6 + 0.25 = 16$$

seconds: 9 seconds of cruise to $r = 9$, 0.75 seconds to brake from v_C to v_K, at least 6 seconds of creep to the dock zone, and 0.25 seconds of final braking. Hence no switching law over these four modes that uses the beacon but not the range channel can satisfy both robust safety over $Z \in [10, 12]$ and the 15-second deadline. The noisy range channel is therefore required by the certified mission itself.

The deadline also propagates to a quantitative sensor specification. Replacing the chosen $\delta = 0.5$ in (T) by a general sharpness width gives

$$T_{\text{dock}} \le Z_{\max} + 3(2 + \delta) - 4.75 = 13.25 + 3\delta.$$

Thus $T_{\max} = 15$ permits

$$\delta \le \delta_{\max} = \frac{7}{12} \simeq 0.583.$$

For the symmetric Gaussian construction $\delta = 2z_{1-\alpha/2}\sigma$, this becomes a hardware-level noise constraint. At $\alpha = 10^{-4}$ the deadline alone permits $\sigma \lesssim 0.075$, while the more conservative design choice $\delta = 0.5$ used here requires $\sigma \lesssim 0.064$. The mission deadline therefore determines an admissible estimator sharpness, which in turn determines a sensor-noise specification at the chosen risk level.

The fixed-speed calculation already includes one full sensing/control period of latency. In an accelerating cruise mode the safe lower-bound test would be replaced by a condition such as

$$\underline{d}_k \ge d_{\text{safe}} + v\Delta + \frac{1}{2}a\Delta^2 + \frac{(v + a\Delta)^2}{2b}.$$

6 Explicit local certificates

The revised geometry permits genuinely local safety certificates. The barrier and progress witnesses are proof-level data and may depend on the latent true clearance d; the energy budget below is deliberately different, because it is made operationally computable from (r, v, q) and the known support bound $Z_{\max} = 12$.

Safety as predicted terminal clearance

The cruise segment is certified by the estimator contract. If $\chi = \mathsf{clear}$ on A_k^-, the following interval remains above clearance 1.75. If braking is selected on the two-sided good event, (D) gives $1.75 < d_s \le 2.5$.

For brake mode define

$$h_B(d, v) = d - d_B - \frac{v^2 - v_K^2}{2b}.$$

Along Brake,

$$\dot{h}_B = -v - \frac{v\dot{v}}{b} = 0.$$

At brake entry with $v = v_C = 1$,

$$h_B(d_s, 1) > 1.75 - 1.03125 - 0.46875 = 0.25.$$

Thus the robot reaches v_K while still strictly before the dock-zone trigger.

For a creep interval, include the clock and define

$$h_K(d, \theta) = d - d_{\text{safe}} - v_K(\Delta - \theta) - \frac{v_K^2}{2b}.$$

Along Creep,

$$\dot{h}_K = -v_K + v_K = 0.$$

At the start of every creep interval with $\kappa = \mathsf{approach}$ we have $d > d_B$, hence

$$h_K(d, 0) > d_* - d_{\text{safe}} - v_K\Delta = 0.4375.$$

At the first sampling instant with $\kappa = \mathsf{dockZone}$, the preceding sample had $d > d_B$, so

$$d > d_B - v_K\Delta = 0.96875.$$

The final-brake barrier is

$$h_D(d, v) = d - d_{\text{safe}} - \frac{v^2}{2b},$$

and again

$$\dot{h}_D = 0$$

along $\mathsf{DockBrake}$. At entry with $v = v_K$,

$$h_D > 0.96875 - 0.5 - 0.03125 = 0.4375.$$

The terminal clearance therefore lies in

$$d_* - \varrho < d_{\text{term}} \le d_*,$$

which is exactly $\mathcal{D}_\varrho$. The safety proof is local: estimator calibration discharges the cruise-to-brake obligation, and each deterministic tail mode carries an invariant predicted-terminal-clearance margin.

