STARDRIVE

A Lock-Window Pulse-Scheduling Architecture for Conventional Propulsion

André Dupke

Scale-Time Dynamics Research

André Dupke
Hamburg, August 7, 2026

scaletimedynamics.com
info@scaletimedynamics.com
Orcid | Academia | YouTube | Facebook

Version: 1.0

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

ABSTRACT

Every propulsion system decomposes into a reservoir that momentum is exchanged with, a coupling that performs each exchange, and a schedule that decides when. This note is a design theory for the third component. The STAR Drive (Sampled Transport by Alias Ratcheting) is a conventional, nonlinear phase-locked pulse-scheduling architecture: an external-reference control loop whose stroboscopic reduction is the sine circle map, with the map parameter Ω identified physically as the gate-to-reference frequency ratio. All directionality lives in the momentum quantum $\vec{\chi}$ supplied by the reservoir coupling; the gate contributes scalar rate structure, a computable static hold interval, and an irreversible command ledger.

Two properties of the map govern the design. A lock window's location is not its rational label: the operating point is the measured window centre, and at the anchor coupling some parameter values $\Omega = p/q$ lie outside their own windows. And a rational winding fixes the event count per superperiod, not the event phases within it: the locked train is exactly uniform only at named operating points and otherwise carries a deterministic, guard-dependent timing floor. Every command is a guarded threshold crossing of the single full lifted

phase, committed to a monotone ledger. Because the reference period must be predicted causally, prediction error scales the realized operating point and coupling gain jointly in the baseline rate-blended realization, $(\Omega_n, \kappa_{\mathrm{eff},n}) = \rho_n(\Omega_n^{\mathrm{cmd}}, \kappa_n)$, so the command is selected by a margin-constrained robust rule with separate mission-level budgets on leaving the stationary tongue, losing invertibility, and losing dynamic lock.

What the architecture predicts directly is a commanded event-rate staircase; a thrust staircase is an additional actuator-dependent claim, and because the staircase rides the command stream, momentum-discriminating controls and a metrology standard for phase-based thrust stands are specified. No new momentum source is proposed and no engineering advantage is claimed in advance: whether the coupled gate outperforms a direct digital p/q scheduler, PWM, or an instrumented digital PLL is the architecture's first research question, judged on a preregistered unwrapped, order-paired event-phase RMSE under a hard event-count constraint, beside named falsifiers, invalidators, and a retirement criterion.

Contents

CLAIM-STATUS TAGS

Substantive claims carry a bracketed status tag, following the discipline of the parent framework [1]. [D] - definition, convention, or internal derivation (including derivations valid within the ideal model of §3.2). [S:En] - supported by experiment n of the Scale-Time Theory 11.0 Simulation Evidence Pack, release 1.0, fixed seed 20599373 [2]; constructive classical simulation evidence only. [S:A1] - deterministic analysis A1: window boundaries, centres, and unguarded locked event phases of the noiseless sine circle map at the anchor coupling $\kappa = 0.95$ (Appendix A). [S:A2] - deterministic analysis A2: guard-branch structure, κ-continuation, robust window intersections and prediction tolerances, and the margin-ceiling/feasible-set computations with their analytic setpoint anchors (Appendix A) - not, by itself, a unique operational command (§3.2 element 4, S§3). Stage 0b is closed in this release: A1 and A2 ship as their own routines with verification tests in the Stage 0b artifact package (a1_routine.py, a2_routine.py, test_stage0b.py; 36 regression tests passing, of which one class recomputes A1/A2 quantities in-process rather than reading the stored results, so a source regression cannot pass unnoticed), with regenerated results.json and a 286-entry manuscript manifest in which every [S:A1] and [S:A2] value in this note traces to its generating artifact and matches. The environment is pinned (pyproject.toml, requirements.lock), and a clean-room regeneration on the same recorded runtime and platform reproduces results.json bit-identically apart from its timestamp. Integration of the artifact package into the released evidence pack accompanies the DOI minting of this note. [C] - conjecture or untested extension, paired with a validation stage in §7; every application of the ideal model to physical hardware and every claimed engineering benefit is [C]. [I] - interpretive device carrying no theoretical weight.

The abstract is a tag-free summary: every claim in it inherits the status of its counterpart in the body.

1. INTRODUCTION AND DESIGN PRINCIPLES

Phase-tracking engineering rests on two questions: where is the phase, and will it stay there. The first is answered by estimation, the second by coupling, and the two are not interchangeable [1][4]. Propulsion has the same pair, dressed differently: how much impulse has been delivered, and will the delivery schedule hold. This note takes the correspondence literally. Under discrete sampling, an accumulated phase-like quantity splits into a bounded wrapped part and an unbounded integer:

$$\tilde{\Phi}(t_N); = ; \Phi(t_N); + ; 2\pi K_N, \qquad K_N \in \mathbb{Z}$$

For an external observer of a wrapped stream that integer is a hidden alias branch: EN-001 proves it absent from the data, recoverable only by assumption [3]. For a controller integrating its own lifted oscillator it is a constructed winding counter: directly available, exact, and maintained. The two are related mathematically and opposite epistemically, and the ladder is kept terminologically separate throughout - alias branch k for the hidden ambiguity; winding counters $K_N^{\mathrm{full}}, K_N^{\mathrm{frac}}$ for the controller's constructed counts (§3.2); command ledger C_N for the irreversible record of issued commands; verified firing count K_N^{del} only after actuator feedback. Delivered impulse is then an event count with three obligations: something must pay for each event (the reservoir), something must schedule it (the gate), and something must count it - commanded by construction, delivered by verification. The founding accounting identity is

$$\text{thrust} = (\text{momentum per event}) \times (\text{event rate})$$

The first factor belongs to the reservoir and coupling, is governed by ordinary conservation, and is not touched here; this note is a design theory for the second. The name records the mechanism: a subcritical coupling ratchets the wrapped phase mismatch between gate and reference into deterministic integer advances, each committed advance releasing one impulse quantum - "alias branch" remaining reserved throughout for the externally unidentifiable integer. [D]

Position, honestly stated. STAR is not a propulsion source. It is a phase-locked scheduling architecture for pulsed actuators - a recognizable relative of injection locking [5] and first-order PLLs [6]. Its proposed benefit is controlled event timing, a computable hold interval, disturbance recovery, and a command-side ledger; whether those properties outperform conventional digital and phase-locked scheduling under matched conditions is the architecture's first research question, not its premise, and the primary metric is accordingly a timing metric (§3.5). This note consumes three prior results: EN-001's non-injectivity theorem and four disambiguating assumption classes [3]; experiments E1–E5 of the Simulation Evidence Pack [2]; and the frame-servo formulation of [4]. Their propulsion-side assembly, and every claimed benefit of it, is [C] with a stage in §7. [D][C]

Three principles close the introduction; the architecture is forced by them, not assembled by preference. [D]

I. The gate must be actuated. A schedule that does not enter the generator of the dynamics cannot change momentum: for any law $\dot{x} = f(x, t)$ independent of the observation schedule, $\Delta p_{\text{sched}} \equiv 0$ (the schedule enters the observation map, not the generator). A gate that only observes is a stroboscope; a gate that commands is a scheduler. The first required component is an actuator inside the loop.

II. Direction lives in the coupling, not in the gate. The gate state and its winding rate are scalars, invariant under spatial inversion; no thrust direction can be constructed from them. All directionality enters through the per-event momentum quantum $\vec{\chi}$, set entirely by the reservoir coupling. Gate engineering and thrust-vector engineering are independent problems; thrust reversal is implemented on the $\vec{\chi}$ side (opposed actuators, gimballing), never by driving the gate backwards, and the gate's design domain is $\Omega \geq 0$, $W_0 \geq 0$.

III. The gate must be coupled. At $\kappa = 0$, exact recurrence occurs only at a measure-zero set of rational ratios with no restoring basin (§2.4): recurrent, not stable. At $\kappa > 0$ those points inflate into finite windows with a computable static hold interval, so the coupling - a servo correction bounded per interval by $\kappa/2\pi$, a classical control action with no exotic content - is a required layer, not an optimization. At the evidence anchor it is not weak: $\kappa = 0.95$ is near-critical invertible, just below the gain at which invertibility is lost, with consequences for window width, event-phase uniformity, and margin (§2). One scoping remark, developed in §3.5: Principle III concerns holding a schedule against an analog, physical detuning; where the schedule is synthetic - referenced to the controller's own clock - a digital counter has no detuning and needs no basin, and the coupled gate must earn its place against that baseline.

2. SUBSTRATE - WHAT THE ARCHITECTURE CONSUMES

This section is self-contained and establishes exactly six facts the architecture uses: wrapped observation is non-injective; sampling depth improves estimation but provides no holding; a coupled circle map opens finite rational lock windows; window location is distinct from the rotation-number label; the rotation number fixes event count, not event timing; and static hold, acquisition, and tracking are three different quantities.

2.1 Wrapped observation and the ratio–branch circularity

A sampled observer of a periodic process with true normalized advance ν (cycles/sample) records only $\nu_{\text{alias}} = \nu - \text{round}(\nu) \in (-\frac{1}{2}, \frac{1}{2}]$, leaving the admissible set

$$\mathscr{A}(\nu_{\text{alias}}) = \nu_{\text{alias}} + k \mid k \in \mathbb{Z}.$$

EN-001 proves the map non-injective in the strong sense: the likelihood is invariant under integer shifts of ν at every record length and noise level; the branch k is absent from the data and must be supplied by assumption - a known start, a second sample rate, a fiducial, or an external prior [3]. [D]

Proposition 1 (Ratio–branch circularity). Since ν is non-identifiable, so is the sampling ratio $R_A = 1/|\nu|$: each admissible branch implies a different ratio. An observer cannot determine its own sampling ratio from the wrapped stream alone; any regime claim is an assumption supplied from outside the data. □ *[D]*

The constructive reading: a controller knows its own cadence by construction - it sets it - but not the unaliased frequency of the process it reads, and therefore not the ratio. Where the mode frequency is external, the ratio requires an independently justified band limit, prior, fiducial, continuity assumption, or multi-rate constraint (§4). A phase-based thrust-stand sensor is in the same position and worse placed within it: the unaliased frequency of the motion it reads is the very quantity under measurement (§5). [D]

2.2 The sampling regimes

The estimation-side substrate model is minimal: a deterministic rotating phasor, additive Gaussian readout noise of per-quadrature scale η, stroboscopic sampling at R_A samples per cycle, initial phase unknown [2]. Four canonical regimes, stated for the first alias interval $0 < |\nu| \le 1$: [D]

