



André Dupke

Hamburg, August 7, 2026

scaletimedynamics.com

info@scaletimedynamics.com

[Orcid](#) | [Academia](#) | [YouTube](#) | [Facebook](#)

Version: 1.0

© 2026 André Dupke. Licensed under CC BY 4.0

DOI: doi.org/10.5281/zenodo.21833505

STT Framework Link: [Scale-Time Dynamics](#)

STAR DRIVE

Sampled Transport by Alias Ratcheting

A Lock-Window Pulse-Scheduling Architecture for Conventional Propulsion

STAR DRIVE - TECHNICAL SUPPLEMENT

This supplement carries the complete controller specification, the derivations and design rationale behind the scheduler's stated rules, and the extended deterministic tables that EN-002 summarizes. Nothing here modifies the main note; every definition has its canonical home in EN-002 and is referenced, not restated. Claim-status tags follow EN-002. Section references §n point into the main note; S§n are internal.

Contents

STAR DRIVE - TECHNICAL SUPPLEMENT	1
S§1 Scheduler rationale: ledger, threshold, & budget separation	3
S§2 Timing references, alignment, and the four-term attribution	5
S§3 The robust command optimizer: staged construction and cost functional	6
S§4 The lock-loss statistic as a trigger–recovery process	10
S§5 Realization alternatives and prediction-tolerance derivations	12
S§6 Thrust factorization: pathwise and ensemble readings	14
S§7 Extended A1/A2 tables	15

S§1 Scheduler rationale: ledger, threshold, & budget separation

Why the ledger is tick-level, not endpoint. The endpoint form

$$\Delta C_n = \max(0, \lfloor x_{n+1} - \delta_g \rfloor - M_n)$$

is an identity only if $x(t)$ is monotone on the interval. Under tick-level noise the phase can cross a threshold, fire, and retreat before the interval ends - reaching $x = 5.2 + \delta_g$ from $M_n = 4$, then ending at $x_{n+1} = 4.8 + \delta_g$ - whereupon the endpoint expression records no command although one was irrevocably issued. The monotone M of the tick-level rule keeps that command and correctly refuses to reissue threshold 5 on the next approach. This matters even though sustained backward motion is a declared fault: a command emitted before the fault must not vanish from the ledger. [D]

Why the integer part sits inside the max. Because $mn \in \mathbb{Z}$,

$$\lfloor x_{n+1} - \delta_g \rfloor = m(n+1) + \lfloor z_{n+1} - \delta_g \rfloor,$$

so with $\widetilde{M}_n = M_n - mn$ the per-interval rule is identically

$$\Delta C_n = \max(0, m + \lfloor z_{n+1} - \delta_g \rfloor - \widetilde{M}_n):$$

the reduced-lift form is an algebraic restatement of the full-lift rule, not a second counter. A rule that issues m commands unconditionally and adds fractional commits on top agrees with this whenever the lift advances - and issues phantom commands during a backward excursion exceeding one cycle. Counts, guard surface, and firing times must derive from one scheduler; falsifier 4 of §7.1 is the test most likely to fire if an implementation quietly reverts to two command streams. Throughout, $m = \lfloor W^* \rfloor$ is the fixed integer translation of the selected winding target (§3.2 element 2), constant within an operating mode; it is never the floor of the instantaneous map parameter. [D]

Why count failure is threshold identity, not the instantaneous ledger difference.

Budgeting on $D_r = C_r - C_r^{(0)}$ conflates timing with counting: a correctly indexed event ar-

iving slightly early makes $D_r = 1$ between its realized and reference crossing times and 0 thereafter, so an instantaneous test reports an extra event exactly when the early-event budget is supposed to be measuring a timing error - and then halts timing evaluation at the first mismatch, destroying the measurement it gates. Matching by the committed integer ℓ makes the separation exact: over the scored horizon, with $\mathcal{L}_R$ and $\mathcal{L}_R^{(0)}$ the threshold sets of run and causal deterministic reference,

$$N_{\text{extra}}^{\text{cmd}} = \left| \mathcal{L}_R \setminus \mathcal{L}_R^{(0)} \right|, \quad N_{\text{miss}}^{\text{cmd}} = \left| \mathcal{L}_R^{(0)} \setminus \mathcal{L}_R \right|,$$

and for every shared threshold the difference $t_\ell^{\text{cmd}} - t_\ell^{(0)}$ is purely a timing error at any magnitude. Because both ledgers commit consecutive thresholds from the same armed origin, the two sets are integer intervals and their symmetric difference is a tail: after the declared settlement interval the count test reduces exactly to the terminal comparison C_R vs $C_R^{(0)}$, and every discrepancy interior to the horizon is timing. A run that goes one command ahead early and is caught by the reference before the horizon ends commits the same threshold set and scores as a pure timing deviation. The former pairing tolerance survives only as an application timing-acceptance deadline $\tau_{\text{deadline}} \geq \tau_{\text{early}}$; it no longer defines event identity in any sense. Operational first passages are formulated on the full ledger C_r , since a tick-level z_r is not automatically the same algebraic quantity as the fiducial-indexed analysis object. [D]

Why the budgets do not reduce to one another.