Progress

Put

$$d_{\min} = d_* - \varrho = 0.9375, \qquad L(d) = d - d_{\min}.$$

Away from the terminal docking region $L \geq 0$. Every nonterminal clock interval has average forward speed at least 0.125, the minimum occurring during the final brake from $v_K = 0.25$ to zero. Hence

$$L' \leq L - 0.125\Delta.$$

We may choose

$$\gamma = 0.1$$

and obtain

$$L' \leq L - \gamma\Delta.$$

Because $d_0 = Z \leq 12$,

$$L_0 \leq 12 - 0.9375 = 11.0625,$$

so the generic counting argument gives the integer bound

$$N \leq \left\lfloor \frac{11.0625}{0.1 \cdot 0.25} \right\rfloor = 442.$$

This proof witness is intentionally coarse; on the two-sided estimator event, the deadline calculation (T) gives the much stronger $T_{\text{dock}} < 15$.

An operational energy potential

Let the mode power rates be

$$p_C = 1, \qquad p_B = 1.2, \qquad p_K = 0.2, \qquad p_D = 0.8.$$

To avoid exporting a ghost depending on the latent Z, define the computable worst-case clearance coordinate

$$\widehat{d} = Z_{\max} - r = 12 - r.$$

Put

$$\widehat{V}_D(v) = p_D \frac{v}{b}.$$

For creep define

$$\widehat{n}_K(r) = \max\left(0, \left\lceil \frac{12 - r - d_B}{v_K \Delta} \right\rceil\right),$$

$$\widehat{V}_K(r) = p_K \Delta\, \widehat{n}_K(r) + \widehat{V}_D(v_K).$$

For a reachable brake state put

$$\widehat{V}_B(r, v) = p_B \frac{v - v_K}{b} + \widehat{V}_K\left(r + \frac{v^2 - v_K^2}{2b}\right),$$

and on cruise states

$$\widehat{V}_C(r) = 12 - r.$$

A continuing cruise step satisfies the exact equation

$$\widehat{V}_C(r) = p_C \Delta + \widehat{V}_C(r + v_C \Delta).$$

At a certified switch from cruise to brake, $d_s > 1.75$ and hence $\widehat{d} > 1.75$. The same ceiling estimate as before gives

$$\widehat{V}_B(r, 1) \le 0.8\widehat{d} - 0.05 < \widehat{d} = \widehat{V}_C(r).$$

The definitions of $\widehat{V}_B, \widehat{V}_K, \widehat{V}_D$ give the local equalities on continuing brake/creep/final-brake segments. If the exact dock beacon fires earlier than the worst-case $Z_{\max}$ budget predicts, the transition $K \to D$ only drops the remaining budget. Thus every certified branch satisfies

$$\widehat{V}(x_M, q) \ge C_M + \widehat{V}(x_M^+, q'), \tag{V}$$

where C_M is the motion subsystem's exported branch-energy allowance.

At the initial state

$$\widehat{V}_C(0) = 12.$$

Choosing

$$e_0 = 13.5, \qquad e_{\text{res}} = 1$$

gives

$$e_0 - \widehat{V}_C(0) = 1.5 \ge e_{\text{res}}.$$

For orientation, exhaustive arithmetic over the certified brake-entry range gives actual total energy at most 11.5; the interface potential is deliberately more conservative.

Dwell time

Every nonterminal flow segment is clocked by the closed guard $\theta = \Delta$ and resets θ to zero. Hence

$$\tau_k = \Delta = 0.25$$

for every nonterminal branch, so Zeno execution is structurally impossible.

Local check

The concrete margins are

obligation	uniform lower margin on certified branches
brake predicted-terminal margin at entry	> 0.25
creep predicted-terminal margin	> 0.4375
final-brake predicted-terminal margin	> 0.4375
average progress rate $- \gamma$	0.025
$\tau - \tau_{\min}$	0 $(\tau_{\min} = 0.25)$.

The energy inequalities are exact by construction except for the deliberately conservative cruise-to-brake comparison. Safety and deadline now use different sides of the estimator contract, while the runtime energy budget uses no latent geometric quantity.

7 Motion–battery wiring: the budget is an interface port

The motion and battery components are coupled by energy demand, so they should not be tensored as though independent. The clean contract exposes the future-energy budget as an explicit interface quantity.

