

FUSION DENSITY UNDER AMBIGUITY

 Scale-Time Dynamics

PRESERVING USEFUL DENSITY UNDER
UNRESOLVED FRINGE COUNTS

André Dupke

Andalucía, October 2, 2026

scaletimedynamics.com

info@scaletimedynamics.com

Orcid | Academia | YouTube | Facebook

Version: 1.0

©2026 - Scale-Time Dynamics LLC - Licensed under CC BY 4.0

DOI: doi.org/10.5281/zenodo.23105860

Fusion Density Under Ambiguity

Preserving Useful Density Under Unresolved Fringe Counts

A Density-Risk Decision for Two-Color Interferometer–Polarimeter Systems
in Magnetic-Confinement Fusion

ABSTRACT

Interferometers that measure the line-integrated electron density of magnetically confined fusion plasmas observe a wrapped phase. The whole number of fringes accumulated since the start of a discharge is absent from the data and has to be restored by assumption; when refraction, pellets or disruptions interrupt the signal, the restored count can jump by whole fringes, and the density supplied to plasma control inherits the error. Current practice treats such an event as integer recovery: select a branch, correct the count, report the density.

This note proposes a decision placed on the quantity the control loop consumes. For the two-color interferometer–polarimeter topology that ITER uses as its primary density-control diagnostic, the estimator marginalizes the posterior over integer hypotheses and continuous uncertainty, reports the density risk $q_t = P(|N - \hat{N}_t| > \epsilon \mid y_{1:t})$ through a conservative bound $\bar{q}_t \geq q_t$, and declares the density available when $\bar{q}_t \leq \alpha$, with or without a resolved fringe count. The note supplies the measurement model, the slip lattices of the DIII-D prototype and of the ITER design, a joint decision margin in the metric of density and path priors, and an observability accounting that counts nuisance drifts and priors explicitly.

In a Gaussian model of the prototype with assumed priors and phase noise, the most probable fringe count is wrong in 29 % of decisions and the count confidence stayed below 99 % in every simulated decision, so a rule that waits for a resolved count withholds every density; the marginalized estimate keeps every density available and lowers the rate of total density errors beyond a 0.05-fringe tolerance from 6.6×10^{-5} to 1.7×10^{-5} . For the ITER wavelength pair the estimate gains nothing, and withholding uncertain densities lowers the error rate about a hundredfold. With the estimate held fixed, the density-risk gate and a count-confidence gate accepted identical decisions in both configurations. The demonstrated benefit is therefore one of estimation: marginalization preserves useful density despite unresolved counts, while an advantage of gating on the density over gating on the count remains a research question. All results are model-conditional and depend on stated search bounds.

Operational reliability, real-time feasibility and any consequence for fusion performance remain open. The note therefore specifies how others can test the decision on measured data: implementation tests, raw-signal replay and bench optics with independent truth; comparison at matched availability against hard branch selection, integer-aperture acceptance, deployed correctors and machine-learning baselines, with estimate and gate compared separately; named falsifiers, invalidators and a retirement criterion. It is offered for evaluation by groups with access to interferometer–polarimeter records.

Content

CLAIM-STATUS TAGS	6
1. INTRODUCTION AND POSITION	7
2. SUBSTRATE: WHAT THE APPROACH CONSUMES	9
2.1 Wrapped observation and the absent integer	9
2.2 Residuals on wrapped and lifted phases	9
2.3 The joint decision margin	9
2.4 Parity accounting	10
3. MEASUREMENT MODEL: THE TIP TOPOLOGY	11
3.1 Instrument and phase	11
3.2 Two-color inversion and the slip lattice	11
3.3 The polarimeter as an absolute channel	12
3.4 Temperature, timing and chromatic drift	13
3.5 Detection and isolation in this register	14
4. THE DENSITY-RISK DECISION	15
4.1 Posterior over integer hypotheses	15
4.2 Density risk, its bound, and availability	16
4.3 Four comparisons	17
4.4 Computation	17
4.5 Relatives and what is proposed	18
5. A MODEL-CONDITIONAL ILLUSTRATION	19
5.1 Assumptions	19
5.2 Branch-induced displacement and total error	19
5.3 Four comparisons	20
5.4 Search bounds and the operating domain	22
5.5 Reach of the illustration	22
6. EVALUATION PROTOCOL	23
6.1 Data, replay and ground truth	23
6.2 Baselines and matched availability	23
6.3 Metrics	24
6.4 Requested from partner groups	24
7. VALIDATION LADDER, FALSIFIERS AND INVALIDATORS	25
7.1 Model falsifiers	25
7.2 Experimental invalidators	26
7.3 Retirement criterion	27