The burst budget $P(\max(\Delta C_n - m) > 1) \leq \alpha_{\text{burst}}$ is not the extra-event budget: an unintended event landing $\overset{n}{\text{in}}$ an interval that should have carried none contributes one commitment, passes the burst test, and still leaves the ledger one command ahead. The endpoint condition $\sup \Delta z_n < 1$ is a conservative deterministic sizing screen only - an endpoint advance of $\overset{n}{\text{a}}$ full cycle need not produce a second crossing, and a noisy path can cross and retreat with endpoint advance below one. The early-event budget is not implied by any condition of the form $\sup e_n \geq \delta_g$: the intended event already sits on the guarded surface, so an arbitrarily small $\overset{n}{\text{positive}}$ trajectory error fires early without the accumulated error ever reaching δ_g . And the single-step quantile $Q_{1-\alpha}(\Delta z_n) < 1 - \delta_g$ sits below even the screens: converting a one-step quantile into a horizon-level statement requires a union bound or an extreme-value model, and since process noise, detector error, quantization, and delay are correlated across intervals, neither supremum is obtainable from the one-step distribution. Additive Gaussian noise admits no finite band excluding all excursions,

so both budgets are horizon-limited by construction. Timing is scored on matched pairs only; an unmatched event contributes no timing sample - consistent precisely because pairing is by identity and independent of timing, so no lead magnitude converts a timing error into a count error. [D]

S§2 Timing references, alignment, and the four-term attribution

The causal deterministic reference $t_\ell^{(0)}(\delta_g)$ is produced by the same realized fiducial sequence $\{T_n\}$, the same causal prediction rule $\hat{T}_{n|n}$, the same initial state, nominal parameters, and guard, with detector noise, process noise, quantization, and latency removed. Holding cadence and predictor fixed is what makes comparison against it isolate execution error. [D]

The stationary operational reference. Isolating what causal prediction costs needs a second reference, taken at the operating parameters actually commanded, not at the A1 anchor: A1 is fixed at $\kappa = 0.95$, $\delta_g = 0$, and the hold-optimal centre, so a difference measured against it would silently absorb the operational guard, a scheduled gain, a displaced setpoint, and the different stationary orbit these imply. Define $\tilde{\phi}_\ell^{\text{stat}}(\Omega^{\text{cmd}}, \kappa, \delta_g)$ - the event phases of the stationary known-period orbit at the candidate parameters, computed by the same deterministic orbit routine that produces A1 and A2. A period- q orbit is defined only up to a cyclic shift, so the reference needs an absolute alignment convention, fixed by the arming state: with $\ell_{\text{arm}} := M_{n_{\text{arm}}}$, application event j relates to threshold identity by

$$j(\ell) = \ell - \ell_{\text{arm}},$$

and the stationary orbit is cyclically aligned so its first post-arm crossing is $j = 1$, on the same lifted branch and from the same arming phase as the realized run. For a time-varying controller, the stationary reference for threshold ℓ is evaluated at the candidate active in the interval in which ℓ was committed; a setpoint move inside a scored superperiod is logged and reported as a separate setpoint-update residual, keeping the attribution unique. [D]

The attribution identity. Written entirely in phase (terms in seconds and cycles cannot be summed), converting through the declared $\tilde{\phi}$ of §3.1:

$$\begin{aligned}
\tilde{\phi}(t_\ell^{\text{del}}) - \tilde{\phi}_{j(\ell)}^{\text{app}} &= \underbrace{\tilde{\phi}(t_\ell^{\text{cmd}}) - \tilde{\phi}(t_\ell^{(0)})}_{\text{execution and noise}} \\
&+ \underbrace{\tilde{\phi}(t_\ell^{(0)}) - \tilde{\phi}_\ell^{\text{stat}}}_{\text{causal prediction and nonautonomy}} \\
&+ \underbrace{\tilde{\phi}_\ell^{\text{stat}} - \tilde{\phi}_{j(\ell)}^{\text{app}}}_{\text{stationary architecture / application mismatch}} \\
&+ \underbrace{\tilde{\phi}(t_\ell^{\text{del}}) - \tilde{\phi}(t_\ell^{\text{cmd}})}_{\text{actuator latency and jitter}}
\end{aligned}$$

four terms owned by the scheduler, the period predictor, the structural STAR floor of §2.5 at the operating point, and the actuator - all reported at Stage 1. The A1 anchor is retained as a separate diagnostic, $\tilde{\phi}_\ell^{\text{stat}} - \tilde{\phi}_\ell^{\text{A1}}$, measuring what guard, scheduled gain, and setpoint displacement cost relative to the zero-guard anchor reference - a useful quantity, but not prediction cost and not to be labelled as such. The safety budget of §3.2 element 3 remains in seconds as a first-passage statement on command times.

The guard's lead-time margin against the unguarded surface, $\tau_{g,n} = \delta_g / (f_{\text{gate},n} + u_n / \hat{T}_{n|n})$ in seconds ($\delta_g / (\Omega_n + u_n T_n / \hat{T}_{n|n})$ in external cycles), is valid on a fixed event-to-interval branch; where the guard carries a crossing past the next fiducial, the delay spans multiple intervals and is accumulated piecewise at their respective rates. Both the guarded and unguarded references are reported. At the guarded null the deterministic reference commits no thresholds, so no realized command can be classified as early: any command there is an extra-event failure, governed by α_{extra} or the dedicated α_{null} . [D]

S§3 The robust command optimizer: staged construction and cost functional

Objective separation. Robustness margin and timing accuracy are different objectives - §2.5 establishes that raising κ_n trades one against the other - and the primary metric is event-phase RMSE, not boundary distance. Margin therefore enters as a feasibility floor, never as the command law. Write the worst-case realized boundary distance

$$\mathcal{M}_n(\omega, \kappa) = \min_{\rho \in [\rho_n^-, \rho_n^+]} \min \left[\rho \omega - \Omega_{p/q}^-(\rho \kappa), \Omega_{p/q}^+(\rho \kappa) - \rho \omega \right]$$