Regime	Condition	Behaviour	Measured anchor
Structural identity	$R_A = 2$	Opposite rotations are the same sampled relation; sign unrecoverable in principle	Overlap 1.000, discrimination error 0.500 exact; direction recovery $P \approx 0.51, 0.50$ at $\eta = 0.35, 0.7$ **[S:E1][S:E4]**
Statistical zone	$2 < R_A \lesssim R_{\mathrm{stat}}(\eta)$	Noise-broadened ambiguity	$P \ge 0.95$ by $R_A \approx 2.16$ ($\eta = 0.35$), 2.60 ($\eta = 0.7$); error $< 5\,\%$ by $R_A \approx 2.21$ ($\eta = 0.5$) **[S:E1][S:E4]**
Alias reversal	$1 < R_A < 2$	Direction **confidently reversed**: a branch misassignment read as data	$\nu = 0.7 \Rightarrow \nu_{\mathrm{alias}} = -0.3$ **[S:E1]**
Stroboscopic freeze	$R_A = 1$	Motion renders as stillness; event slips invisible	Demonstrated at $\Delta\varphi = 2\pi$/sample **[S:E1]**

The regimes recur across higher branches (sign inversion on every band $k+\frac{1}{2} < |\nu| < k + 1$; freezes at $\nu \in \mathbb{Z}$). Reversal and freeze are not merely low-information - one is confidently wrong, the other manufactures stillness - and both return as design exclusions (§4, rules 3–4) and measurement hazards (§5). [D][S:E1]

2.3 Resolution is smooth and structure-relative

Across $R_A \in [2,512]$, per-cycle phase-estimation error falls smoothly at the Cramér–Rao limit, $\varepsilon_\phi \sim C_\varepsilon\, \eta / \sqrt{R_A}$ (measured slope -0.504 vs $-1/2$; residuals $\leq 5.5\,\%$): no depth value - integer, half-integer, or dyadic - carries special stability. Depth is a smooth dial, never a resonance; this concerns the R_A axis only, while the detuning axis of §2.4 has rational plateau structure by design. The depth needed to resolve a residual structure $\Delta\varphi_{\mathrm{res}}$ at noise η is budgeted, not discovered: $R_A^* \sim C_{\mathscr{E}}\, \eta^2 / |\Delta\varphi_{\mathrm{res}}|^2$, with anchors $R^* \approx 20.9,\ 99.9,\ 330.1$ for $\eta = 0.2,\ 0.45,\ 0.8$ at $\Delta\varphi_{\mathrm{res}} = 0.1257$ rad/cycle. [S:E2][S:E3][D]

2.4 The coupled lock: windows, centres, and two detunings

Deep sampling answers where is the phase; it does not answer will it stay there. Without coupling, phase-overlay recurrence is exact only at zero-width rational parameter values - a dense set of measure zero, with no attracting basin and no noise robustness [S:E2][S:E5]. Finite coupling changes the problem. On the lifted phase $x_n \in \mathbb{R}$,

$$x_{n+1} \;=\; x_n \;+\; \Omega \;+\; \frac{\kappa}{2\pi}\,\sin\!\big(2\pi x_n\big) \;+\; \xi_n, \qquad \theta_n \;=\; x_n \bmod 1,$$

the standard sine circle map of mode-locking theory [8][9], with Ω the bare advance per interval (physically a frequency ratio, §3.1), $\kappa \geq 0$ the coupling gain, ξ_n process noise; the restoring-servo convention of [4] is the same family under reflection. The map is invertible for $0 \leq \kappa < 1$; this note's design domain is restricted to that range, anchored by evidence at $\kappa = 0.95$. Three winding quantities are kept distinct: the deterministic winding number $W_0(\Omega, \kappa)$ (noiseless limit), the finite-run estimate $\widehat{W}_N = (x_N - x_0)/N$, and the stochastic mean winding $\overline{W}$ where the limit exists. Two exact symmetries are load-bearing: translation, $W_0(\Omega + 1, \kappa) = W_0(\Omega, \kappa) + 1$, making integer windows exact translates; and reflection, $W_0(1 - \Omega, \kappa) = 1 - W_0(\Omega, \kappa)$, making p/q and $(q-p)/q$ windows mirror images about $\Omega = 1/2$. [D]

At $\kappa > 0$, finite lock windows open in Ω wherever W_0 is rational. Three objects must be kept apart: the rotation number p/q (the plateau's value - it fixes the thrust level and nothing else); the window $[\Omega^-_{p/q}, \Omega^+_{p/q}]$ with width $w_{p/q}$ and centre $\Omega^{\mathrm{c}}_{p/q} = \frac{1}{2}(\Omega^- + \Omega^+)$; and the two detunings $\delta\Omega^{\mathrm{rate}}_{p/q} = \Omega - \frac{p}{q}$ (rate mismatch) and $\delta\Omega^{\mathrm{win}}_{p/q} = \Omega - \Omega^{\mathrm{c}}_{p/q}$ (window-relative).

Every hold condition in this note is a statement about $\delta\Omega^{\mathrm{win}}$. The two coincide only where $\Omega^{\mathrm{c}}_{p/q} = p/q$ - exact for integer windows (translation) and for $1/2$ (reflection), and otherwise not generic. Measured at $\kappa = 0.95$ (widths from the E5 scan [S:E5]; boundaries and centres from A1 [S:A1]):

Winding p/q	**Window** $[\Omega^-, \Omega^+]$	**Width** $w_{p/q}$	**Centre** $\Omega^c_{p/q}$	$\Omega^c - p/q$
0/1	[−0.151197, +0.151197]	0.302394	0.000000	0 (exact)
1/4	[0.274710, 0.287946]	0.013236	0.281328	+0.031328
1/3	[0.337193, 0.363743]	0.026550	0.350468	+0.017135
1/2	[0.466409, 0.533591]	0.067181	0.500000	0 (exact)
2/3	[0.636257, 0.662807]	0.026550	0.649532	−0.017135
3/4	[0.712054, 0.725290]	0.013236	0.718672	−0.031328
1/1	[0.848803, 1.151197]	0.302394	1.000000	0 (exact)

A1 reproduces every released E5 width to within the E5 scan resolution ($\Delta\Omega \approx 4.2 \times 10^{-4}$; locked fraction 0.449) and confirms the deterministic integer-lock half-width $\kappa/2\pi = 0.151197$. This is not a refinement. For the interior windows the displacement $|\Omega^c - p/q|$ exceeds the half-width, so the parameter value $\Omega = p/q$ lies outside its own window: commanding $\Omega = 1/4$ locks to $W_0 = 5/24$, and $\Omega = 1/3$ to $W_0 = 8/25$ - a stable lock at the wrong thrust level (classification and plateau confirmation: A1 artifact). Both $\Omega^c_{p/q}$ and $w_{p/q}$ are κ-dependent and must come from a map scan at the operating coupling, never from the rational label. Across the measured family, widths are ordered by denominator, $w_{1/2} > w_{1/3} = w_{2/3} > w_{1/4} = w_{3/4}$ (the equalities are the reflection symmetry); a strict ordering over all rationals is not established and is a Stage-1 scan question. [S:E5][S:A1] [D]

2.5 Event phases inside the superperiod

p/q fixes the event count per superperiod, not generally the event phases inside it.

On the plateau the gate runs a stable period-q orbit whose per-interval advances are unequal, so the p crossings per superperiod are not generally uniformly spaced - and event phase is this note's primary performance variable (§3.5). The phases are not even a property of the map alone: they additionally require the declared intra-interval realization of

§3.2, under which the locked train is exactly computable. Under that realization, exact uniformity occurs throughout integer windows (fixed-point orbit, constant advance m, verified to machine precision), at the symmetry centre of half-integer windows only (off-centre, the lock survives and the spacing does not: relative deviation grows as $\approx 0.63\,|\delta\Omega^{\mathrm{win}}|$, reaching 2.6–2.8 % of the gap at the hold limit), and trivially for pure $1/\tilde{q}$ operation at $m = 0$ (one event per superperiod, not preserved under integer translation). Other measured operating points carry deterministic shape errors of roughly 3 % of the nominal gap at denominator three and 5 % at denominator four, nearly independent of the integer part m: [S:A1][D]

Operating point (window centres, $\kappa = 0.95$)	**Events per superperiod**	**RMS shape deviation (% of gap)**
$m/1$, any m (anywhere in window)	m per interval	0 (exact, guard-independent)
$m + 1/2$ at $\Omega = m + \frac{1}{2}$ exactly	$2m + 1$ per 2	0 (exact, guard-independent)
$1/\tilde{q}$ at $m = 0$ (anywhere in window)	1 per $\tilde{q}$	0 (trivially exact)
$\tilde{q} = 3$ family (2/3, 1 + 1/3, 1 + 2/3, 2 + 1/3, 2 + 2/3)	$m\tilde{q} + \tilde{p}$	3.0–3.2 %
$\tilde{q} = 4$ family (3/4, 1 + 1/4, 1 + 3/4, 2 + 1/4, 2 + 3/4)	$m\tilde{q} + \tilde{p}$	5.2–5.3 %

The error is guard-dependent and piecewise quadratic in δ_g on fixed event-to-interval branches, $\mathrm{RMSE}_{\mathrm{shape}}(\delta_g)^2 = \mathrm{RMSE}_{\mathrm{shape}}(0)^2 + 2\delta_g C + \delta_g^2 V$, with branch breakpoints where an event moves to the next interval; at 2/3 the law is exactly affine, $0.023952 + 0.120\,\delta_g$, so a routine guard $\delta_g = 0.1$ raises the floor by half again. The uniform families have $C = V = 0$ and stay exactly uniform at every guard. Operational timing floors must therefore be quoted at the operational guard, and the guard itself is fixed by the first-passage budgets of §3.2, not chosen for shape. One isolated off-centre uniform point exists in the 2/3 window at $m = 0$ ($\Omega^* = 0.655081$), costing 42 % of the hold half-width; no $m \geq 1$ counterpart was found in bounded scans, as the dimension count predicts (S§7). On the half-integer plateau the hold-optimal and phase-optimal points coincide at the centre - unlike 2/3, where they are 0.418 half-widths apart (rule 9). Along the coupling axis the trade runs the other way: raising κ widens the window and worsens the intrinsic nonuniformity, both strictly across nine sampled values $\kappa \in [0.40, 0.95]$ (Appendix A.2); coupling gain is not a parameter to be maximized. Detailed translated schedules, phase-optimal points, and guard coefficients: Appendix A and S§7. [S:A1][S:A2][D]

2.6 Hold, acquisition, and tracking are three different questions

The windows of §2.4 are static hold intervals - detuning ranges over which the lock exists for stationary Ω. Two further questions are not answered by them:

$$\text{static hold} \neq \text{initial acquisition} \neq \text{dynamic tracking}.$$