8. SCOPE AND LIMITS	28
9. CONCLUSION	29
REFERENCES	30
APPENDIX A. NOTATION AND TERMS	33
Terms	34
APPENDIX B. REPRODUCTION SCRIPT AND MANIFEST	37

CLAIM-STATUS TAGS

Substantive claims carry a bracketed status tag, following the discipline of the parent framework [1].

[D] Definition, convention, or derivation within a stated model.

[A] Explicit modeling assumption on which later statements depend.

[M] Model-conditional result: a derivation that holds under the stated assumptions.

[N] Number computed for this note; Appendix B reproduces it. Rare-event rates carry exact intervals or one-sided bounds.

[R] Result reported in the cited external literature.

[C] Conjecture, proposal or untested application, paired with a stage of the validation ladder in Section 7.

[I] Interpretive device without theoretical weight.

The abstract and the conclusion are tag-free summaries: every claim in them inherits the status of its counterpart in the body.

1. INTRODUCTION AND POSITION

Electron density is one of the quantities a tokamak control system must hold: it sets the operating point, governs fueling and approaches to the density limit, and enters disruption avoidance. ITER's five-channel Toroidal Interferometer and Polarimeter, TIP, is the primary diagnostic for real-time density control [2]. Its interferometer measures a phase proportional to the line-integrated density, wrapped into one fringe; the accumulated fringe count carries most of the measurement and is restored from the history of the signal. Fringe jumps, whole-fringe errors in that count, occur during edge-localized modes (ELMs), fast density rises, pellet injection and disruptions, when density gradients refract the beam, reduce the signal amplitude and break phase continuity [3][4]. For TIP itself, ray tracing predicts negligible refraction from ELMs and significant refraction once line-averaged densities reach the mid to upper 10^{21} m^{-3} range of disruption mitigation [2]. Correctors in use compare two wavelengths, consult a polarimeter, run phase counters and model-based observers, and, recently, machine-learning classifiers [3][5][6][7][8][9][10]. [R]

These correctors share one structure, integer recovery: decide the count, then report the density that follows from it. The control loop, however, consumes the density, and in a two-color instrument count and density can part company. Some count errors move the density by a small fraction of a fringe while moving the inferred optical path by many micrometers. An unresolved count can then coexist with a trustworthy density, and a rule that withholds the density until the count is resolved discards usable measurements. [M]

Position, stated plainly. EN-003 proposes one decision rule and one evaluation protocol for one class of diagnostic. It works within established measurement physics, leaves the burn, confinement, materials and magnets of fusion devices untouched, and its claims end at the density measurement. Its numbers come from a Gaussian model with assumed priors and noise: they demonstrate a mechanism and leave operational reliability to measured data. The closest relatives come from satellite navigation, where integer least squares, integer-equivariant estimation and integer-aperture estimation handle carrier-phase ambiguities [11][12][13][14]. The element proposed here is the placement of the decision on a tolerance for the density, with the benefit of a better estimate and the benefit of a better gate reported separately. [D][C] In the model of Section 5 the demonstrated benefit lies in the estimate: marginalization preserves useful density despite unresolved counts, while the density-risk gate withheld the same decisions as a gate on the count. [N]