(script $\mathcal{M}$, distinct from the ledger's M_n). [D]

Staged construction. Defining the margin ceiling over a set that already contains the floor would be circular and would leave the ceiling undefined exactly when the floor is set too high. The construction is therefore staged. First the pre-margin reachable set - window intersection, gain cap, slew limits, and hard rollout feasibility, but no margin condition:

$$\mathcal{G}_n = \left\{ (\omega, \kappa) : \omega \in \mathcal{J}_n(\kappa), \kappa \rho_n^+ \leq \kappa_{\text{safe}}, |\omega - \Omega_{n-1}^{\text{cmd}}| \leq s_\Omega, |\kappa - \kappa_{n-1}| \leq s_\kappa, \text{HF}_{n,H_c}(\omega, \kappa) = 1 \right\},$$

with $\mathcal{J}_n(\kappa)$ the robust window intersection of §3.2 element 4 and HF_{n,H_c} the explicit hard-feasibility predicate: the candidate's rollout delivers the target finite-horizon event count $K_{H_c}^{\text{pred}} = K_{H_c}^{\text{target}}$, respects the minimum pulse separation, the actuator rate and recovery limits of §3.2 element 2, and the application deadline τ_{deadline} , and the candidate's required coupling range $[\rho_n^- \kappa, \rho_n^+ \kappa]$ lies inside the A2-verified window library. Then the ceiling actually available on it, $\mathcal{M}_n^{\text{ceil}} = \max_{\mathcal{G}_n} \mathcal{M}_n$; and only then the feasible set

$$\mathcal{F}_n = \left\{ (\omega, \kappa) \in \mathcal{G}_n : \mathcal{M}_n(\omega, \kappa) \geq \mathcal{M}_{\min}^{\text{design}} \right\}.$$

The floor $\mathcal{M}_{\min}^{\text{design}} > 0$ is a fixed safety policy chosen once, before the mission, calibrated offline against a declared design envelope $\mathcal{D}$ of admissible cadence, prediction-band, and gain conditions: $\mathcal{M}_{\min}^{\text{design}} < \inf_{\zeta \in \mathcal{D}} \mathcal{M}^{\text{ceil}}(\zeta)$. The runtime ceiling is a diagnostic and fallback target, never an input that lowers the floor. If $\mathcal{M}_n^{\text{ceil}} < \mathcal{M}_{\min}^{\text{design}}$ then $\mathcal{F}_n = \emptyset$ and the controller takes the declared suspend/reacquire/fault path - it does not relax the floor, which would preserve non-emptiness by discarding the requirement the floor enforces. An adaptive floor is admissible only as a declared degraded mode with a hard lower bound, its own risk budget, stated entry and exit conditions, and mode-transition logging. $\mathcal{G}_n = \emptyset$ and $\mathcal{F}_n = \emptyset$ are distinct faults - no reachable point holds the window at all, versus none with the demanded margin - and only the second could, under a declared degraded mode, be relieved short of reacquisition. [D]

Command law.

$$(\Omega_n^{\text{cmd}}, \kappa_n) = * \operatorname{argmin}_{(\omega, \kappa) \in \mathcal{F}_n} J_n(\omega, \kappa), \quad J_n = \frac{R_n(\omega, \kappa)}{R_{\text{ref}}} + \lambda_\Omega \left(\frac{\omega - \Omega_{n-1}^{\text{cmd}}}{s_\Omega} \right)^2 + \lambda_\kappa \left(\frac{\kappa - \kappa_{n-1}}{s_\kappa} \right)^2 + \lambda_P \frac{P_{\text{control}}(\omega, \kappa)}{P_{\text{ref}}},$$

each term normalized dimensionless before weighting, with declared scales $R_{\text{ref}}, P_{\text{ref}}$ (naturally the rule-9 timing budget and the nominal control effort at the previous

setpoint). The gain enters J_n on the same footing as ω - never chosen first to maximize margin and then held fixed, which is the rejected "maximize the intersection" ordering. [D]

The timing term is a computable functional, not a name. R_n is declared with six pre-registered choices: the target it scores against (the application target, since the controller serves the mission, not map fidelity); nominal, expected, or worst-case in ρ ; the horizon H_c , at least one rational superperiod so a full event pattern is scored; the guard δ_g at which phases are evaluated, on the branch containing it; how the current orbit state is propagated; and whether candidate phases come from the A2 library, a local surrogate, or an online rollout. Over a horizon the uncertain object is a ratio trajectory: different sequences $\rho = (\rho_n, \dots, \rho_{n+H_c-1})$ from the same scalar range give different orbits, counts, phases, gains, branch crossings, and transients. The functional therefore ranges over a declared path set $\mathcal{P}_{n,H_c} = \{\rho : \rho_k \in [\rho_k^-, \rho_k^+], \text{ plus declared temporal constraints}\}$ - cadence slew, autocorrelation, monotone drift bounds; without them the set is needlessly pessimistic - with default

$$R_n(\omega, \kappa) = \sup_{\rho \in \mathcal{P}_{n,H_c}} \text{RMSE}_{\text{phase}} \left[\text{rollout}(x_n, \omega, \kappa, \rho, \delta_g; H_c) \right],$$

matching the worst-case convention of $\mathcal{M}_n$. The rollout needs a future setpoint policy: the declared default is move blocking (hold the candidate fixed across H_c); policy rollout (reapply the optimizer at each future update) is more faithful but recursively depends on future prediction bands and is a declared option, not the baseline. The single-ratio form $\max_{\rho} \text{RMSE}_{\text{phase}}$ is a frozen-ratio screening surrogate, not the worst case over a moving cadence. Replacing the supremum by an expectation under a declared predictive law is a legitimate alternative and a different controller, fixed before Stage 1. [D]