Acquisition (capture probability, capture time, pull-in) depends additionally on initial phase, noise, sweep rate, and delay [6]. Tracking - whether lock survives a moving Ω - depends on how fast it moves: per-instant window membership does not imply lock, since the servo's restoring authority per interval is bounded by $\kappa/2\pi$. The quasi-static screen $|\Omega_{n+1} - \Omega_n| \ll w_{p/q}$ is a sizing heuristic, not a proven condition [C]; trackable slew, modulation bandwidth, and pull-out probability are Stage-1 measurements. Tracking matters most, because a drifting physical cadence is the only regime in which the coupled gate might beat a digital scheduler (§3.5). Inside a window the lock is noise-confined over the measured horizon: at drive noise 0.03 cycles/step, a 500-member ensemble at the $1/2$ centre holds its lifted-phase spread at 0.302 cycles after 400 steps, while an incommensurate ensemble diffuses to 0.913 cycles and keeps growing (free diffusion would cross the half-cycle count-loss criterion in ≈ 278 steps). These are finite-run data; the hypothesis that long-term count error is governed by rare window escapes [10] is a Stage-1 first-exit question [S:E5][C]. The compact statement, inherited from [4]: oversampling estimates; coupling protects. [D]

3. THE STAR ARCHITECTURE

3.1 The physical loop and its reduction to the circle map

All phases are in cycles, frequencies in cycles per second. This subsection states the fixed-parameter reference loop; §3.2 element 4 replaces f_{gate} and κ with the scheduled $f_{\text{gate},n}$, κ_n that causal implementation requires. Per controller interval of duration T_n: [D]

$$\phi_{n+1} = \phi_n + f_{\text{ext},n} T_n, \qquad y_n = \text{wrap}\big(\phi_n - \psi_n + \epsilon_n\big), \qquad \psi_{n+1} = \psi_n + f_{\text{gate}} T_n + u(y_n) + \xi_n,$$

external reference phase, phase detector with readout noise ϵ_n, and gate oscillator with process noise ξ_n and the sinusoidal servo $u(y) = (\kappa/2\pi)\sin(2\pi y)$ - Adler's injection-locking form [5][6]. The interval is measured, not assumed: $T_n = t^{\text{fid}}_{n+1} - t^{\text{fid}}_n$ is the realized gap between external fiducials, $f_{\text{ext},n} := 1/T_n$ the interval-averaged reference frequency, and $T_{\min}$ the shortest admissible interval. Sampling stroboscopically at the reference cadence and lifting the gate phase gives

$$x_{n+1} \;=\; x_n + \Omega_n - \frac{\kappa}{2\pi}\sin\big(2\pi(x_n - \epsilon_n)\big) + \xi_n, \qquad \Omega_n \;=\; f_{\text{gate}} T_n \;=\; \frac{f_{\text{gate}}}{f_{\text{ext},n}},$$

the substrate map in the reflected convention: the map parameter is physically a frequency ratio, carrying the two detunings of §2.4 as $\delta\Omega^{\text{win}}_{p/q,n} = (f_{\text{gate}} - \Omega^{\text{c}}_{p/q} f_{\text{ext},n})T_n$. Three declarations complete the reduction. External phase between fiducials is defined, not assumed: the preregistered default is fiducial-linear phase, $\tilde{\phi}(t) = n + (t - t^{\text{fid}}_n)/T_n$ - unwrapped, estimator-free, scheduler-independent, scorable retrospectively from the fiducial record; the declared alternative is a continuously estimated, branch-tracked physical phase, at the price of admitting estimator bandwidth and latency into the primary metric. Whichever is chosen applies identically to STAR and every baseline. Cycle-boundary identification is itself an assumption: strobing once per cycle presupposes a per-cycle fiducial or an already branch-tracked reference (EN-001 Methods C or A/D) - a declared architectural requirement, not an implementation detail. The counter runs on the lifted gate phase x_n: its committed windings are actuator cycles ($W_0 = p/q$ means p firings per q external cycles); the phase-error lift counts slips relative to the reference and is logged as a diagnostic, never fed to the actuator. [D]

3.2 The gated-actuator elements

The drive comprises six elements. Their complete specification - including all derivations, worked failure cases, and the full controller construction - is S§1–S§5; this section defines the machine. [D]

1. Reference cadence.

The measured fiducial interval T_n (time-varying, with floor $T_{\min}$) in the physical-detuning regime; the controller clock in synthetic operation, where per §3.5 the coupled gate claims no advantage.

2. Gate oscillator with integer–fractional decomposition on a reduced lift.

The integer part is fixed by the selected winding target, not by the instantaneous map parameter: with target $W^\star = m + \tilde{p}/\tilde{q}$, set $m := \lfloor W^\star \rfloor$, held constant for the duration of the selected operating mode; define the reduced lift $z_n = x_n - mn$ and the translated parameter $\bar{\Omega}_n := \Omega_n - m \in \mathbb{R}$, not constrained to $[0,1)$. The $1/1$ window of §2.4 spans the integer $\Omega = 1$ asymmetrically yet is one $m = 1$ tongue throughout: ordinary motion within a tongue - including across integer Ω - changes nothing, while a change of the selected tongue family changes m, is a re-decomposition event, and triggers re-arming (element 5). The reduced lift evolves by the conjugate circle map (with scheduled, prediction-scaled parameters in the causal realization, S§5). The decomposition is bookkeeping, not a second command stream: all commands issue from the single full-lift rule of element 3; z is retained because it keeps the servo argument and accumulator bounded and because the map analysis of §2 lives on it. The fractional target $\tilde{p}/\tilde{q}$ is selected by its measured window centre, never its rational label (§2.4). Actuator requirements follow from the worst-case interval: $f_{\text{act, max}} \geq (m+1)/T_{\min}$, plus a declared minimum pulse separation, recovery budget, and queue-or-drop policy with latency accounting.

3. Firing rule: one guarded event surface on the full lift.

Every command, without exception, is the committed crossing of $x = \ell + \delta_g$, $\ell \in \mathbb{Z}$, with declared guard domain $0 \leq \delta_g < 1$, by the full lifted phase - the integer part m produces its m commands per interval as crossings of the same surface, under the same guard, into the same ledger. With M_r the highest committed threshold, evaluated at each controller tick r:

$$L_r = \lfloor x_r - \delta_g \rfloor, \qquad \Delta C_r = \max\big(0, L_r - M_r\big), \qquad M_{r+1} = \max\big(M_r, L_r\big),$$

with C_N the resulting irreversible command ledger: a threshold is committed once and only once, and a committed command is never recalled. The ledger updates on first passages, not interval endpoints - under tick-level noise the phase can cross, fire, and retreat before the interval ends, and the endpoint form $\max(0, \lfloor x_{n+1} - \delta_g \rfloor - M_n)$ is an identity only for monotone intra-interval motion; it is retained in the deterministic analysis only (rationale and worked cases: S§1). Commands are matched to the deterministic reference by threshold identity ℓ: a shared threshold with displaced timing is a timing error; a threshold present in only one schedule is a count error. No tolerance mediates between the two categories, and no timing magnitude converts one into the other. Count, burst, early-event, and null failures are budgeted as separate finite-horizon first-passage quantities - none follows from a one-step quantile (S§1):

$$P\left(N_{\text{extra}}^{\text{cmd}} \geq 1\right) \leq \alpha_{\text{extra}}, \quad P\left(N_{\text{miss}}^{\text{cmd}} \geq 1\right) \leq \alpha_{\text{miss}}, \quad P\left(\max_n(\Delta C_n - m) > 1\right) \leq \alpha_{\text{burst}}, \quad P\left(\min_\ell \left[t_\ell^{\text{cmd}} - t_\ell^{(0)}(\delta_g)\right] < -\tau_{\text{early}}\right) \leq \alpha_{\text{early}}$$

where $N_{\text{extra}}^{\text{cmd}}, N_{\text{miss}}^{\text{cmd}}$ count unmatched thresholds against the causal deterministic reference (same fiducials, same predictor, same initial state and parameters, stochastic errors removed), and the early-event budget runs over thresholds committed in both schedules. The guard buys a computable lead-time margin $\tau_{g,n} = \delta_g / \left(f_{\text{gate},n} + u_n / \widehat{T}_{n|n}\right)$ against the unguarded surface, valid on a fixed event-to-interval branch. Timing is scored on matched pairs only; the full four-term attribution of delivered-event phase error (execution, causal prediction, stationary architecture, actuator latency) and its alignment conventions are S§2. [D] (budget values and attainment: Stage 1) [C]

4. Event generator: the declared digital realization.

The stroboscopic map supplies counts per interval, not pulse times within it, and different realizations sharing the same map produce different event phases - so the realization is part of the architecture. Baseline (rate-blended servo): the scheduler runs the lifted accumulator at fixed tick Δt; detector and servo update once per interval at the fiducial; the correction $u_n = -(\kappa_n/2\pi)\sin(2\pi(z_n - \epsilon_n))$ is applied as a constant frequency offset on zero-order hold:

$$x_{r+1} \;=\; x_r + \left(f_{\text{gate},n} + \frac{u_n}{\widehat{T}_{n|n}}\right)\Delta t + \xi_r, \qquad t_r \in \left[t_n^{\text{fid}},\, t_{n+1}^{\text{fid}}\right).$$

The period must be predicted, not known: T_n is unavailable at t_n^{fid}, so the scheduler uses a causal one-interval prediction $\widehat{T}_{n|n}$ (declared default: persistence, $\widehat{T}_{n|n} = T_{n-1}$). With $\rho_n = T_n / \widehat{T}_{n|n}$ and the command law $f_{\mathrm{gate},n} = \Omega_n^{\mathrm{cmd}} / \widehat{T}_{n|n}$, the realized map parameters are jointly scaled:

$$\Omega_n \;=\; \rho_n\, \Omega_n^{\mathrm{cmd}}, \qquad \kappa_{\mathrm{eff},n} \;=\; \rho_n\, \kappa_n.$$

Prediction error therefore moves the system radially through the (Ω, κ) plane: it displaces the detuning as well as the gain, and for a narrow interior window it can leave the intended tongue long before it threatens invertibility ($\kappa_{\mathrm{eff},n} < 1$, i.e. $\widehat{T}_{n|n} > \kappa_n T_n$ - at the anchor gain, a period underestimate of about $5\,\%$ already violates it). The prediction error is multiplicative, not additive: the residual $u_n(\rho_n - 1)$ scales with the servo action, does not vanish merely because the gate is locked (only the symmetry-centred points $\Omega = m$ and $m+\frac{1}{2}$ have $u_n \equiv 0$ - the same condition that makes their trains uniform), and continuously modulates an established orbit; Stage 1 must retain the correlation.