Lineage. The approach comes from the sampling layer of Scale-Time Theory (STT) [1]: the non-injectivity of wrapped observation and the four assumption classes that restore a lost integer (EN-001 [15]); residuals on lifted phases and the separation of estimation from protection (HCOPL, Harmonic Consistency and Oversampled Phase-Lock [16]); ledgers, falsi-

fiers and a retirement criterion (EN-002 [17]). The note is self-contained; reading the earlier documents is optional, and the scale-space interpretation of the parent framework plays no part in it. [D]

Three principles shape the design. [D]

- I. The integer is absent from the data. Every restored count rests on an assumption, and the assumption has to be named [15].
- II. The decision belongs on the consumed quantity. A control loop uses the density, so availability is decided by a tolerance on the density, and the count is reported as a separate quantity.
- III. Reliability is a measured number. A claim of reliability names its data, its baseline, its tolerance and the number at which it fails.

2. SUBSTRATE: WHAT THE APPROACH CONSUMES

2.1 Wrapped observation and the absent integer

An instrument returns $\tilde{\phi} = \text{wrap}(\phi) = \phi - 2\pi k$, with $\text{wrap}(\phi) = [(\phi + \pi) \bmod 2\pi] - \pi$ and an unknown integer k . The likelihood of wrapped data is invariant under integer shifts, so k is absent from the data, and every recovery supplies it from an additional source [15]: continuity from a known start (A), congruence of channels with different scale factors (B), a fiducial or absolute reference (C), and an external coarse prior (D). Practical systems combine a prior with continuity tracking. For nearest-branch tracking the exact structural limit is half a fringe of phase change per tracking step; at the limit rising and falling density become indistinguishable, and beyond it the tracker reports reversed changes [15] [18]. [D][M]

2.2 Residuals on wrapped and lifted phases

For channels that observe one coordinate θ with integer harmonic indices h_m , the wrapped residual $\Delta_{mn} = \text{wrap}(h_n \tilde{\phi}_m - h_m \tilde{\phi}_n)$ cancels every branch integer and therefore stays blind to whole-cycle count errors. After branch selection, lifted phases $x_m = \tilde{\phi}_m + 2\pi \hat{k}_m$ and residuals $S_{mn} = h_n x_m - h_m x_n$ expose inconsistent counts. Coordinated errors $q_m = \ell h_m$, for an integer ℓ , shift the coordinate by whole cycles and stay invisible, so only a prior or an absolute anchor resolves them [16]. [M]

2.3 The joint decision margin

A branch decision compares candidate solutions with priors. When several integers are coupled with continuous parameters, each integer hypothesis fixes a solution; two hypotheses differ by a lattice vector δ , and with the covariance Σ of the decision residual, prior and measurement errors together, their separation is the Mahalanobis distance

$$d = \sqrt{\delta^T \Sigma^{-1} \delta}$$

For Gaussian errors and equal prior weights, two hypotheses are confused with probability

$$\frac{1}{2} \text{erfc} \left(\frac{d}{2\sqrt{2}} \right),$$

and the sum of this term over all competitors bounds the wrong-count probability [19]. The familiar half-fringe rule is the special case of one unresolved branch with all others fixed: $d = 2b/\sigma_r$ for a margin b , a residual spread σ_r , and two neighbors. A competitor that differs mainly in a nuisance direction can stay unresolved while the target quantity stays accurate, and the target's accuracy combines that displacement with the continuous measurement error. [M][R]

Two accounting rules complete the margin. For n repairs per discharge with wrong-repair probabilities p_i , the union bound

$$P_{\text{any}} \leq \sum_i p_i$$

holds under arbitrary dependence, while the independent form $1 - (1 - p)^n$ needs evidence of independence, which clustered jumps during a disruption contradict; for $p = 10^{-3}$ and $n = 50$ the two give 5.0×10^{-2} and 4.9×10^{-2} . A simulation with zero failures in n trials bounds the failure rate at 95 % confidence by $1 - 0.05^{1/n} \approx 3/n$. [M][N]