Hard feasibility precedes the soft metric. A rollout can deliver the wrong finite-horizon count, violate minimum pulse separation, breach actuator rate or recovery limits, or miss the application deadline; a timing RMSE over such a candidate is meaningless and optimizing it would trade away the architecture's own hard constraint. This is the predicate HF_{n,H_c} in the definition of $\mathcal{G}_n$ above; equivalently $J_n = +\infty$ on violators. Controller-family status. Until R_n , $\mathcal{P}_{n,H_c}$ and its temporal constraints, the future-setpoint policy, H_c , the phase source, the state-propagation method, the optimization solver or minimization grid, the weights, the normalizing scales, and the definition of P_{control} (named in J_n but deliberately left to the preregistration) are all fixed, this specifies a controller family: the feasible-set construction is [D], and no quoted setpoint inherits [S:A2] from the solver architec-

ture alone. This release fixes none of them; the unique operational instantiation belongs to the Stage-1 preregistration. [D]

Ties, update dynamics, and library confinement. Selection follows a declared total order, unique by construction: (i) minimum J_n ; (ii) among minimizers, maximum $\mathcal{M}_n$; (iii) minimum distance to the preferred nominal or phase-optimal operating point; (iv) minimum κ ; (v) minimum ω - equivalently, the fixed lexicographic index of the solver grid. These are ordinary setpoint updates, not mode transitions, but not exempt from a declared law. The setpoint law is rate-limited inside $\mathcal{G}_n$, before margin or cost - a pointwise re-solve can otherwise jump commands between intervals, move event phases, force crossings, excite tracking dynamics, invalidate the orbit target, and chatter when the binding ρ endpoint changes. Optimizing over $\mathcal{J}_n$ first and clipping to the slew limit afterwards is unsafe and is not what the staged construction does: the optimum may be unreachable while the previous command has left the current robust intersection, in which case "holding" violates the containment budget. When $\mathcal{F}_n = \emptyset$, holding is admissible only if the previous command still satisfies the current constraints; otherwise commitments suspend, ledgers are preserved, and reacquisition or fault handling begins. Candidates outside the A2-verified library fail HF_{n,H_c} and are excluded, not extrapolated; if that empties $\mathcal{G}_n$ the controller reports the empty-set fault rather than silently extending the library. [D]

Maximin as calibration diagnostic and fallback. $\mathcal{M}_n^{\text{ceil}}$ tells the designer how much margin is on the table before the floor is fixed, and doubles as a declared safety-priority fallback mode (command the $\mathcal{M}_n$ -maximizer over $\mathcal{G}_n$, under mode-transition discipline). One closed form serves as an analytic anchor: for fixed gain κ_n , symmetric band $\rho \in [1 - \varepsilon, 1 + \varepsilon]$, non-binding slew, and an integer tongue, the maximizer is

$$\Omega^{\text{cmd}} = m + \frac{\varepsilon \kappa_n}{2\pi}, \quad \text{margin} = \frac{\kappa_n(1 - \varepsilon^2)}{2\pi} - \varepsilon m,$$

biased above m because the binding case is $\rho = 1 - \varepsilon$, where the window narrows as the point moves down; the margin is non-negative exactly under the rule-8 tolerance condition. This is the fixed-gain static integer-window ceiling, not the general $\mathcal{M}_n^{\text{ceil}}$ (which also optimizes κ and may have slew or library constraints binding); it is an analytic regression test for the A2 optimizer. The gap it exposes is not academic: at $m = 1$, $\varepsilon = 0.03$ it gives $\Omega^{\text{cmd}} = 1.004536$ with margin 0.121061 against 0.116661 for the naive $\Omega^{\text{cmd}} = m$; the 2/3 example of Appendix A.2 shows a sixth of the available margin discarded by commanding the nominal centre. [S:A2][D]

Fallback ladder and mode transitions. In order: clamp or gain-schedule κ_n ; switch to the causal phase-step realization when prediction uncertainty exceeds the cap's reach; failing both, declare a fault and leave the operating window. A realization change is a mode transition, not a continuous update: switching can jump the lift across guarded surfaces, emit commands at the fiducial, change the event-phase pattern, and invalidate the map-fidelity target - so the switch suspends commitments, preserves both ledgers, executes a declared bumpless transfer, and requires reacquisition and re-arming before scored operation resumes. Without that discipline the fallback manufactures exactly the early or additional events the budgets exist to control. The mission-level budgets of §3.2 are guaranteed only on the event that T_n falls inside its band over the horizon: coverage must be simultaneous over the horizon, sequentially valid, or explicitly allocated

$$\sum_n \alpha_n \leq \alpha_{\text{mission}}$$

- a per-interval statement is not sufficient. Two-sided bounds are required because an over-long predicted period shrinks κ_{eff} and the window with it, destroying hold margin without threatening invertibility. The prediction residual $u_n(\rho_n - 1)$ is multiplicative and correlated with the servo state (§3.2); representing it as independent additive noise is a coarse convenience only, and Stage 1 retains the correlation. Since under persistence the prediction error is the cadence's own first difference, it is the same slew quantity rule 10 budgets, and the two are budgeted together. The gain cap is not free: buying prediction headroom at $\kappa = 0.80$ costs about 38 % of the 2/3 hold budget (Appendix A.2). [D] (coverage in practice:) [C]