The robust requirement is consequently two-dimensional. Candidate commands must remain inside the selected stationary tongue for all admitted prediction ratios, satisfy the invertibility cap, remain reachable under the setpoint slew limits, keep their rollout feasible against the hard event-count, pulse-separation, and actuator constraints, and exceed a preregistered worst-case margin floor $\mathscr{M}_{\min}^{\mathrm{design}}$ (fixed offline against a declared design envelope, never re-selected at runtime):

$$\Omega_n^{\mathrm{cmd}} \;\in\; \bigcap_{\rho \in [\rho_n^-,\, \rho_n^+]} \left[\frac{\Omega_{p/q}^-(\kappa_n \rho)}{\rho}, \frac{\Omega_{p/q}^+(\kappa_n \rho)}{\rho} \right], \qquad \kappa_n \rho_n^+ < 1,$$

with two-sided prediction bounds $T_{n|n}^{\pm}$ at stated coverage and the gain capped at $\kappa_n \le \kappa_{\mathrm{safe}}\, \widehat{T}_{n|n} / T_{n|n}^+$. Among feasible candidates the controller minimizes a declared prediction-horizon cost dominated by predicted event-phase error, with rate-limited setpoint updates, deterministic tie-breaking, and candidates confined to the A2-verified window library (excluded, not extrapolated, outside it). An empty feasible set is a fault, not a tuning inconvenience: the controller does not relax the margin floor; it holds only if the previous command still satisfies the current robust constraints, and otherwise suspends new commitments, preserves both ledgers, and enters reacquisition or fault handling. The complete staged construction ($\mathscr{G}_n$, $\mathscr{M}_n^{\mathrm{ceil}}$, $\mathscr{F}_n$), the cost functional and its rollout path set, the tie-break order, the maximin diagnostic with its closed-form integer-window anchor, and the fallback ladder are specified in S§3. What A2 verifies is the feasible-set architecture, the window intersections, the margin ceiling, and the analytic anchors; a unique operational command exists only once the Stage-1 preregistration fixes the remaining declared choices of S§3, and until then this release specifies a controller family. Containment is budgeted over the mission, not per interval, by three distinct quantities - leaving the invertible domain, leaving the stationary tongue of the frozen map, and losing dynamic lock:

$$P\Big(\max_{n<N} \kappa_{\mathrm{eff},n} > \kappa_{\mathrm{safe}}\Big) \le \alpha_{\kappa,\mathrm{mission}}, \qquad P\Big(\exists n < N:\ \Omega_n \notin \big[\Omega^{\pm}_{p/q}(\kappa_{\mathrm{eff},n})\big]\Big) \le \alpha_{\mathrm{param,mission}}, \qquad P\big(\tau_{\mathrm{loss}} < N\big) \le \alpha_{\mathrm{lock,mission}}$$

For narrow interior windows the tongue budget binds first (Appendix A.2). Frozen-map containment is necessary but not sufficient for lock retention: parameters can move faster than the loop follows, so the third budget is a Stage-1 measurement, with loss of lock declared as a trigger–recovery process - a persistent integer branch offset, or a failed recovery within a declared horizon, evaluated over every excursion episode, with onset and decision times logged separately (full definitions: S§4). [D] (dynamic lock retention, coverage in practice:) [C]

In the deterministic case the lift advances linearly within each interval and guarded crossing times are exact,

$$t_{n,j} \;=\; t_n^{\mathrm{fid}} + T_n\,\frac{(\ell_j + \delta_g) - x_n}{\Omega_n + u_n T_n / \widehat{T}_{n|n}} \qquad (\xi_r \equiv 0),$$

well-posed whenever the intra-interval rate is positive; three levels are named throughout - the stationary known-period reference (A1's model), the deterministic causal-predictor realization, and the stochastic tick-level first passage, which has no closed form and is located numerically at tick resolution. Declared alternative (phase-step servo): u_n applied instantaneously at the fiducial. This removes prediction error from the coupling amplitude ($\kappa_{\mathrm{eff}} = \kappa_n$, so the gain-scaling failure mode disappears) but not from the commanded ratio ($\Omega_n = \rho_n \Omega_n^{\mathrm{cmd}}$ as before); it is a valid fallback for invertibility failures, not for loss of the robust intersection, it steps thresholds discontinuously (such commands are assigned the fiducial instant), and the measured comparison (Appendix A.2) shows it helps strictly where invertibility binds and changes little, not always favourably, where drift binds. A realization change is a mode transition - commitments suspend, ledgers are preserved, bumpless transfer executes, reacquisition and re-arming precede scored operation (S§5).

The controller tick must resolve the claimed timing quantities; a tick-convergence test is a required Stage-1 item, and command latency τ_{lat} is measured and logged separately. [D]

5. Initialization and arming.

The counted run begins only after acquire-before-arm: establish the reference and its cycle-boundary source; enter the selected window at its declared operating point (measured centre in stationary operation; the rule-9 phase-optimal point where a timing budget dominates; the rule-11 command point under causal uncertainty); wait through a declared acquisition transient; initialize the commitment threshold $M_{n_{\text{arm}}} = \lfloor x_{n_{\text{arm}}} - \delta_g \rfloor$; arm the actuator; begin counting. The $0/1$ guarded null holds no-fire only after acquisition and arming, and under noise only up to a budget - an extra-event budget α_{null}, not an early-event one, since the null's reference schedule commits no thresholds at all. [D]

6. Ledgers.

Four counters, coinciding only under named conditions: the winding counters $K_N^{\text{full}} = \lfloor x_N \rfloor - \lfloor x_0 \rfloor$ and K_N^{frac} (either can decrease after a backward excursion); the monotone command ledger C_N; and the verified firing count K_N^{del} from actuator feedback. Each command carries its threshold identifier ℓ, and feedback returns it where hardware permits, so dropped and duplicated events are counted separately (N_{drop}, N_{dup}); the net error $E_r^{\text{del}} = K_r^{\text{del}} - C_r$ is a cross-check only, since one drop cancels one duplicate. Scheduler count errors are scored by threshold identity, actuator errors through $N_{\text{drop}}/N_{\text{dup}}$, command timing on t_ℓ^{cmd}, delivered timing on $t_\ell^{\text{del}} = t_\ell^{\text{cmd}} + \tau_{\text{lat},\ell}$. In monotone captured operation $C_N = K_N^{\text{full}} = mN + K_N^{\text{frac}}$ up to boundary effects; divergences are diagnostics. [D]

The ideal STAR model.

Quantitative statements of §§3.3–3.4 marked [D] hold within this explicitly idealized model, and their application to hardware is [C]: constant impulse bit $\vec{\chi}$; perfect execution ($K_N^{\text{del}} = C_N$); no latency or actuator memory; the full-lift guarded scheduler after completed acquire-before-arm. Physically, delivered impulse is

$$\vec{J}_N = \sum_{j=1}^{K_N^{\text{del}}} \vec{\chi}_j,$$

and the gate provides an exact command ledger which becomes an onboard estimate of delivered impulse only through a calibrated impulse-bit model and verified actuator response. [D]

3.3 The thrust factorization through event rates

Rates are defined per physical time and finite-horizon first: over N intervals of total duration

$$T_N^{\Sigma} = \sum_{n<N} T_n,$$

the pathwise rates are $\hat{\lambda}_{\mathrm{del},N} = K_N^{\mathrm{del}} / T_N^{\Sigma}$ and $\hat{\lambda}_{C,N} = C_N / T_N^{\Sigma}$, and the finite-horizon mean thrust of any run is $\overline{\vec{F}}_N = \vec{J}_N / T_N^{\Sigma}$. At any horizon the factorization is an identity on the event-weighted realized impulse bit $\overline{\vec{\chi}}_N = \vec{J}_N / K_N^{\mathrm{del}}$:

$$\overline{\vec{F}}_N \;=\; \hat{\lambda}_{\mathrm{del},N}\, \overline{\vec{\chi}}_N$$

Asymptotic rates are invoked only where the limits exist under stated conditions; the pathwise long-run force factors through the event-conditioned sample-path mean, $\overline{\vec{F}} = \lambda_{\mathrm{del}}\, \vec{\mu}_{\chi}^{\,\mathrm{event}}$, and an ensemble $\mathbb{E}[\vec{F}]$ is legitimate only under ergodicity or a Palm-expectation reading - the ordinary mark mean silently assumes marks i.i.d. and independent of the counting process, which is exactly the timing–impulse correlation this note refuses to erase (details, including the guarded-null boundary case: S§6). In noiseless monotone captured operation at constant T everything collapses to the exact ideal form

$$\langle \vec{F} \rangle \;=\; \frac{\vec{\chi}}{T}\, W_0(\Omega, \kappa)$$

- the accounting identity of §1 made quantitative: the substrate supplies the design theory of W_0; the reservoir coupling supplies $\vec{\chi}$; the factors are independently auditable, and if $\vec{\chi} = 0$ the ideal drive produces a controlled zero while the gate still winds (§5 turns that separability into the control structure). As a description of a physical drive the chain is the Stage-1–3 hypothesis, tested link by link: $\hat{\lambda}_{\mathrm{del},N} \approx \hat{\lambda}_{C,N}$ (execution), $\hat{\lambda}_{C,N} \approx \widehat{W}_N / T$ (monotone capture), $\widehat{W}_N / T \to \overline{W} / T$ (stationarity and ergodicity), and $\overline{W}/T \to W_0/T$ as $(\sigma_\xi, \sigma_\epsilon, \Delta t, \tau_{\mathrm{lat}}, \ldots) \to 0$; the first three finite-horizon and comparable within one run at constant T (the monotone-capture equivalence is a Stage-1 verification target, not an assumed identity); only the last two links are limits, each with its own condition. [D] (physical applicability:) [C]

3.4 Locked operation: what the window buys

Operated inside a rational window, the ideal drive holds the thrust level $F_{p/q} = (|\vec{\chi}|/T)(m + \tilde{p}/\tilde{q})$ against any stationary detuning with $|\delta\Omega^{\mathrm{win}}_{p/q}| < w_{p/q}/2$; the level is set by the rotation number, the operating point by the measured centre, and the two are distinct (§2.4).