2.4 Parity accounting

For a linearized model $\mathbf{y} = H\mathbf{u} + \mathbf{f} + \mathbf{e}$ with M scalar measurements and n_u continuous unknowns $\mathbf{u}$, a matrix Π whose rows span the left null space of H forms parity residuals $\mathbf{r} = \Pi\mathbf{y} = \Pi\mathbf{f} + \Pi\mathbf{e}$, free of $\mathbf{u}$, in a parity space of dimension $M - n_u$ [20]. A fault in channel m is detectable when column m of Π is nonzero and the fault exceeds the noise; two single-channel faults are distinguishable when their columns are non-collinear. With $M - n_u = 1$ every fault signature lies on one line, so faults are detectable and never isolable, and every leave-one-out subset of n_u measurements fits exactly. [M][R]

The count of unknowns includes every continuous state free at that instant, nuisance drifts included. A prior on a state enters as a pseudo-measurement and adds one parity dimension, whose residual is only as reliable as the prior. Over a window of T_w samples with n_f states free at each sample and n_s states shared across the window, the stacked model has $(M - n_f)T_w - n_s$ parity dimensions when its matrix has full rank; this is how slow states supply temporal redundancy [20]. [M]

3. MEASUREMENT MODEL: THE TIP TOPOLOGY

3.1 Instrument and phase

The topology studied combines a 10.59 μm CO₂ interferometer, a second-color interferometer for vibration compensation and a 10.59 μm Faraday polarimeter. The DIII-D prototype of ITER's TIP uses a 5.22 μm quantum-cascade laser (QCL) as second color [6][21]; the ITER design uses 4.6 μm [2]. ITER's Density Interferometer Polarimeter, DIP, a dispersion interferometer combined with a polarimeter, complements TIP with a different phase structure and lies outside this note [22][23]. [R][D]

In the cold-plasma approximation, for probing frequencies far above the plasma frequency, an interferometer at wavelength λ measures

$$\phi = r_e \lambda N, \quad N = \int n_e dl$$

with $r_e \approx 2.818 \times 10^{-15}$ m the classical electron radius and the line integral taken along the full beam path; for a double-pass chord, as in TIP, it includes both passes. One fringe corresponds to $N_F = 2\pi/(r_e \lambda)$: $2.11 \times 10^{20} \text{ m}^{-2}$ at 10.59 μm , $4.27 \times 10^{20} \text{ m}^{-2}$ at 5.22 μm and $4.85 \times 10^{20} \text{ m}^{-2}$ at 4.6 μm . Densities below are quoted in fringes at 10.59 μm . [D][N]

3.2 Two-color inversion and the slip lattice

With an optical path change ΔL , the change in optical path length caused by vibration and other motion of the optics, the two phases are $\phi_i = r_e \lambda_i N + 2\pi \Delta L / \lambda_i$; a displacement Δx of a mirror or retroreflector along a double-pass beam gives $\Delta L = 2 \Delta x$. The vibration-free combination

$$N = \frac{\phi_1 \lambda_1 - \phi_2 \lambda_2}{r_e (\lambda_1^2 - \lambda_2^2)}$$

turns count errors (k_1, k_2) in the lifted phases into a density step and a path step

$$\delta N = \frac{2\pi (k_1 \lambda_1 - k_2 \lambda_2)}{r_e (\lambda_1^2 - \lambda_2^2)}, \quad \delta L = \frac{\lambda_1 \lambda_2 (k_2 \lambda_1 - k_1 \lambda_2)}{\lambda_1^2 - \lambda_2^2}$$

[M] The two phases determine density and path exactly for every pair of integers, so the integers follow only from priors on both continuous quantities. The candidate set is the

full integer lattice; it becomes a finite search set under bounds on the density and path changes plausible across a signal gap. [M][A] Steps in fringes at 10.59 μm and in micrometers of path:

Slip (k_1, k_2)	Prototype δN (fringe)	Prototype δL (μm)	ITER δN (fringe)	ITER δL (μm)
(1, 2)	0.0187	10.39	0.1618	8.88
(3, 7)	-0.5950	38.07	-0.0500	32.30
(10, 23)	-1.7663	124.60	0.0116	105.78
(1, 3)	-0.6324	17.29	-0.3736	14.55
(0, 1)	-0.6511	6.90	-0.5354	5.67
(1, 1)	0.6698	3.50	0.6972	3.21
(1, 0)	1.3209	-3.40	1.2326	-2.46

[N] Each step comes with its negative. The near-null direction, the lattice vector whose ratio k_2/k_1 best approximates λ_1/λ_2 within the path range of the search set, depends on that range. For the prototype it is (1,2), since $\lambda_1 \approx 2\lambda_2$: the slip moves the density by only 0.0187 fringe, $3.94 \times 10^{18} \text{ m}^{-2}$, and the path by 10.4 μm , so it barely corrupts the density and strongly corrupts the vibration estimate. For the ITER pair the same slip moves the density by 0.162 fringe; the near-null role passes to (3,7), with 0.050047 fringe and 32.3 μm , and beyond path steps of 106 μm to (10,23), with 0.0116 fringe and 105.8 μm . For the prototype's (1,2) slip the density-only decision boundary lies at 0.0094 fringe, half of 0.0187, beyond the reach of any realistic density prior, while its 10.4 μm path step makes a path prior effective. The problem is two-dimensional, density and path, which places it in the integer-least-squares class with lattice decorrelation [11]. [M][N][R]

3.3 The polarimeter as an absolute channel

The TIP polarimeter launches collinear right- and left-handed circularly polarized beams at 10.59 μm and measures their phase difference 2ψ [6], with the Faraday angle

$$\psi = C_F \lambda^2 \int n_e B_{\parallel} dl, \quad C_F = \frac{e^3}{8\pi^2 \epsilon_0 m_e^2 c^3} \approx 2.63 \times 10^{-13} \text{ rad T}^{-1}$$

in the cold-plasma approximation. [D][R] As a density prior it reads $N_{\text{pol}} = \psi / (C_F \lambda^2 \bar{B}_{\parallel})$, with the density-weighted parallel field

$$\bar{B}_{\parallel} = \left(\int n_e B_{\parallel} dl \right) / \left(\int n_e dl \right)$$

taken from the equilibrium reconstruction along the chord. Its error budget combines angle noise, the uncertainty of $\bar{B}_{\parallel}$ from equilibrium and chord geometry, polarization effects such as residual Cotton–Mouton coupling, temperature corrections (Section 3.4), and the lag of its narrower bandwidth during transients. [D] The prototype resolves $1.5 \times 10^{19} \text{ T m}^{-2}$ at 100 Hz bandwidth, a phase resolution of 0.05° , against $1.5 \times 10^{18} \text{ m}^{-2}$ at 1 kHz for its vibration-compensated interferometer, quoted with 1.9° of phase noise [6]. [R] For an assumed $\bar{B}_{\parallel}$ of 1.5–2 T, the polarimeter noise corresponds to 0.036–0.047 fringe of density. [A][N] The angle stays far below a wrap, so the polarimeter supplies an absolute channel whose accuracy rests on $\bar{B}_{\parallel}$ and calibration. [M]