Let

$$R = \widehat{V}(x_M, q), \qquad R' = \widehat{V}(x_M^+, q').$$

The motion subsystem guarantees both

$$C_{\mathrm{act}} \le C_M$$

and, from (V),

$$R' \le R - C_M.$$

The battery subsystem performs the exact update

$$e' = e - C_{\mathrm{act}}.$$

Therefore

$$e' - R' \ge e - C_{\mathrm{act}} - (R - C_M) = e - R + (C_M - C_{\mathrm{act}}) \ge e - R.$$

Hence the invariant

$$\boxed{e - R \ge e_{\mathrm{res}}}$$

is preserved. This is now a genuinely operational certificate port: $\widehat{V}$ is computed from (r, v, q) and the public support bound $Z_{\max}$, while the battery knows only the imported budget R, the branch allowance C_M, and its own charge e. By contrast, the barrier functions h_B, h_K, h_D and the ranking witness L may depend on the latent true clearance and are ghost proof data. This operational/ghost distinction will matter for the refinement architecture in Section 12.

In contract notation,

$$M : A_M \Rightarrow \big(G_M \wedge C_{\mathrm{act}} \le C_M \wedge R' \le R - C_M\big),$$

while

$$B : (C_{\mathrm{act}} \le C_M,\ R' \le R - C_M) \Rightarrow (e' - R' \ge e - R).$$

This is the local assume–guarantee pattern familiar from contract-based design [4], now attached to runner composition.

8 Two-sided estimator risk: safety and deadline

The two sides of (E) support different mission guarantees. For a fixed installed Z, define the critical safety query

$$k_*(Z) = \left\lceil \frac{Z - 2}{v_C \Delta} \right\rceil.$$

Since $Z \in [10, 12]$,

$$k_*(Z) \in \{32, 33, \dots, 40\}.$$

If the robot has remained in cruise until this index, its true clearance is at most 2. If $A^-_{k_*}$ then holds, $\underline{d}_{k_*} \le d_{k_*} \le 2$, so braking is selected. If an earlier sharpness error has already caused braking, safety is already reduced to the deterministic tail certificates and the value of $A^-_{k_*}$ is immaterial. Sampling the range channel at every clock instant ensures the event is nevertheless defined. Hence, for each fixed installation,

$$\boxed{\Pr_Z(\mathsf{collision\ before\ docking}) \le \varepsilon^-_{k_*(Z)}.} \tag{RS}$$

No union over all previous queries is needed for safety.

Deadline satisfaction has the opposite dependence. Let

$$J_Z = \{k : d_k > 2.5 \text{ at the } k\text{th cruise query}\}.$$

If A_k^+ holds for every $k \in J_Z$, then the controller cannot brake while true clearance exceeds 2.5. Combined with $A_{k_*}^-$, braking therefore begins in the interval (D), and the calculation (T) proves $T_{\mathrm{dock}} < 15$. Consequently

$$\boxed{\Pr_Z(T_{\mathrm{dock}} > 15 \text{ or collision}) \le \varepsilon^-_{k_*(Z)} + \sum_{k \in J_Z} \varepsilon_k^+.} \tag{RT}$$

There are at most 38 sharpness-sensitive queries, so if

$$\varepsilon_k^- \le \bar{\varepsilon}_-, \qquad \varepsilon_k^+ \le \bar{\varepsilon}_+,$$

then uniformly over installations

$$\Pr_Z(\text{failure of safe docking by the deadline}) \le \bar{\varepsilon}_- + 38\bar{\varepsilon}_+.$$

This is a union bound, not a product estimate; shared latent geometry or a stateful estimator can correlate errors strongly across queries.

The nine possible values of k_* are

$$1 + \frac{Z_{\max} - Z_{\min}}{v_C \Delta} = 9.$$

Thus the width of the geometry uncertainty appears directly in the set of possible critical safety queries. If one works instead in the Bayesian joint law and does not condition on a fixed installation, a union over these possible critical indices gives the corresponding prior-averaged bound. The concrete theorem keeps the stronger per-installation reading (RS)–(RT).