The $0/1$ window is the guarded null - no-fire within the half-width $\kappa/2\pi$ after arming, false events budgeted by α_{null} - and the $m \geq 1$ integer points are its exact translates, the widest firing operating points. What is established, and by what, is separated deliberately: [D] [S:E5][S:A1] (assembled loop:) [C]

Property	Static result (A1/A2, E5)	Dynamic status
Event rate	p/q exactly, throughout the stationary tongue	Realized rate under moving cadence: Stage-1 measurement
Hold interval	Deterministic window widths and centres, measured	Necessary, not sufficient, for lock retention
Event phases	Operating-point- and guard-dependent floors (§2.5)	Measured primary metric (§3.5)
Slip probability	Not supplied by static windows	Stage-1 first-exit statistic

Delivered impulse inherits the gate's count statistics: finite-horizon confinement versus immediate free diffusion is demonstrated (§2.6); the long-term slip hypothesis is Stage 1. Under the ideal i.i.d.-marks model the impulse covariance splits as

$$\mathrm{Cov}(\vec{J}_N) = \mathbb{E}[K_N^{\mathrm{del}}]\mathrm{Cov}(\vec{\chi}) + \mathrm{Var}(K_N^{\mathrm{del}})\, \vec{\mu}_\chi \vec{\mu}_\chi^{\mathsf{T}}$$

- the count component is the architecture's; the shot term is the actuator's, and rate-dependent actuator behaviour adds a cross-covariance that must be calibrated out (S§6). Countability is architectural: the ledger advances on a deterministic timetable in lock, so the schedule audits itself; the impulse estimate is as good as the actuator calibration. [D] (estimate fidelity on hardware: Stage 3) [C]

3.5 Relatives, baselines, and the primary metric

The gate is recognizably a first-order PLL: Adler's equation in the loop [5], the sine circle map in reduction [8][9], standard acquisition-versus-hold distinctions [6]. That kinship defines the burden of proof. Where the schedule is synthetic - referenced to the controller's own clock - direct digital methods (a p/q event scheduler, DDS phase accumulator, pulse-density modulation, fractional divider [7]) already generate exact rational schedules with exact counts and no analog phase diffusion; against that baseline the coupled gate offers nothing, and this note claims nothing for it there. The candidate regime is a physical detuning: firing that must acquire and hold an external or analog cadence - spin phase, orbital phase, a structural-vibration node, a drifting analog reference - where Ω is unknown and time-varying and hold, acquisition, basin recovery, and graceful degradation are the currencies. Whether the coupled gate outperforms a well-engineered, fully instrumented digital PLL there is an open empirical question, posed as such: [D][C]

Scheduler	Reference source	Response to physical cadence drift	Event-phase structure	Characteristic failure mode
Direct digital p/q	Own clock	None (open-loop w.r.t. cadence)	Exact synthetic grid	Silent drift against physical cadence
PWM / pulse-density	Own clock	None	Duty-cycle grid, quantization jitter	Same, plus pattern noise
Instrumented digital PLL + event counter (slip detection and telemetry matched to STAR's)	External, tracked	Linear-loop tracking, cycle slips	Loop-filter dependent	Cycle slips under disturbance, logged by its own monitor

STAR rate-blended (baseline)	External, tracked	Basin + robust setpoint rule	Deterministic, guard-dependent floors (§2.5)	Tongue exit under prediction error
STAR phase-step (sensitivity case)	External, tracked	Basin; no gain scaling	Discontinuous at fiducials	Ratio drift unmitigated

Targets. A comparison needs a target, and §2.5 forbids assuming there is only one. Phase is measured in the unwrapped external coordinate $\tilde{\phi}$ (§3.1), applied identically to every arm; the origin is fixed at arming; target events are indexed $j = 1,2,\ldots$ The application target $\{\tilde{\phi}_j^{\text{app}}\}$ is the externally specified train the mission wants - preregistered, scheduler-independent, benchmark form uniform, $\tilde{\phi}_j^{\text{app}} = \tilde{\phi}_{\text{arm}} + j\,q/p$. The map-fidelity target is the ideal STAR lock itself - the locked orbit's crossing phases repeated per superperiod - and asks whether the implementation reproduces the mathematics, not whether the mathematics is what the mission wanted. Against the uniform benchmark the two coincide only at the exactly uniform operating points of §2.5; elsewhere a perfectly implemented gate carries a deterministic floor by construction (roughly 3–5 % of the nominal gap at the measured points, larger at the operational guard), stated here before Stage 1 because it is a structural handicap on the architecture's own primary metric. Shape error and origin error are related exactly by $\text{RMSE}_{\text{app}}^2 = \text{RMSE}_{\text{shape}}^2 + \bar{e}^2$; A1 values are shape errors and hence lower bounds on the preregistered score. A static pre-run alignment calibration is permitted only if declared in advance and applied identically to every baseline. [D][S:A1][S:A2]

Metric hierarchy. Events pair to targets in order - the j-th delivered event with the j-th target - because nearest-in-time pairing is precisely how a whole-cycle slip is concealed. With $e_j = \tilde{\phi}(t_j^{\text{del}}) - \tilde{\phi}_j^{\text{target}}$ unwrapped:

Primary (preregistered): Long-horizon

$$\text{RMSE}_{\text{app}} = \sqrt{\frac{1}{J}\sum_j e_j^2}$$

against the application target at the fixed arming origin. Unwrapped by construction - an event one full cycle late must not score zero. The wrapped residual is a secondary within-cycle jitter measure, never a substitute.

Hard constraint: No missing and no additional delivered events over the horizon, at a preregistered tolerance (exactly zero; a rate per 10^n events; or a confidence bound on per-event error probability). A violating scheduler is not scored on the primary metric.

Secondary: $\mathrm{RMSE}_{\mathrm{shape}}$; delivered-count RMSE against target; map-fidelity RMSE (STAR-only implementation check).

Propulsion-level: Integrated impulse error at matched mean rate and actuator model.

The ordering is forced: count error alone is passed by a scheduler firing the right number of events at entirely the wrong phases; wrapped error alone is passed by a scheduler that has slipped whole cycles. The primary metric must resist both. The Stage-1 baseline suite (§7) runs all schedulers above at matched mean rate, power, impulse distribution, disturbance, jitter, bandwidth, and computational cost. [D]

4. DESIGN REQUIREMENTS

The eleven rules, stated normatively. Each traces to a substrate measurement or declared budget; none is re-derived here (derivations: §2, §3, S§3, S§5). [D]

	Rule	Requirement	Governing expression	Source / test
1	Static hold	Operating point plus worst-case uncompensated drift stays inside the selected stationary window, on the invertible domain $0 \le \kappa < 1$	$\lvert\delta\Omega^{\mathrm{win}} p/q\rvert + \Delta\Omega\mathrm{worst} < w_{p/q}/2$; integer half-width $\kappa/2\pi$	[S:E5] [S:A1]; moving cadence: rule 10
2	Window selection	Use the simplest rational window delivering the required level; command the measured centre $\Omega^{c}_{p/q}$ (stationary), the rule-9 phase-optimal point, or the rule-11 robust point - never the rational label, which can lie outside its own window	$\Omega^{c}_{p/q} \neq p/q$ generically	[S:E5][S:A1] [S:A2]
3	Anti-reversal clearance	Every sensing channel feeding the servo holds its raw-channel sampling ratio above the statistical zone	$R \gtrsim R_{\mathrm{stat}}(\eta) \approx 2.2–2.6$ for $\eta \le 0.7$	[S:E1][S:E4]
4	Anti-freeze exclusion	No sensing channel at a freeze coincidence (integer true advance per sample), where slips are invisible	$R \not\approx 1$ and $\nu \notin \mathbb{Z}$	[S:E1]
5	Resolvability budget	Sensing resolves the finest residual structure it acts on; depth computed, not assumed	$R^{*} \sim C\mathscr{E}\eta 2/\lvert\Delta\varphi\mathrm{res}\rvert 2$	[S:E3]

6	No magic depths	No design may “tune to” a privileged sampling-depth value; resolution is smooth in R_A (rational plateaus in Ω are the exploited structure, not an exception)	ε_ϕ smooth in R_A	[S:E2][S:E5]
7	Guard budget	δ_g selected against the separate first-passage budgets of §3.2 element 3; single-step quantile screens size, never prove	$(\delta_g, \tau\,\text{early}, \alpha\,\text{extra}, \alpha\,\text{miss}, \alpha\,\text{burst}, \alpha\,\text{early}, \alpha\,\text{null})$; $0 \le \delta_g < 1$	Stage 1 **[C]**
8	Physical-units mapping	Tolerable physical drift from the normalized window at the measured centre	$\lvert f\text{gate} - \Omega^c p/q f\text{ext}\rvert < wp/q/(2T_n)$; integer case $\kappa/(2\pi T_n)$	[D]
9	Event-phase uniformity	Timing is budgeted separately from rate; sit at an exactly uniform operating point (integer window; half-integer centre, where hold- and phase-optimal coincide; $1/\tilde{q}$ at $m = 0$) or enter the measured floor at the operational guard	floors $\approx 3\,\%$ ($\tilde{q} = 3$), $5\,\%$ ($\tilde{q} = 4$) of gap, guard-dependent per the branch law of §2.5	[S:A1][S:A2]
10	Dynamic tracking	Budget trackable slew, modulation bandwidth, pull-out probability, reacquisition time; the screen $\lvert\Omega n + 1 - \Omega_n\rvert \ll wp/q$ is heuristic only. The budget on which the case is won or lost	servo authority $\le \kappa/2\pi$ per interval	Stage 1 **[C]**

11	Robust command under prediction error	Two-sided prediction bounds at stated coverage; nonempty robust intersection; gain cap; margin-constrained cost-minimizing setpoint with fallback ladder (§3.2 element 4, S§3)	$(\Omega_n, \kappa \text{eff}, n) = \rho_n(\Omega^{\text{cmd}} n, \kappa_n)$; $\kappa_n \leq \kappa \,\text{safe} \widehat{T} n \,/ T^+_{n\|n}$	$\varepsilon \lesssim (w/2)/\Omega^{\text{c}}$ | **[S:A2]**; cov

Two notes complete the rules. Sensing channel versus map update: the ratio in rules 3–5 is the raw phase-estimation channel, not the once-per-cycle Poincaré update, whose branch continuity is supplied by the internal lift (Method A by construction); what is forbidden is inferring a branch from one wrapped observation per cycle without lift, prior, or fiducial. The quoted E1–E4 thresholds transfer directly only where the detector matches the substrate's phasor-plus-Gaussian model; otherwise they are benchmarks requiring recalibration. Construction versus inference: by Proposition 1 the controller knows its cadence but not the external frequency; regime claims must cite the constructed cadence and the declared external-frequency assumption (EN-001's classes, operating inside the drive), never the data alone. Measured robustness tolerances for rule 11 are collected in Appendix A.2, running from the guarded null - whose window containment is invariant, leaving invertibility as its only cap - down to under one percent at 3/4, which for most moving cadences rules that window out. Lowering the gain does not automatically buy robustness (it narrows windows and shifts centres), and $\kappa = 0.95$ is an evidence anchor, not an operating recommendation. [D][S:A2]

5. EXPERIMENTAL SIGNATURE, CONTROLS, & METROLOGY

5.1 The staircase and its interpretation

Because $W_0(\Omega, \kappa)$ is a mode-locking staircase - the incomplete devil's staircase of subcritical coupling - at fixed $\kappa \in (0,1)$ [9], the commanded event rate of a swept gate is a staircase by construction: plateaus at rational values, on parameter intervals not generally centred on those rationals, with the measured width hierarchy and integer half-width tracking $\kappa/2\pi$. This is a statement about the stationary fixed-parameter map; under the causal realization, prediction uncertainty shifts and blurs plateaus, and the sharp staircase is the reference against which that blurring is measured. In the gate itself the staircase is a model check, not a discovery - plateau positions included: a gate whose plateaus sit at $\Omega = p/q$ at $\kappa = 0.95$ would not be reproducing the substrate map. [D][S:E5][S:A1] What the architecture predicts directly is the commanded-rate staircase in W_0, C_N, and λ_C; a thrust staircase requires in addition that the impulse bit not vary with timing, pressure, temperature, or operating point, and a deviation there indicts the actuator model, not the gate. [C]

The caution governing interpretation: the staircase is imposed on the command stream, and everything driven by that stream inherits it - power draw, valve actuation, heating, fields, vibration all step when the rate steps, so a command-correlated systematic can reproduce plateau positions, hierarchy, and even κ-scaling. A staircase in a force channel demonstrates that the gate is operating as designed; by itself it does not demonstrate momentum transfer. [D]

5.2 Momentum-discriminating controls

Run as part of the sweep, not after it [11]: matched dummy load (electrically and thermally matched, no momentum exchange - the primary null); thrust-axis reversal (genuine momentum transfer reverses sign; most artifacts do not); instrument orientation reversal (device-fixed vs lab-fixed); direct momentum balance (reservoir side must close within budget); opposed symmetric actuators (command signature present, net momentum absent by construction); scrambled-winding control (identical power spectra and mean rates, plateau structure broken); blinded sweeps and environmental logging. The gate-side null ($\kappa \to 0$ at matched mean rate) separately checks plateau provenance: with coupling removed, plateaus must vanish from the schedule - a surviving schedule staircase would falsify the gate model itself. A bare disconnect is not a dummy load: it changes exactly the systematic environment that must be matched. [D][C] (Stage 2.)