3.4 Temperature, timing and chromatic drift

At reactor temperatures the electron response turns relativistic. To first order in $\tau_e = T_e/(m_e c^2)$, the interferometric phase carries the factor

$$1 - \frac{3}{2} \tau_e$$

inside the line integral and the Faraday rotation the factor $1 - 2\tau_e$; at $T_e = 25 \text{ keV}$ these are corrections of 7.3 % and 9.8 %. Second-order terms add about 0.3–2 % to the interferometric phase and 0.5–3 % to the Faraday rotation at 20–50 keV [24]. A design study for the Japanese demonstration power plant, JA DEMO, finds deviations of about 7 % and 10 % on chords near the magnetic axis and recovers the density to within 1 % from the difference between interferometer and polarimeter signals [25]. [R][N] Left uncorrected, and for similar weighting along the chord, the corrections offset a polarimeter-based density prior from the interferometer by about $\tau_e/2$ of the line density, 0.12 fringe at 10^{21} m^{-2} and 25 keV, comparable to the lattice steps of Section 3.2. For ITER and DEMO the measurement model therefore carries temperature-weighted corrections. [M][N]

The channels arrive with different bandwidths and delays, so the estimator propagates each measurement to a common time base through its own response, and during a pellet the density change within the polarimeter's averaging window enters the prior error. [A] Window temperature, pressure and humidity change the optical path differently at the two wavelengths [23][26]. On the TIP prototype, heated BaF₂ and ZnSe test pieces shifted the vibration-compensated phase by 0.0121° and -0.275° per $^\circ\text{C}$ and per cm of thickness, so the temperatures of the secondary-confinement and primary vacuum windows have to stay within about 1°C and 2°C to meet ITER's density requirements passively [26]. [R] The vi-

bration-free combination passes a chromatic path difference δL_{12} into the density as $2\pi \delta L_{12} / (r_e(\lambda_1^2 - \lambda_2^2))$, 0.012 fringe per 0.1 μm for both pairs, the size of the prototype's near-null step. The estimator carries this drift as a slow state with its own prior. [M][N]

3.5 Detection and isolation in this register

With the chromatic drift calibrated and other nuisance inputs fixed, the register holds three scalar measurements, two interferometer phases and the polarimeter angle, for two continuous unknowns, density and path. Its parity space is then one-dimensional, the difference between polarimeter and interferometer density: the register detects density-relevant slips and assigns them to a density class, and the faulty channel stays unnamed, because each leave-one-out pair fits exactly. When the drift is a free state at each instant, three measurements meet three unknowns and the instantaneous parity space vanishes; an added path sensor restores one dimension and a drift prior, as pseudo-measurement, a second. Held constant over a window of T_w samples, the drift leaves $T_w - 1$ parity dimensions. [M][N]

The discrete structure of slips helps only where the lattice separates competing attributions by more than the residual noise: a CO₂ slip (1,0) and a double QCL slip (0, -2) differ in the residual by 0.0187 fringe for the prototype pair and by 0.162 fringe for the ITER pair. With the drift calibrated, an added path sensor gives two parity dimensions, yet with noise-weighted signatures the CO₂ and polarimeter faults meet at only about 7° in the example of Appendix B, so isolation power has to be computed for the actual noise levels and fault sizes. [M][N]

4. THE DENSITY-RISK DECISION

4.1 Posterior over integer hypotheses

Let S be the search set of integer hypotheses, bounded by the largest density change $N_{\max}$ and path change $L_{\max} = v_{\max} \Delta t_{\text{gap}}$ plausible across a gap of duration Δt_{gap} for a maximum rate of optical path change $v_{\max}$. For each hypothesis k in S , the lifted phases give a solution x_k for density and path. With a prior of mean x_{prior} and covariance Σ_p , and a measurement covariance Σ_m , the hypothesis weights are

$$w_k \propto \rho_k \exp \left(-\frac{1}{2} r_k^T \left(\Sigma_p + \Sigma_m \right)^{-1} r_k \right), \quad r_k = x_k - x_{\text{prior}}$$

with prior weights ρ_k , and under hypothesis k the state has a Gaussian posterior with mean

$$m_k = x_{\text{prior}} + \Sigma_p \left(\Sigma_p + \Sigma_m \right)^{-1} r_k$$

and covariance $C = \Sigma_p - \Sigma_p \left(\Sigma_p + \Sigma_m \right)^{-1} \Sigma_p$. [M]