9 A local certified hybrid step

The repaired instance illustrates the general local pattern. A hybrid branch begins from (x_M, q, e), executes a flow segment, and ends at (x_M^+, q', e'). A certified branch carries:

(H1) a barrier inequality during continuous flow and barrier-preserving reset;

(H2) a motion-side potential inequality $R' \le R - C_M$;

(H3) a battery accounting inequality $C_{\mathrm{act}} \le C_M$;

(H4) a structural dwell bound $\tau \ge \tau_{\min} > 0$;

(H5) a progress inequality $L' \le L - \gamma\tau$ away from docking;

(H6) an estimator-interface obligation whose lower side protects safety and whose upper/sharpness side protects the mission deadline.

These obligations have different variance and need not be forced into one scalar quantale. The two-sided estimator contract is the clearest example: making the lower bound larger helps the deadline but can hurt safety if calibration is lost. The comparison with differential dynamic logic is useful here: dL also decomposes hybrid correctness into proof obligations for discrete and differential program fragments [5]; the Sweedler objective is instead to organise the implementation interfaces, residual capabilities, certificates, and behavioural abstraction in one categorical architecture.

10 Behavioural abstraction

After wiring the system, internal state contains at least

$$S = X_M \times Q \times \Pi \times \mathbb{R}_{\geq 0} \times F \times X_{\text{clock}},$$

together with the latent environmental parameter in the underlying probabilistic model. Mission observations expose only enroute, docked, and failed, with failed explicitly wired to collision or energy-depletion status. Thus safety and resource failure are genuinely expressible in the mission observation vocabulary.

If the fully wired implementation yields a coalgebra

$$\xi : S \longrightarrow \{T_M, R_{\text{ext}}\}(S)$$

and the measuring comonad has a final coalgebra

$$\omega : B_{M,R} \longrightarrow \{T_M, R_{\text{ext}}\}(B_{M,R}),$$

then there is a behaviour map

$$\text{beh} : S \longrightarrow B_{M,R}.$$

The raw witnesses h_B, h_K, h_D, L, π and the operational budget $\widehat{V}$ do not generally descend: behaviourally equivalent states may carry different margins, energy budgets, rankings, or posteriors. What can descend are properties expressible in the chosen observation vocabulary, for example

$$\Pr(\mathsf{failed}_{\text{collision}} \text{ before } \mathsf{docked}) \leq \varepsilon_{\text{safe}},$$

together with the timed behavioural property $T_{\text{dock}} \leq 15$ on the joint estimator good event. The runtime energy invariant is stronger: it is transported by the operational budget port rather than inferred from mission observations alone.

When a quantitative witness must be transported along beh, conservative quantification remains available. For an available-slack quantity one may define

$$\bar{c}(b) = \inf\{c(s) : \text{beh}(s) = b\},$$

with the dual supremum for upper-bound obligations. This is the concrete form, in the present example, of the predicate-transport adjunctions of Part V. The qualification is essential: if failure were not observable, a collision-free and a colliding state could lie in the same behavioural class and no nontrivial safety property could descend.

11 End-to-end theorem for the repaired instance

We can now state the non-vacuous end-to-end result in a form that separates exactly what calibration buys from what estimator sharpness buys.

Theorem schema, instantiated. Fix an installed wall position

$$Z \in [10, 12]$$

and let $\Pr_Z$ denote probability over the sensor/estimator randomness. Assume the estimator satisfies (E) at every clock sample. Write

$$\mathcal{A}_Z = A^-_{k_*(Z)}, \qquad \mathcal{H}_Z = \mathcal{A}_Z \cap \bigcap_{k \in J_Z} A^+_k.$$

Take the controller, saturated plant modes, clock $\Delta = 0.25$, local safety margins h_B, h_K, h_D, ghost progress function

$$L(d) = d - 0.9375,$$

operational energy potential $\widehat{V}$, initial motion state

$$r_0 = 0, \qquad v_0 = 1, \qquad q_0 = C,$$

and battery

$$e_0 = 13.5, \qquad e_{\mathrm{res}} = 1.$$

Then:

(a) on $\mathcal{A}_Z$, the continuous trajectory satisfies $d(t) > d_{\mathrm{safe}}$ until termination and stops inside $\mathcal{D}_\varrho$;

(b) on $\mathcal{A}_Z$,

$$e_k - \widehat{V}(x_{M,k}, q_k) \geq e_{\mathrm{res}}$$

at every sampled hybrid state;

(c) every nonterminal flow segment lasts exactly Δ, hence Zeno execution is impossible independently of estimator correctness;

(d) on $\mathcal{A}_Z$, the ranking witness gives finite termination with the generic bound $N \leq 442$; on the stronger event $\mathcal{H}_Z$, estimator sharpness gives the mission deadline

$$T_{\mathrm{dock}} \leq 14.75 < 15;$$

(e) for collision, energy failure, or failure of eventual docking,

$$\Pr_Z(\text{failure}) \leq \varepsilon^-_{k_*(Z)};$$

more specifically the three properties in (a), (b), and the termination part of (d) hold whenever $\mathcal{A}_Z$ holds;

(f) for successful safe docking by the deadline,

$$\Pr_Z(\textsf{safe docking by } 15)^c \leq \varepsilon^-_{k_*(Z)} + \sum_{k \in J_Z} \varepsilon^+_k.$$

The proof is compositional. Estimator calibration alone prevents a late brake at the critical clearance; this is enough for safety, resource adequacy, and eventual docking. Estimator sharpness is used only to exclude unnecessarily early braking and hence only for timeliness. The brake, creep, and final-brake margins are invariant along their respective flows. The operational potential $\widehat{V}$ supplies the motion-side budget consumed by the battery proof. The clock gives the dwell bound, while L gives a separate finite-progress witness. No exact landing, independent-error assumption, or knowledge of the installed Z by the controller is used.

The deadline is what makes the probabilistic component genuinely load-bearing. As shown above, even a range-free controller allowed to exploit the exact beacon must brake no later than $r = 9$ to retain robust safety for the possible installation $Z = 10$. With no acceleration mode, that loss of speed is irreversible; for $Z = 12$ even the fastest continuation using the available modes takes at least $16 > 15$ seconds. The estimator resolves precisely the uncertainty that separates robust safety from timely docking.

Thus the case study establishes the conditional end-to-end statement

two-sided probabilistic range certification
+discrete stateful control
+certified hybrid flow
+operational resource accounting
$\Downarrow$
safe, energy-adequate, non-Zeno docking by a deadline.

12 Cross-ambient interpretation and certified runners

The example should not be read as saying that all four components naturally live in one ambient category. Rather, it suggests native Sweedler worlds

$$E \in \mathrm{Run}(\mathcal{D}_{\mathrm{prob}}), \qquad C \in \mathrm{Run}(\mathcal{D}_{\mathrm{disc}}), \qquad M \in \mathrm{Run}(\mathcal{D}_{\mathrm{sm}}), \qquad B \in \mathrm{Run}(\mathcal{D}_{\mathrm{res}}),$$

with typical candidates

$$\mathcal{D}_{\mathrm{prob}} = [\mathrm{QBS}, \mathrm{QBS}]_{\mathrm{acc}}, \qquad \mathcal{D}_{\mathrm{disc}} = [\mathbf{Set}, \mathbf{Set}]_{\mathrm{acc}},$$

$$\mathcal{D}_{\mathrm{sm}} = [\mathrm{SmoothSet}, \mathrm{SmoothSet}]_{\mathrm{acc}},$$

and an enriched C-Set endofunctor category for resource reasoning.

The certification layer should not be a functor

$$\mathrm{Run}(\mathcal{D}) \to \mathsf{Cert}$$

that magically extracts witnesses from a bare runner: a runner does not determine its barrier, potential, ranking function, or probabilistic calibration proof. A more natural architecture is a displayed/refinement structure

$$\pi_{\mathcal{D}} : \mathrm{CertRun}(\mathcal{D}) \longrightarrow \mathrm{Run}(\mathcal{D}),$$

whose objects above a runner are chosen certificates and proofs of their compatibility. This is close in spirit to the general categorical view of refinement systems as functors developed by Melliès–Zeilberger [7].