5.3 Phase-readout validity

This standard applies to stands whose observable is periodic, wrapped, interferometric modulo a fringe, or reconstructed from sampled oscillation without verified band-limiting and branch tracking; stands with calibrated linear displacement sensors and verified band-limited acquisition are governed by conventional practice [11]. Within scope, propulsion testing is the class wrapped-phase non-injectivity fools most effectively: small, slow, quasi-periodic displacement, read through a sampled phase-like observable, near the noise floor, by an experimenter with a preferred answer. The two silent failure modes are those of §2.2 - confidently wrong (alias-reversal band: inverted sign at full confidence) and invisible (freeze: impulse rendered as stillness, and a null manufactured as easily as a signal). The instrument sets its cadence but never thereby knows the unaliased frequency of the motion it reads - on a thrust stand, the very quantity under measurement. EN-001's four assumption classes translate into required instrument features, each present or its absence declared: [D][S:E1][S:E4]

EN-001 method	Assumption supplied	Thrust-stand requirement
A - known-start continuity	in-band start, smooth evolution	Measured baseline interval bounding drift and slew; continuous trace with inter-sample change $< \frac{1}{2}$ cycle
B - two-rate congruence	two rates, bounded physical range	Two independently clocked acquisition paths, documented transfer functions, common time base
C - fiducial timing	detectable marker	Independent index channel bypassing the phase stream
D - external coarse prior	second modality within half a wrap	A sensor of different physical type

Method B is the most decisive: a physical displacement agrees across rates; an unresolved alias branch generally does not. The combined workflow D + A (prior seeding plus continuity tracking) is the recommended baseline configuration - recovering true motion at RMS ≈ 0.008 cycles/sample in EN-001's demonstrated time-varying case, a consumed EN-

001 numerical result [3] carrying that note's evidence status rather than this note's [D] - with the run's full assumption ledger published beside the data. [D] (recommendation only)

5.4 The null-first decision rule

Controls run before the claim, and outcomes read as a matrix:

Observation	Interpretation
Staircase in command stream and schedule telemetry only	Gate verified; no force claim
Staircase survives matched dummy load	Command-correlated systematic; force reading invalid until separated
Signal reverses with thrust axis, momentum books close, dummy load null	Consistent with intended momentum transfer within the stated uncertainty budget
Any branch-identification control absent or failed (§5.3)	Measurement invalid; resolve before any claim

Publishing the controls beside the sweep is the recommended standard. [D][C]

6. SCOPE AND LIMITS

What the drive is. A phase-locked pulse-scheduling layer for a conventional actuator against a conventional reservoir: timing structure, a static hold interval, a command-side ledger, an audit discipline. Its first testable domain is reproducibility and countability of gated impulse; its first research question is the baseline comparison of §3.5. [D]

What the drive is not. Not a momentum source: within the ideal model STAR cannot change mean thrust at fixed mean event rate - the scheduler redistributes timing and adds a basin, a ledger, and a control structure; it amplifies nothing. Timing-dependent per-event impulse changes are ordinary actuator dynamics, measured separately. Not a specific-impulse or thrust-to-power improvement; comparisons are at matched mean power. Not yet demonstrated to outperform a digital scheduler or PLL anywhere. And it contains no reactionless mode: sustained thrust surviving the dummy-load null while failing axis reversal is a systematic to be found, and genuine reservoir-free thrust would contradict the accounting identity the architecture is built on. [D]

Where it would pay, if the baselines are beaten. Physical cadence acquisition and hold combined with missions dominated by impulse predictability: station-keeping under tight deadbands, precision formation flight, long low-thrust transfers where schedule noise compounds into navigation error, and applications where a trustworthy onboard ledger reduces reliance on external tracking. [C]

Inherited limits. The evidence pack is one-mode, one-observer, Gaussian readout noise, finite-run ensembles, with E5's circle-map coupling on the invertible domain; it supports exactly the [S:En]-tagged claims, extended by the deterministic analyses A1 and A2 released in the Stage 0b artifact package. Multi-mode structure, distributed actuation, correlated noise, dynamic acquisition and lock retention, long-horizon slips, the (Ω, κ) plane beyond the measured family, super-critical coupling, and actuator dynamics all lie beyond the demonstrated regime and belong to the validation ladder. [D]

7. VALIDATION LADDER, FALSIFIERS, AND INVALIDATORS

Validation proceeds in stages, each gated by the previous: [D]

Stage	Purpose	Primary outputs	Failure condition
0a (complete)	Released substrate	E1–E5 anchors of §2 and Appendix A	Mismatch with released evidence
0b (complete - this release)	Deterministic completion	A1/A2 routines, tables, 36 verification tests, regenerated results.json, 286-entry manuscript manifest, pinned environment	Closed: all A1/A2-based values, including rules 1, 2, 8, 9, 11, regenerate from the artifact package in a clean room
1	Closed-loop simulation with baselines	Preregistered primary metric and count constraint; acquisition, tracking, slip, guard-budget, estimator, and convergence measurements	Baseline superiority absent (→ §7.3); model falsifier fires (→ §7.1)
2	Bench with momentum controls	Full §5 control set on hardware; Stage-1 comparison repeated	Dummy-load or reversal failure (→ §7.2)
3	Applied gating	Impulse-estimate fidelity and ledger verification on a conventional actuator	No practical advantage

Stage 1 simulates the full loop of §3.1 with the declared realization of §3.2 (both realizations run; one declared, one sensitivity case), explicit reservoir, four-ledger bookkeeping, acquire-before-arm, and exact total-momentum conservation as a mandatory internal check. Its measurement list instantiates what the architecture has already defined, and is specified against S§1–S§5: the schedule staircase including plateau positions against symmetry-fixed values; dynamic acquisition (capture probability and time); long-horizon first-exit and slip statistics; the guard budgets as first-passage quantities by threshold identity; the period estimator, its induced effective-gain variation, and all three mission containment budgets, including the measured $P(\tau_{\text{loss}} < N)$; the monotone-capture equivalence in

finite-horizon form; dynamic-tracking quantities during motion, not only in steady lock; locked event phases against A1 with the attribution of §3.2 element 3; a controller-tick convergence test; and the command/delivery split so a dropped pulse is never charged to the gate. Centrally: the baseline suite of §3.5 at matched conditions, judged on the preregistered primary metric under the hard count constraint, against a superiority margin and confidence level fixed in the preregistration document before any run. [C]

7.1 Model falsifiers

Any one of these, reproduced, sends the architecture back to Stage 0. Under stationary known-period, fixed-parameter, low-noise conditions with windows exceeding scan resolution and transients discarded: (1) no plateaus in the schedule at $\kappa \in (0,1)$; (2) wrong widths, positions, or scaling at the preregistered windows - the measured ordering $w_{1/2} > w_{1/3} = w_{2/3} > w_{1/4} = w_{3/4}$, the symmetry-fixed centres $\Omega^{c}_{m/1} = m$, $\Omega^{c}_{1/2} = 1/2$, or the $\kappa/2\pi$ integer half-width scaling fails to reproduce (a blurred or displaced staircase in the causal nonautonomous implementation is a predicted consequence of period uncertainty, not a substrate failure); (3) no locked/unlocked contrast under E5-matched conditions; (4) wrong locked event phases against the full preregistered A1/A2 tables at the declared operating points and operational guard, beyond tick-convergence tolerance - the falsifier most likely to fire if an implementation has quietly reverted to two command streams; (5) privileged depth values - any reproducible non-smooth feature in estimation error as a function of R_A beyond the demonstrated structure (rational plateaus in Ω are expected and exempt). [D][S:E2]

7.2 Experimental invalidators

These invalidate a measurement or force interpretation until resolved: (1) signal changes under Method B's two-rate test with transfer functions accounted for - alias contamination; (2) staircase survives the matched dummy load - command-correlated systematic; (3) no reversal with the physical thrust axis - whatever is measured, it is not momentum transfer; (4) reservoir-side momentum books do not close within budget; (5) command-correlated contributions inseparable by the §5.2 control set - inconclusive; redesign controls before re-running. [D]

7.3 Redundancy retirement

One outcome ends the project without any falsification: if at Stages 1–2 the coupled gate fails the preregistered superiority test on the primary metric under the count constraint, the scheduler is redundant engineering and is retired in favour of the baseline; secondary metrics inform the post-mortem but do not rescue the architecture. Three quantities must stop being placeholders before Stage 1, in a separate preregistration document: the superiority margin, the confidence level, and the count-constraint tolerance (in one of the three admissible forms of §3.5) - fixed before any run, with the architecture's own deterministic floor (§3.5) declared there as expected performance rather than discovered in the comparison. The note's surviving contribution then reduces to its accounting discipline and the metrology standard of §5, which stand independently. [D]

8. CONCLUSION

STAR is not a propulsion source. It is a phase-locked pulse-scheduling architecture whose quantitative content is one auditable chain - from delivered thrust through pathwise event rates to the deterministic winding number - with each link either measured, defined, or explicitly deferred to a named validation stage. Its substrate results are deliberately unglamorous: lock windows whose locations are measured rather than read off rational labels; event trains uniform only at named operating points and otherwise carrying a computable, guard-dependent floor; a causal implementation whose period-prediction error dilates the operating point and coupling jointly and is governed by a margin-constrained robust rule with mission-level containment budgets; and an irreversible command ledger generated by a single guarded surface on the full lift.