S§4 The lock-loss statistic as a trigger-recovery process

"Loss of lock" is declared, not left to intuition; for a noisy nonautonomous map, plausible readings disagree, and a single infimum over OR-ed clauses cannot work - a deviation persisting L intervals would declare loss before reacquisition had any opportunity to succeed, so a recovery clause could never perform its distinction. Two quantities separate the roles. Let

$$d_n = \text{wrap}(x_n - x_n^{(0)}), \quad b_n = \text{round}(x_n - x_n^{(0)})$$

be the within-branch deviation from the causal deterministic reference and the integer branch offset wrapping discards, with the departure band constrained to $0 < \theta_{\text{exc}} < \frac{1}{2}$ and half-cycle ties rounded downward so $x_n - x_n^{(0)} = b_n + d_n$ holds exactly with $d_n \in (-\frac{1}{2}, \frac{1}{2}]$. Onset and decision are separated, because an L -sample run cannot be known before its last sample. Episode r 's excursion has

$$\tau_{\text{exc,on}}^{(r)} = \inf \{n : |d_k| > \theta_{\text{exc}} \text{ for } k = n, \dots, n + L - 1\}, \quad \tau_{\text{exc,dec}}^{(r)} = \tau_{\text{exc,on}}^{(r)} + L - 1;$$

reacquisition starts at the decision time, which by itself declares nothing. Loss is declared afterwards by two named events. A persistent branch offset,

$$\tau_b = \inf \{n : b_k \neq 0 \ \forall k = n, \dots, n + H_b - 1\}$$

(onset; its decision time is $\tau_b + H_b - 1$). And a recovery failure, evaluated over every excursion episode - a mission whose first excursion recovers and whose later one does not must not score as never having lost lock. Recovery is searched only after the trigger is causally available,

$$\tau_{\text{rec}}^{(r)} = \inf \{k > \tau_{\text{exc,dec}}^{(r)} : |d_k| \leq \theta_{\text{rec}}\};$$

episode r fails if $\tau_{\text{rec}}^{(r)} > \tau_{\text{exc,dec}}^{(r)} + H$; if it recovers, the next episode opens only after the recovered state, $\tau_{\text{exc,on}}^{(r+1)} = \inf \{n > \tau_{\text{rec}}^{(r)} : |d_k| > \theta_{\text{exc}} \text{ for } k = n, \dots, n + L - 1\}$, so excursions are counted as separate episodes rather than merged. Then

$$\tau_{\text{rec-fail}} = \inf_r \{\tau_{\text{exc,dec}}^{(r)} + H : \text{episode } r \text{ fails}\}, \quad \tau_{\text{loss}} = \min(\tau_b, \tau_{\text{rec-fail}}).$$

Finite records. A record of finite length cannot always decide an episode: if the recovery horizon H extends past the last sample, or a branch-offset run is still active at the final sample, the episode is undetermined, not failed. The evaluator reports such episodes explicitly and returns a verdict of undetermined rather than lost - a truncated record must neither manufacture a lock loss nor conceal one, and a mission-level $\alpha_{\text{lock,mission}}$ estimated over truncated records must account for the undetermined mass rather than silently scoring it as survival.

Parameters $(\theta_{\text{exc}}, \theta_{\text{rec}}, L, H, H_b)$ are preregistered, with $\theta_{\text{rec}} < \theta_{\text{exc}}$ providing hysteresis. Reporting convention: onset and decision times are logged for every event class - excursion $(\tau_{\text{exc,on}}, \tau_{\text{exc,dec}})$, branch-offset loss $(\tau_b, \tau_b + H_b - 1)$, and recovery failure (onset $\tau_{\text{exc,on}}^{(r)}$, decision $\tau_{\text{exc,dec}}^{(r)} + H)$ - and mission budgets that must be evaluated causally use the decision times. The branch criterion is stated on b_n persisting rather than on an instantaneous threshold-set difference, consistently with S§1: one ledger temporarily running ahead is a timing difference, not a count failure. Between them, d_n and b_n lose no branch information. The budget $P(\tau_{\text{loss}} < N) \leq \alpha_{\text{lock,mission}}$ cannot be derived from A1 or A2 - it involves slew, noise, state history, acquisition basin, prediction error, and latency together - and is a Stage-1 measurement. Where EN-002 says a window "cannot be guaranteed", the claim is frozen-map parameter containment and nothing more. [D] (dynamic lock retention:) [C]

S§5 Realization alternatives and prediction-tolerance derivations

Causal reduced map. Under the causal realization the operational reduced map carries the scheduled, prediction-scaled parameters,

$$z_{n+1} = z_n + (\rho_n \Omega_n^{\text{cmd}} - m) - \frac{\kappa_{\text{eff},n}}{2\pi} \sin(2\pi(z_n - \epsilon_n)) + \xi_n, \quad \kappa_{\text{eff},n} = \rho_n \kappa_n,$$

with $m = \lfloor W^* \rfloor$ fixed per operating mode (§3.2 element 2), coinciding with the fixed-parameter analysis form exactly when $\hat{T}_{n|n} = T_n$ and the gain is unscheduled; analysis quantities (windows, orbits, A1/A2) are computed on the first form, the implementation runs the second. Sampling the tick-level accumulator at the next fiducial returns $x_{n+1} = x_n + f_{\text{gate},n} T_n + u_n T_n / \hat{T}_{n|n} + \xi_n$: the circle map is recovered exactly only for a known-period or constant reference, and otherwise prediction error appears as the fluctuating effective gain plus a process disturbance. The measured κ -scan gives roughly 7 % variation in $w_{2/3}$ per 5 % of gain, so period-prediction error jitters window boundaries directly. [D][S:A2]

Integer-window tolerance, exact form.