Likewise let

$$\mathrm{CertIfc} \longrightarrow \mathrm{Ifc}$$

be certified interfaces over ordinary interfaces. The native semantics supplies a boundary map

$$\partial_i : \mathrm{Run}(\mathcal{D}_i) \longrightarrow \mathrm{Ifc},$$

while certified components admit extraction maps

$$\widehat{\Gamma}_i : \mathrm{CertRun}(\mathcal{D}_i) \longrightarrow \mathrm{CertIfc}$$

lying over ∂_i. The estimator exports a calibrated clearance contract, motion exports barrier/progress and budget contracts, and the battery exports reserve preservation. Heterogeneous components can then meet at certified interfaces without being flattened into one semantic fibre.

The worked example also shows that certificate data itself has two statuses. The barriers and ranking function are *ghost witnesses*: they live only in the refinement proof. The energy budget $R = \widehat{V}(x_M, q)$ is *operational certificate data*: it is exposed across the motion–battery

boundary and consumed by another component at runtime. Only the latter kind must be projected to CertIfc and wired. A satisfactory definition of $\mathrm{CertRun}(\mathcal{D}) \to \mathrm{Run}(\mathcal{D})$ should therefore distinguish ghost refinement data from operational certificate ports rather than treating every witness as an interface value.

The two-sided sensing contract adds a second structural lesson. Calibration and sharpness are obligations of opposite variance on the same exported bound. This is further evidence that a CPS certificate is naturally a structured product of ordered/modal obligations, not one scalar resource.

A plausible long-term abstraction is an indexed or fibred system over a category Amb of admissible ambient worlds, with fibre $\mathrm{Run}(\mathcal{D})$ or $\mathrm{IntRun}(\mathcal{D})$ over $\mathcal{D}$, together with certified refinements and common interface projections. Lerman–Schmidt's networks of hybrid open systems provide a useful nearby comparison: their interconnection maps build hybrid systems from open subsystems [6]; the additional Sweedler ambition is to combine such wiring with theory/residual structure, measuring comonads, and certificate transport.

For the probabilistic fibre there is encouraging ambient structure: quasi-Borel spaces admit the concrete-sheaf presentation underlying their Grothendieck-quasitopos structure and support higher-order probability [8, 9]. This does not by itself establish that $[\mathrm{QBS}, \mathrm{QBS}]_{\mathrm{acc}}$ has all the duoidal and measuring-comonad structure required here; that remains a separate check.

13 What is established and what remains programme-level

The following pieces of the case study are explicit or standard local arguments:

- the finite stateful controller runner;
- a non-reversing hybrid plant with clocked control modes;
- the two-sided probabilistic estimator contract, with separate calibration and sharpness obligations and an explicit per-installation probability reading;
- local predicted-terminal-clearance barriers h_B, h_K, h_D, progress ranking L, operational energy potential $\widehat{V}$, and concrete parameters satisfying the local inequalities on the good-event branches;
- a genuinely load-bearing range-estimation problem: robust safety and a 15-second mission deadline cannot both be certified by any switching law over the four available modes using the beacon but not the range channel, uniformly over $Z \in [10, 12]$;
- the corrected motion–battery contract with a computable operational budget port $R = \widehat{V}(x_M, q)$, separated from ghost barrier/ranking witnesses;
- behavioural observability of collision, energy failure, and docking, together with the witness/property distinction.

The following remain architectural or theorem-level work:

- a precise indexed category or double category of ambient Sweedler worlds;
- the definition and composition theory of $\mathrm{CertRun}(\mathcal{D}) \to \mathrm{Run}(\mathcal{D})$;
- certified-interface extraction compatible with runner composition and with the predicate-transport adjunctions of Part V;

- general cross-ambient composition and its relation to IntRun and DOTS-style interfaces;
- compatibility of certification with the Beh : HRun $\to$ CofProf semantics and with trace/feedback;
- verification of the required duoidal/measuring-comonad hypotheses for $[\mathrm{QBS}, \mathrm{QBS}]_{\mathrm{acc}}$;
- existence and useful description of the final coalgebra of $\{T_M, R_{\mathrm{ext}}\}$ in the probabilistic setting used for relative behaviour.