Two boundaries bear restating because they are where the architecture is easiest to oversell. A rational winding fixes the event count per superperiod, not the event phases inside it - the timing performance that motivates the design is a separate, measured quantity. And the gate directly predicts an event-rate staircase; a thrust staircase is an additional actuator-dependent claim that stands or falls with the constant-impulse-bit model.

What is measured is the gate's statics and finite-horizon noise behaviour; what is proposed is the assembly, staked on two questions the note does not prejudge: whether measured thrust tracks the commanded staircase and survives the momentum-discriminating controls, and whether the coupled gate earns its place against digital schedulers and PLLs on the preregistered event-phase metric in the one regime where it plausibly could - acquisition, hold, and tracking of a moving physical cadence. A clean loss retires the scheduler with no harm done. Two deliverables stand regardless: the accounting discipline (momentum per event times event rate, every counter distinct and auditable), and the metrology standard for phase-based thrust stands, which is overdue independently of this architecture.

Star Drive is offered as the parent framework requires: precise enough to fail on a bench, in a named place, at a stated number.

REFERENCES

[1] Dupke, A. (2026). Scale-Time Theory 11.0 - The Pre-Geometry of Scale Space (Version 11.0). Scale-Time Dynamics Research. Zenodo. https://doi.org/10.5281/zenodo.21221364

[2] Dupke, A. (2026). Scale-Time Theory 11.0 - Simulation Evidence Pack, release 1.0. `stt11_simulation.py`, fixed seed 20599373, figures `figE1`–`figE5` and `figM3`–`figM5`, `results.json`. Scale-Time Dynamics Research.

[3] Dupke, A. (2026). Relative-Ratio Bookkeeping for Kinematic Aliasing: An Assumption-Driven Framework for Wrapped-Phase Recovery (Engineering Note EN-001). Zenodo. https://doi.org/10.5281/zenodo.18904052

[4] Dupke, A. (2026). Wrapped Phase, Protected Qubits: Minimal Phase-Channel Model for Syndrome-Guided Branch Tracking in Quantum Error Correction (Version 2.0.1). Scale-Time Dynamics Research. Zenodo. https://doi.org/10.5281/zenodo.21760092

[5] Adler, R. (1946). A study of locking phenomena in oscillators. Proceedings of the IRE, 34(6), 351–357. https://doi.org/10.1109/JRPROC.1946.229930

[6] Gardner, F. M. (2005). Phaselock Techniques (3rd ed.). Wiley.

[7] Egan, W. F. (2000). Frequency Synthesis by Phase Lock (2nd ed.). Wiley.

[8] Pikovsky, A., Rosenblum, M., & Kurths, J. (2001). Synchronization: A Universal Concept in Nonlinear Sciences. Cambridge University Press.

[9] Jensen, M. H., Bak, P., & Bohr, T. (1983). Complete devil's staircase, fractal dimension, and universality of mode-locking structure in the circle map. Physical Review Letters, 50(21), 1637–1639. https://doi.org/10.1103/PhysRevLett.50.1637

[10] Hänggi, P., Talkner, P., & Borkovec, M. (1990). Reaction-rate theory: fifty years after Kramers. Reviews of Modern Physics, 62(2), 251–341. https://doi.org/10.1103/RevModPhys.62.251

[11] Polk, J. E., Pancotti, A., Haag, T., King, S., Walker, M., Blakely, J., & Ziemer, J. (2017). Recommended practice for thrust measurement in electric propulsion testing. Journal of Propulsion and Power, 33(3), 539–555. https://doi.org/10.2514/1.B35564

APPENDIX A - SUBSTRATE EVIDENCE SUMMARY

All stochastic results regenerate with fixed seed 20599373; A1 and A2 are deterministic and regenerate from their fixed solver specifications (`a1_routine.py`, `a2_routine.py`). All are checked by automated verification. E1–E4 use the minimal phasor-estimation model of §2.2; E5 adds the sine circle map on the invertible domain, with widths measured on the noiseless map and ensemble statistics finite-run under drive noise. [D]

The provenance map below identifies what each experiment supplies; the consumed numbers have a single home in the body and are not restated here.

Experiment	Supplies	Canonical home
E1 - spin-direction statistics	Structural-identity coin flip, alias reversal, stroboscopic freeze, statistical-zone exit depths	§2.2 table; rules 3–4; §5.3
E2 - resolution and passive recurrence	Cramér–Rao-limited smoothness in R_A; measure-zero passive recurrence	§2.3; §2.4; rule 6; §7.1 falsifier 5
E3 - resolvability	Structure-relative depth thresholds with $\eta 2$ scaling	§2.3; rule 5
E4 - structural identity and rate uncertainty	Exact ± identity at $R_A = 2$; edge-zone discrimination error	§2.2 table; §5.3
E5 - coupling and lock windows	Released window widths, locked fraction, finite-run confinement-vs-diffusion ensemble	§2.4 table; §2.6; §3.4; rules 1–2; §5.1

A.1 Analysis A1 - window boundaries, centres, & unguarded event phases

Method. The released E5 scan thresholds a finite-iteration winding estimate on a uniform grid ($\Delta\Omega = 1/2400$) and records widths only. A1 recomputes window boundaries from the same map at the same coupling by locating the existence of a period-q orbit: root candidates of $g = f^q(x) - x - p$ are detected either by a sign change between adjacent grid nodes or by a node with $|g| \leq$ tol, then refined, and the boundaries are solved from the simultaneous saddle-node conditions $g = 0$ and $\partial_x f^q(x) - 1 = 0$ to $\sim 10^{-10}$, with residuals recorded for both equations ($\leq 2.4 \times 10^{-14}$ achieved across all fourteen edges) and the dense-grid multiplier check retained as independent bracketing (EN-002-EP, S§7). It computes locked event phases in closed form under the declared realization (linear intra-interval lift, $\delta_g = 0$, transient $\geq 2 \times 10^4$ intervals), and an orbit that fails to meet its convergence tolerance within its iteration budget is raised as a recorded failure rather than returned. No stochastic element, no new parameter. A1 characterizes the stationary, known-period, zero-guard reference map at $\kappa = 0.95$: it supplies nominal targets, and realized windows and phases under a moving cadence are Stage-1 quantities. [S:A1][D]

Orbit-pair verification. Inside a tongue the period-q orbit comes in a stable/unstable pair that merges at the saddle-node edges, and the scheduler follows the stable one. A1 verifies this in every window at three parameters - the centre and ± half the half-width - refining all roots of g, grouping them into cycles, classifying each by its multiplier, and confirming by direct iteration which cycle the attractor occupies. All seven windows carry a genuine pair with a stable attractor at all three probes; the attractor multipliers at the centres are 0.05 ($m/1$), 0.0975 (1/2), 0.1434 (1/3, 2/3) and 0.1834 (1/4, 3/4), against unstable partners above unity. [S:A1]

The window table is given in §2.4 and is not repeated. Locked event phases at the $m = 0$ window centres:

p/q	Per-interval advances Δ_n	Event gaps (external cycles)	Uniform q/p	RMS shape deviation
1/2	0.5, 0.5	2 (exact)	2	0
1/3	0.42247, 0.34908, 0.22845	3 (exact)	3	0
1/4	0.27920, 0.13307, 0.20396, 0.38377	4 (exact)	4	0
2/3	0.57753, 0.65092, 0.77155	1.4521, 1.5479	1.5	0.023952
3/4	0.72080, 0.86693, 0.79604, 0.61623	1.3215, 1.1961, 1.4824	1.3333	0.067657

The advances are strongly nonuniform even where the $m = 0$ train is not: one event per superperiod cannot be unevenly spaced, which is why that uniformity does not survive integer translation ($\tilde{q} = 3$ translates measure 3.0–3.2 % of the gap, $\tilde{q} = 4$ translates 5.2–5.3 %; full translated tables, half-integer detuning sensitivity, and interior uniform-point structure: S§7). RMS shape deviation is

$$\min_b \sqrt{\frac{1}{J}\sum_j (e_j - b)^2}$$

- a lower bound on the preregistered $\mathrm{RMSE}_{\mathrm{app}}$, which is taken at the fixed arming origin. [S:A1]

A.2 Analysis A2 - guard branches, ϰ-continuation, and robust tolerances

A2 extends the same deterministic map along the two axes A1 does not touch - the guard δ_g and the coupling κ - and computes the robust-controller quantities of rule 11. Every A2 number is conditional on the anchor-map identity A1 establishes. [D]

Guard-branch coefficients (window centres, $\kappa = 0.95$; quadratic law of §2.5, coefficients valid on the branch containing $\delta_g = 0$ up to the listed breakpoint):

Case	$\mathrm{RMSE}_{\mathrm{shape}}(0)$	C	V	**Branch valid to**	**At** $\delta_g = 0.1$
$m/1$; $m + 1/2$ centre; $1/\tilde{q}$ at $m = 0$	0	0	0	zero-error law valid at every admissible guard (branch crossings occur without changing the value)	0
2/3	0.023952	+0.002877	0.014425	0.4210	0.035962
3/4	0.067657	+0.000912	0.009159	0.3815	0.069652
1 + 2/3	0.018450	+0.000001	0.000721	0.4210	0.018647
1 + 3/4	0.030227	−0.000055	0.000750	0.3815	0.030170
1 + 1/3	0.022452	+0.000107	0.002055	0.0015	0.022577 (measured beyond branch)

$1 + 1/4$	0.041612	+0.000437	0.003866	0.0022	0.041993 (measured bey

The closed form reproduces direct computation to six decimals within each branch (verified by the A2 regression tests); a quadratic quoted outside its branch is simply the wrong curve (the $1 + 1/3$ and $1 + 1/4$ crossings sit within $\sim 2 \times 10^{-3}$ of a fiducial, so any practical guard changes their branch - those columns are directly measured). The shape-optimal guard $-C/V$ is admissible only where $C < 0$ within the declared guard domain $0 \leq \delta_g < 1$ (only $1 + 3/4$, for a $0.2\,\%$ improvement): the guard is a real but weak phase-shaping parameter, fixed by the budgets of rule 7. [S:A2][D]

Full branch coverage over the guard domain. The table above reports the branch containing $\delta_g = 0$. A2 additionally enumerates every branch across $0 \leq \delta_g < 1$ for each case - stepping from breakpoint to breakpoint, recording each branch's validity interval and its own coefficients relative to its own reference guard, and verifying the closed form inside each - yielding four branches at $\tilde{q} = 3$, five at $\tilde{q} = 4$, two at $m + 1/2$ and one at $m/1$, with the law verified on all of them (worst in-branch error $< 5 \times 10^{-7}$). The operational guard is not fixed in this release: rule 7 makes δ_g a first-passage budget quantity, preregistered before Stage 1. What Stage 0b supplies instead is the complete branch map plus a verifier that identifies the branch containing any candidate guard and confirms the coefficients quoted for it; running that verifier on the preregistered value is a Stage-1 entry condition (EN-002-EP). [S:A2][D]