The integer tongue boundaries $\Omega = m \pm \kappa/2\pi$ are affine in κ with half-width proportional to it, so ρ cancels from the half-width and the rate-blended robust condition reduces to $m\epsilon/(1 - \epsilon^2) \leq \kappa_n/2\pi$. Writing $c = \kappa_n/(2\pi m)$,

$$\varepsilon_{\max} = \frac{\sqrt{1 + 4c^2} - 1}{2c} = \frac{2c}{\sqrt{1 + 4c^2} + 1},$$

of which the familiar $\varepsilon < \kappa_n/(2\pi m)$ is the small- ε approximation (14.79 % against 15.12 % at $m = 1, \kappa_n = 0.95$). Consequences: the guarded null's window containment is exactly invariant ($\Omega^{\text{cmd}} = 0$ scales to 0), leaving the effective-gain/invertibility cap as its only prediction limit in the rate-blended realization; sensitivity grows linearly in m , so integer translation buys thrust level and pays in prediction robustness; the first-order interior screen is $\varepsilon \lesssim (w_{p/q}/2)/\Omega_{p/q}^c$, the radial displacement being proportional to the operating point itself. [D][S:A2]

Phase-step realization. With u_n applied whole at the fiducial there is no $u_n/\hat{T}_{n|n}$ term, hence no gain scaling: invertibility reduces to $\kappa_n < 1$ and that failure mode disappears. Prediction error remains in the commanded ratio - $f_{\text{gate},n} = \Omega_n^{\text{cmd}}/\hat{T}_{n|n}$, so $\Omega_n = \rho_n \Omega_n^{\text{cmd}}$ - and the phase-step map is

$$x_{n+1} = x_n + \rho_n \Omega_n^{\text{cmd}} - (\kappa_n/2\pi)\sin(2\pi(x_n - \epsilon_n)) + \xi_n,$$

with robust condition

$$\Omega_n^{\text{cmd}} \in \bigcap_{\rho \in [\rho_n^-, \rho_n^+]} \left[\frac{\Omega_{p/q}^-(\kappa_n)}{\rho}, \frac{\Omega_{p/q}^+(\kappa_n)}{\rho} \right],$$

nonempty exactly when $\varepsilon \leq w_{p/q}/(2\Omega_{p/q}^c)$ - the rate-blended screen, exact here because the window no longer moves with ρ ; at the guarded null both $\Omega = 0$ and κ_n are invariant, so the null is not prediction-limited within the admitted positive- ρ domain. It is a valid fall-back for gain-scaling and invertibility failures, not for loss of the robust intersection: an empty intersection from radial Ω displacement stays empty. A threshold can be stepped over with no crossing time; such commands are assigned the fiducial instant, injecting a discontinuity into event phases. Causality, crossing-time exactness, and event-phase smoothness are three axes on which the realizations differ, and causality may reverse the baseline choice where the cadence is strongly unpredictable - which is why Stage 1 runs both. The intra-interval servo (correction recomputed at tick rate from a held estimate) changes the stroboscopic map itself unless the detector genuinely updates at that rate; an implementation states its realization, and Stage 1 reports the others as sensitivity cases. [D]

Tick budget and convergence. The tick must resolve the claimed timing quantities: $\Delta t (m + 1)/T_{\min}$ small against the smallest claimed event-phase difference, tick quantization entering the guard budget through Δz_n . The Stage-1 convergence test refines Δt and requires crossing times, command counts, capture statistics, and event-phase RMSE to converge rather than depend on the tick; it presupposes the declared realization and does not substitute for it. Queue overflow, delayed execution, and dropped commands propagate into K_N^{del} - accounting events, not hardware footnotes. [D]

S§6 Thrust factorization: pathwise and ensemble readings

Writing $\mathbb{E}[\vec{F}] = \vec{\mu}_\chi \lambda_{\text{del}}$ with the ordinary mark mean is not generally valid for the physical model: the impulse bit may depend on event timing, reservoir pressure, preceding pulses, saturation, valve state, or the cadence itself - the dependences EN-002 insists on elsewhere - and the product form silently assumes marks identically distributed and independent of the counting process. The assumption-light asymptotic objects are pathwise: the event-conditioned sample-path mean

$$\vec{\mu}_\chi^{\text{event}} = \lim_N K_N^{\text{del},-1} \sum_j \vec{\chi}_j$$

and the pathwise long-run force $\overline{\vec{F}} = \lambda_{\text{del}} \vec{\mu}_\chi^{\text{event}}$, both sample-path limits. An ensemble statement is legitimate by exactly two routes, declared in any Stage-1 report: ergodicity, under which the long-run limits are deterministic and $\mathbb{E}[\vec{F}]$ is redundant rather than wrong; or a Palm expectation $\vec{\mu}_\chi^0 = \mathbb{E}^0[\vec{\chi}]$ of the stationary marked point process, for which $\mathbb{E}[\vec{F}] = \lambda_{\text{del}} \vec{\mu}_\chi^0$ holds by construction. The ordinary mark mean coincides with either only under the i.i.d.-and-count-independent assumption - part of the ideal model, and the same assumption behind the variance split of §3.4, whose clean command/shot separation is itself ideal: rate-dependent actuator behaviour (blow-down, thermal throttling) correlates $\vec{\chi}_j$ with the schedule and adds a cross-covariance to be calibrated out. [D]