A further engineering extension would close the motion–battery interaction into genuine feedback by allowing voltage sag or state of charge to limit admissible actuation. The present chapter uses series wiring $M \to B$; a bidirectional coupling would exercise the trace/feedback material of Part XI more directly.

14 Conclusion

The repaired robot is small enough that the local pieces can be calculated explicitly, but rich enough to exercise much of the Sweedler systems programme. Its estimator, controller, hybrid motion, and battery subsystem need not be forced into a single ontology. They retain native semantics, communicate through interfaces, and carry additional certificate data in a refinement layer over runners.

The worked example also sharpens the overall architecture. Safety, energy, and progress are not automatic annotations generated from implementation semantics; they are witnesses chosen above that semantics. Motion exports a quantitative budget which becomes an operational interface proof obligation for the battery, while barrier and ranking data remain ghost witnesses in the refinement layer. The uncertain wall position together with the deadline makes probabilistic sensing genuinely necessary: calibration protects safety while sharpness protects timeliness, and their risks compose differently. Behavioural abstraction preserves only properties expressible at the chosen interface, so collision and energy failure must be observable if their guarantees are to descend.

The resulting picture is

native heterogeneous semantics
$\downarrow$
stateful runners and open interface wiring
$\downarrow$
displayed/refined certificate data
$\downarrow$
closed hybrid execution
$\downarrow$
behavioural abstraction and end-to-end guarantee.

This provides a natural resting point for the enquiry. The case study now demonstrates a non-vacuous local synthesis while exposing, rather than concealing, the theorem-level work required for a general heterogeneous Sweedler framework for certified cyber–physical systems.

References

[1] S. Prajna, A. Jadbabaie, and G. J. Pappas, "A framework for worst-case and stochastic safety verification using barrier certificates," *IEEE Transactions on Automatic Control* 52(8), 1415–1428, 2007. DOI: 10.1109/TAC.2007.902736.

[2] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control* 62(8), 3861–3876, 2017. DOI: 10.1109/TAC.2016.2638961.

[3] K. H. Johansson, M. Egerstedt, J. Lygeros, and S. Sastry, "On the regularization of Zeno hybrid automata," *Systems & Control Letters* 38(3), 141–150, 1999. DOI: 10.1016/S0167-6911(99)00059-6.

[4] A. Benveniste, B. Caillaud, D. Nickovic, R. Passerone, J.-B. Raclet, P. Reinkemeier, A. L. Sangiovanni-Vincentelli, W. Damm, T. A. Henzinger, and K. G. Larsen, *Contracts for System Design*, *Foundations and Trends in Electronic Design Automation* 12(2–3), 124–400, 2018. DOI: 10.1561/1000000053.

[5] A. Platzer, "Differential dynamic logic for hybrid systems," *Journal of Automated Reasoning* 41, 143–189, 2008. DOI: 10.1007/s10817-008-9103-8.

[6] E. Lerman and J. Schmidt, "Networks of hybrid open systems," *Journal of Geometry and Physics* 149, 103582, 2020. DOI: 10.1016/j.geomphys.2019.103582.

[7] P.-A. Melliès and N. Zeilberger, "Functors are type refinement systems," *ACM SIGPLAN Notices* 50(1), 3–16, 2015. DOI: 10.1145/2775051.2676970.

[8] C. Heunen, O. Kammar, S. Staton, and H. Yang, "A convenient category for higher-order probability theory," *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 1–12, 2017. DOI: 10.1109/LICS.2017.8005137.

[9] C. Heunen, O. Kammar, S. Moss, A. Ścibior, S. Staton, M. Vákár, and H. Yang, "The semantic structure of quasi-Borel spaces: algebra, logic, and recursion," *Probabilistic Programming Languages, Semantics, and Systems (PPS 2018)*, 2018.