Coupling-axis trade-off (Newton edge continuation with the same boundary/orbit solver as A1; anchor consistency verified by regression test; all four columns increase strictly across the sampled set - width is bought at the price of nonuniformity; continuous monotonicity not claimed). One canonicalization finding from the artifact: the shape deviation at the exact $2/3$ solver centre is 0.023953, while the 0.023952 of the guard table and §2.5 is the evaluation at the 6-dp rounded centre 0.649532 - both are recorded in `results.json`, the manifest labels each occurrence, and the guard-table values remain quoted at the rounded centre, the operational convention for a commanded setpoint:

κ	$w_{2/3}$	dev 2/3	$w_{3/4}$	dev 3/4
0.40	0.002196	0.013313	0.000515	0.032725
0.60	0.007209	0.018480	0.002448	0.047344
0.80	0.016437	0.022185	0.007137	0.059911
0.90	0.022862	0.023446	0.010921	0.065218
0.95	0.026550	0.023953	0.013236	0.067657

(Full nine-row table: \`results.json\` / S§7.) Buying prediction headroom by lowering the gain costs hold width - $w_{2/3} = 0.01644$ at $\kappa = 0.80$ against 0.02655 at 0.95, about $38\,\%$ - and the low-κ end is a hold-budget statement against mission drift, not a general verdict. [S:A2]

Robust operating-point tolerances at $\kappa_n = 0.95$ (bisection on ε until the robust intersection empties; integer rows analytic, $\varepsilon_{\max} = 2c/(\sqrt{1+4c^2}+1)$, $c = \kappa_n/2\pi m$; derivation S§5):

Operating point	**Screen** $(w/2)/\Omega^c)$	$\varepsilon_{\max}$, **rate-blended**	**Binding constraint**	$\varepsilon_{\max}\$, phase-step$
0/1 guarded null	-	$5.263\,\%$ - window containment invariant for all $\rho > 0$; the cap is effective-gain/invertibility	invertibility	not prediction-limited (window and gain both invariant)
1/1	$15.12\,\%$	$5.263\,\%$	invertibility	$15.120\,\%$
2/1	$7.56\,\%$	$5.263\,\%$	invertibility	$7.560\,\%$
3/1	$5.04\,\%$	$5.027\,\%$	window drift	$5.040\,\%$
5/1	$3.02\,\%$	$3.021\,\%$	window drift	$3.024\,\%$
1/2	$6.72\,\%$	$5.263\,\%$	invertibility	$6.718\,\%$
1/3	$3.79\,\%$	$4.112\,\%$	window drift	$3.788\,\%$

2/3	2.04 %	1.959 %	window drift	2.044 %
1/4	2.35 %	2.900 %	window drift	2.352 %
3/4	0.92 %	0.858 %	window drift	0.921 %

Tolerance falls roughly as $(w/2)/\Omega^{c}$ - equally wide windows at larger Ω are markedly less robust; the binding constraint crosses from invertibility to window drift near $m = 3$; the guarded null's window containment is exactly invariant because $\Omega^{\mathrm{cmd}} = 0$ scales to 0 - its operational $\varepsilon_{\max}$ is therefore set entirely by the effective-gain cap in the rate-blended realization, and no prediction limit remains under phase-step. The phase-step column helps strictly where invertibility binds and changes little, not always favourably, where drift binds. [S:A2][D]

Optimizer anchor (one example). For fixed gain, symmetric band $\rho \in [1-\varepsilon, 1+\varepsilon]$, non-binding slew, and an integer tongue, the margin-maximizing command is $\Omega^{\mathrm{cmd}} = m + \varepsilon\kappa_n/2\pi$ with margin $\kappa_n(1-\varepsilon^2)/2\pi - \varepsilon m$ - an analytic regression anchor for the A2 optimizer, reproduced numerically by the test suite. At the $2/3$ window with $\varepsilon = 0.015$: admissible interval [0.646920, 0.653145]; ceiling at $\Omega = 0.650079$ with worst-case margin 0.003112; interval midpoint 0.003066; nominal centre 0.649532 only 0.002573 - a sixth of the available margin discarded by the naive choice, which is why the margin floor is calibrated against a computed ceiling rather than guessed. [S:A2]

A.3 Reproducibility manifest

The Stage 0b artifact package ships with this note: separate A1 and A2 routines (`a1_routine.py`, `a2_routine.py`) with fixed solver tolerances, independent regression tests (`test_stage0b.py`, 36 tests, including a class that recomputes A1/A2 quantities in-process so a source regression cannot pass by validating stored output), a pinned environment (`pyproject.toml`, `requirements.lock`; CPython 3.14.7, NumPy 2.5.1, SciPy 1.18.0, pytest 9.1.1), a machine-readable output schema (`stage0b-1.1`), regenerated `results.json` entries alongside E1–E5, and the repository manifest (`manifest.json`) identifying every manuscript value and its generating artifact - 286 entries, one canonical value per manuscript number, scan-accuracy duplicates labelled. The regeneration workflow refuses phase caches written under a different source, runtime or platform specification, writes every artifact through a temporary file and atomic rename, and exits nonzero on any manifest mismatch or artifact-level failure; a clean-room run on the same recorded runtime and platform reproduces `results.json` bit-identically apart from its timestamp. The release condition of EN-002-EP is met: both routines, their verification tests, and the regenerated `results.json` ship together, and Stage 0b is closed (§7). The pack tests nothing beyond the tagged claims: no momentum exchange, no actuator response, no dynamic acquisition or lock retention, no assembled thrust loop - every claim of that kind is [C] with a stage in §7. [D]

APPENDIX B - NOTATION

Symbols appearing in this note; supplement-only symbols are defined in EN-002-S. Units: phases in cycles, times in seconds.

Symbol	Meaning
ν, ν_{alias}, k, $\mathscr{A}$	True and wrapped normalized advance; hidden alias branch; admissible branch set
R_A, Rstat(η), R*A	Samples per mode cycle (raw estimation channel in rules 3–5); statistical clearance depth; resolvability threshold
η, $\varepsilon\phi$, $\Delta\varphi$res	Readout-noise scale; per-cycle estimation error; residual structure to resolve
$f$ext, $f$gate, $t^{\text{fid}}n$, T_n, Tmin, T_N^{Σ}	Reference and gate frequencies; fiducial instants; measured interval; shortest interval; run duration
ϕ_n, ψ_n, y_n, $u(\,\cdot\,)$, ϵ_n, ξ_n	Reference and gate phases; detector output; servo action; readout noise; process noise
x_n, $z_n = x_n - mn$, θ_n	Full lifted gate phase; reduced lift (bookkeeping only); wrapped phase
Ω; $W^\star = m + \tilde{p}/\tilde{q}$; $m = \lfloor W^\star \rfloor$; $\bar{\Omega}_n = \Omega_n - m$	Frequency ratio; selected winding target; fixed integer translation, held per operating mode (tongue-family changes re-arm); translated parameter, unconstrained
p/q, $\Omega^{\pm}p/q$, wp/q, $\Omega^{\text{c}}_{p/q}$	Rotation number (plateau value); window boundaries, width, measured centre
$\delta\Omega^{\text{rate}}$, $\delta\Omega^{\text{win}}$	Rate mismatch $\Omega - p/q$; window-relative detuning $\Omega - \Omega^{\text{c}}$ - the operative hold variable

κ, κ_n, κsafe, κeff, n	Coupling gain (domain [0,1]); integer half-width $\kappa/2\pi$); commanded gain; robust cap; effective gain $\rho_n\kappa_n < 1$
δ_g, $\ell + \delta_g$, L_r, M_n, ΔC_n	Guard band, domain $0 \leq \delta_g < 1$; guarded event surface; tick-level threshold; highest committed threshold; per-interval commands
C_N, K_N^{full}, K_N^{frac}, K_N^{del}, K_N^{target}	Irreversible command ledger; winding counters; verified delivered count; target count
N^{cmd}extra, N^{cmd}miss, Ndrop, Ndup, E_r^{del}	Unmatched-threshold counts vs the causal reference; actuator drop/duplicate counts; net execution error (cross-check only)
αextra, αmiss, αburst, αearly, α_{null}	Finite-horizon first-passage budgets of §3.2 element 3
$\alpha\kappa$, mission, αparam, mission, $\alpha_{\mathrm{lock,mission}}$	Mission budgets: invertible domain, stationary tongue, dynamic lock (τ_{loss}: S§4)
$t^{\mathrm{cmd}}\ell$, $t^{\mathrm{del}}\ell$, $t^{(0)}\ell$, τlat, τearly, τdeadline	Command, delivered, and causal-reference crossing times; latency; allowed lead; application deadline (does not define identity)
$\widehat{T}_n n$	$T^{\pm}n \mid n$, ρ_n, $\rho^{\pm}n$ | Causal period prediction (default persistence) and bounds; prediction ratio $T_n/\widehat{T}_{n\mid n}$ and bounds
$\Omega^{\mathrm{cmd}}n$, $\mathscr{M}^{\mathrm{design}}$min, $s\Omega$, $s\kappa$	Commanded ratio; preregistered margin floor; setpoint slew limits (full controller: S§3)
Δt, Δ_n	Controller tick; realized per-interval full-lift advance (event spacing $1/\Delta_n$; guard delay δ_g/Δ_n)
W_0, $\widehat{W}_N$, $\overline{W}$	Deterministic winding number; finite-run estimate; stochastic mean winding

$\tilde{\phi}$, $\tilde{\phi}$arm, $\tilde{\phi}^{\mathrm{app}} j$, $\tilde{\phi}_j^{\mathrm{map}}$, e_j	Unwrapped external phase (fiducial-linear default); arming origin; application and map-fidelity targets; unwrapped order-paired error
RMSE app, RMSE shape, RMSE_K	Primary metric; shape (offset-removed) error, $\mathrm{RMSE\,app}^2 = \mathrm{RMSE\,shape}^2 + \bar{e}2$; count RMSE
$\vec{\chi}, \vec{\chi}_j, \bar{\vec{\chi}}N$, $\overrightarrow{\mu}^{\mathrm{event}}\chi$, $\overrightarrow{\mu}_\chi$	Nominal impulse bit; realized per-event impulse; event-weighted realized bit; event-conditioned mean; ordinary mark mean
$\vec{J}_N$, $\overline{\vec{F}}N$, $\widehat{\lambda}C, N$, $\widehat{\lambda}\mathrm{del}, N$, λ_C, $\lambda\,\mathrm{del}$	Delivered impulse; finite-horizon mean thrust; pathwise command and delivered rates; asymptotic rates (where limits exist)
$C\varepsilon$, $C\mathscr{E}$	Estimator and resolvability coefficients