Boundary case: at the guarded null $\lambda_{\text{del}} = 0$ and $\vec{\mu}_\chi^{\text{event}}$ is undefined (an average over an empty set); in the deterministic ideal null $\overline{\vec{F}} = 0$ directly from $\vec{J}_N = \vec{0}$. Under stochastic null operation the finite-horizon identity remains fully operative: a false event is an α_{null} failure, but its impulse enters $\vec{J}_N$ and K_N^{del} normally, so the null's thrust is zero only to the extent the budget is met. Constant- T reductions are pathwise to match ($r_{\text{del}} = \lim K_N^{\text{del}}/N$, ensemble $r_{\text{del}}^{\text{ens}}$ named separately and equal only under ergodicity), the monotone-capture equivalence $\hat{\lambda}_{C,N} \approx \widehat{W}_N/T$ holds when backward excursions carry negligible rate and is a

Stage-1 verification target, and $W_\infty = \overline{W}$ requires ergodicity, not stationarity - a sample-path estimate does not converge to an ensemble expectation by stationarity alone. Recovering W_0 in the final link of the §3.3 chain requires detector noise, quantization, tick resolution, latency, and every implementation error to vanish together, not σ_ξ alone. [D]

S§7 Extended A1/A2 tables

Event phases under integer translation (window centres, $\kappa = 0.95$, $\delta_g = 0$; total winding $m + \tilde{p}/\tilde{q}$ emits $m\tilde{q} + \tilde{p}$ events per superperiod): [S:A1]

W	Events/superperiod	Uniform gap	RMS shape deviation	% of gap
1/1, 2/1, 3/1	m per interval	$1/m$	$\leq 3.3 \times 10^{-17}$	0
$m + 1/2$ at $\Omega = m + \frac{1}{2}$	$2m + 1$ per 2	$2/(2m + 1)$	$\leq 1.2 \times 10^{-16}$	0
$1 + 1/3$	4	0.75	0.022452	3.0 %
$1 + 1/4$	5	0.8	0.041612	5.2 %
$1 + 2/3$	5	0.6	0.018450	3.1 %
$1 + 3/4$	7	0.5714	0.030227	5.3 %
$2 + 1/3$	7	0.4286	0.013614	3.2 %
$2 + 1/4$	9	0.4444	0.023580	5.3 %
$2 + 2/3$	8	0.375	0.011976	3.2 %
$2 + 3/4$	11	0.3636	0.019359	5.3 %

The uniformity criterion at $m \geq 1$: events inside interval n are spaced $1/\Delta_n$, so the train is uniform iff the per-interval advance is constant. Integer windows have a fixed-point reduced orbit (advance exactly m at every interior $\bar{\Omega}$; verified to machine precision out to $\bar{\Omega} = 0.145$). On the $1/2$ plateau the two advances satisfy $D_0 + D_1 = 1$; equality forces $\sin(2\pi z_0) = 0$ and hence $\bar{\Omega} = 1/2$ exactly - uniformity is a property of the operating point, not the plateau. Uniformity imposes $\tilde{q} - 1$ conditions on one parameter at $m \geq 1$, so only $\tilde{q} = 2$ is generically solvable, and that solution is the symmetry centre. Deviation, absolute

and relative: integer translation shortens the gap and reduces absolute deviation without improving the relative one. [S:A1][D]

Half-integer detuning sensitivity ($1/2$ window, $w/2 = 0.033591$): [S:A1]

$\delta\Omega^{\text{win}}$	Advances D_0, D_1 ($m = 1$)	dev $m = 1$	dev $m = 2$	dev $m = 3$	Relative
0	1.5, 1.5	0	0	0	0
± 0.010	1.478719, 1.521281	0.004314	0.002513	0.001776	$\approx 0.65\%$
± 0.020	1.455788, 1.544212	0.008648	0.005146	0.003662	$\approx 1.3\%$
± 0.030	1.426421, 1.573579	0.013729	0.008410	0.006038	$\approx 2.1\%$
± 0.033590 (hold limit)	-	0.017583	0.011063	0.008002	2.6 – 2.8 %

Relative deviation is nearly independent of m and grows as $\approx 0.63 |\delta\Omega^{\text{win}}|$, steepening near the edge; integer windows show no such sensitivity. Off-centre operation on the half-integer plateau costs timing in addition to margin - a second reason to budget $\delta\Omega^{\text{win}}$ tightly. [S:A1]

Interior uniform-point structure ($\kappa = 0.95$): root-solving the $2/3$ gap-difference condition gives $\Omega^* = 0.655080852$, gaps 1.5,1.5 to residual 4×10^{-16} , sitting $+0.418 (w/2)$ off centre with hold margin 0.007726 against 0.013275 at the centre - uniformity costs 42 % of the static hold budget.

The $3/4$ window is searched by branch-covering global minimisation, not by a single bounded call. The objective is piecewise smooth in Ω , kinking wherever an event changes its interval assignment, and the exact indicator of such a change is the cyclic signed distance of the event closest to a fiducial - under fiducial-linear phase, interval boundaries sit at integer phase, so an assignment change is precisely a zero of that signed distance (a counts-based signature is blind to changes that permute cyclically and must not be used). The signed indicator is rotation-invariant because event-list indices themselves rotate at a fiducial crossing. Every local minimum of the unsigned sampled margin localizes a candidate; its signed event–fiducial equality is then root-solved, and only those solved zeros par-

tition the window. A bounded minimisation runs on each branch; all branch minima are recorded as competing minima; isolated curvature spikes are retained as a diagnostic cross-check and never used as authoritative cuts. At 801 sample points the 3/4 window resolves into two branches - a boundary an earlier coarser sampling missed - and the global minimum is unchanged at 0.062432 at $\Omega = 0.722371$ (window endpoints 0.0898, 0.0719), interior and stationary ($|\text{central gradient}| < 10^{-3}$, curvature $\approx +950$), with no zero detected, matching the dimension count ($\tilde{p} - 1$ conditions on one parameter at $m = 0$).

At $m \in \{1, 2\}$, 401-point uniform scans - labelled as scans in the artifact, not as branch-covering minimisations - of the whole 2/3, 3/4, and 1/3 windows find no uniform point: minima 0.018166, 0.030065, 0.021969 at $m = 1$ and 0.011911, 0.019331, 0.013530 at $m = 2$. Global nonexistence is not claimed anywhere: the 3/4 result is global over its declared branch partition, and the $m \geq 1$ results are scan minima. Hold-optimal, phase-optimal, and level-selection are three different choices (rule 9). [S:A1][D]

Coupling continuation, full scan (w and shape deviation at the 2/3 and 3/4 centres): the nine-row table over $\kappa \in \{0.40, 0.50, 0.60, 0.70, 0.75, 0.80, 0.85, 0.90, 0.95\}$ is reproduced in the evidence pack; the five-row summary appears in EN-002 Appendix A.2, and all four columns increase strictly across the sampled values. These are scan outputs at reduced accuracy; the anchor values at $\kappa = 0.95$ are canonical from the A1/A2 boundary solver (EN-002 Appendix A.2, EN-002-EP).

Intermediate rows: $w_{2/3} = 0.004237/0.011240/0.013690/0.019490$ and $\text{dev} = 0.016061/0.020542/0.021424/0.022858$ at $\kappa = 0.50/0.70/0.75/0.85$; $w_{3/4} = 0.001221/0.004364/0.005633/0.008893$ and $\text{dev} = 0.040255/0.053918/0.056993/0.062636$ at the same values. [S:A2]

Boundary detection, tangency, and the orbit pair. A window edge is a saddle-node bifurcation: the stable and unstable period- q orbits merge and $(f^q)'(x^*) = 1$. The dense-grid check (4×10^5 x -points) is exploratory detection and bracketing only: it finds the multiplier at the surviving root equal to 1 within 7.5×10^{-3} worst case (2.6×10^{-5} best) - commensurate with a grid check, not with six-decimal boundary values. Final boundaries therefore come from the direct solve of the saddle-node system $f_\Omega^q(x) - x - p = 0$, $\partial_x f_\Omega^q(x) - 1 = 0$, with residual tolerances recorded for both equations (released run: $\leq 2.4 \times 10^{-14}$ on all fourteen edges). Detection caveat carried into the artifact specification: a pure sign-change test misses roots landing exactly on grid nodes (as the 1/2 orbit's points $z = 0, \frac{1}{2}$ do at its centre) and misses tangent roots by construction - a root candidate

is detected either by a sign change between adjacent nodes or by a node with $|g| \leq \text{tol}$, then refined and independently verified.

The **orbit pair itself is verified inside every window**, at the centre and at $\pm$ half the half-width: all roots of g are refined, grouped into cycles under the map, and classified by multiplier, and the attracting branch is identified by iterating from a generic seed and matching the limit set (match distance $< 10^{-6}$ in every case). Every window carries a stable/unstable pair with the attractor on the stable branch - centre multipliers 0.05 ($m/1$), 0.0975 ($1/2$), 0.1434 ($1/3, 2/3$), 0.1834 ($1/4, 3/4$), against unstable partners 1.95, 1.9524, 1.9093, 1.8652 respectively. [S:A1][D]

Guard-law branch mechanics. The affine per-event relation $e_j(\delta_g) = e_j^{(0)} + \delta_g/\Delta_{n(j)}$ holds while event j stays in interval $n(j)$; it leaves at the breakpoint $\delta_{g,j}^{\text{bp}} = x_{n+1} - \ell_j$, whereupon the superperiod may be cyclically re-indexed and C, V must be recomputed - hence the piecewise-quadratic global law with the branch table of EN-002 Appendix A.2, on the declared guard domain $0 \leq \delta_g < 1$. Because breakpoints are absolute in δ_g , the branch structure is enumerated by walking from one breakpoint to the next until the domain is exhausted: at $\kappa = 0.95$ this gives four branches for the $\tilde{q} = 3$ cases, five for $\tilde{q} = 4$, two for $m + 1/2$ and one for $m/1$, each with its own validity interval, its own (R_0, C, V) relative to its own reference guard, and an in-branch verification of the closed form (worst error $< 5 \times 10^{-7}$). Each branch carries its own shape-optimal candidate $-C/V$, admissible only where $C < 0$ and the candidate falls inside that branch and the domain. A branch change need not change the value: the uniform families have breakpoints too, but their trains stay exactly uniform because their zero-error law is globally valid - one event per superperiod is spaced by periodicity whatever interval carries it, and constant-advance orbits shift all events identically.

Beyond-branch checks: evaluating the quoted (C, V) outside their branch for $1 + 1/3$ and $1 + 1/4$ would give 0.023368/0.027445 and 0.043100/0.048391 at $\delta_g = 0.1/0.3$ - overestimates of 3–10 % against the directly measured 0.022577/0.024837 and 0.041993/0.044174. [S:A2][D]