These expressions assume Gaussian errors, with the prior error $x_{\text{prior}} - x$, for the true state x , independent of the measurement error ξ_m of the solutions x_k ; the residual covariance Σ of Section 2.3 is then $\Sigma_p + \Sigma_m$. Shared optics, measurements reused in both the prior and the likelihood, or a prior supplied by an observer that has already processed the same phases correlate the two errors. With their cross-covariance

$$\Sigma_{pm} = E \left[\left(x_{\text{prior}} - x \right) \xi_m^T \right],$$

the residual covariance and the gain become

$$\Sigma = \Sigma_p + \Sigma_m - \Sigma_{pm} - \Sigma_{pm}^T, \quad G = \left(\Sigma_p - \Sigma_{pm} \right) \Sigma^{-1}$$

the weights use Σ^{-1} in place of

$$\left(\Sigma_p + \Sigma_m \right)^{-1},$$

and the component posteriors become $m_k = x_{\text{prior}} + G r_k$ and $C = \Sigma_p - G \Sigma G^T$; for $\Sigma_{pm} = 0$ they reduce to the expressions above. A heavy-tailed likelihood changes the component posteriors as well as the weights: inserted into the weights alone, it leaves m_k and C tied to the Gaussian model it replaced. The component posteriors then follow from the same likelihood, for example through a scale mixture of Gaussians such as the Student-t distribution, or by numerical integration, and the density risk of Section 4.2 sums over them. [M]

An explicit out-of-set hypothesis o , with prior weight ρ_o and a broad likelihood, receives the posterior weight w_o ; the in-set likelihoods then keep their full Gaussian normalization, and the weights normalize over S and o together. Normalizing over an incomplete set alone gives confident wrong answers whenever the true count lies outside it, so a rising out-of-set probability triggers search expansion, and abstention follows when the expanded set still fails to explain the data. In practice, measured residual distributions replace the assumed ones, with weights and component posteriors computed under the same model, and the resulting probabilities are calibrated on discharges kept apart from the test set. [D] [C]

4.2 Density risk, its bound, and availability

The weights satisfy

$$\sum_{k \in S} w_k + w_o = 1.$$

.With $m_{k,N}$ the density component of m_k , $s^2 = C_{NN}$ and Φ the standard normal distribution function, the density risk of an estimate $\hat{N}_t$ is

$$q_t = P\left(\left|N - \hat{N}_t\right| > \epsilon \mid y_{1:t}\right) = \sum_{k \in S} w_k \eta_k + w_o \eta_o, \quad \eta_k = \Phi\left(\frac{\hat{N}_t - \epsilon - m_{k,N}}{s}\right) + \Phi\left(\frac{m_{k,N} - \hat{N}_t - \epsilon}{s}\right)$$

for a stated tolerance ϵ , with $y_{1:t}$ all measurements up to time t and η_o the probability of an error beyond tolerance under the out-of-set hypothesis, whose density the data leave unconstrained. Every η lies between 0 and 1, so evaluating the components of a subset $K \subseteq S$ and counting the remaining posterior mass as failure gives the conservative bound

$$\bar{q}_t = \sum_{k \in K} w_k \eta_k + \left(1 - \sum_{k \in K} w_k\right) \geq q_t$$

which holds for every K and every $\hat{N}_t$ and counts w_o in full. The estimate $\hat{N}_t$ is chosen to make $\bar{q}_t$ small (Section 4.4), and the density is available when $\bar{q}_t \leq \alpha$: availability gates on the bound, while calibration concerns the reported probabilities themselves, and Section 7.1 tests the two separately. The estimator reports three quantities separately: the probabilities of the integer hypotheses, for count recovery; the density with its bound $\bar{q}_t$, for density accuracy; and the path state with its own uncertainty, for vibration-state accuracy. An unresolved count can therefore coexist with an available density. [D]

4.3 Four comparisons

The decision mixes two potential benefits, a better estimate and a better gate, and the evaluation keeps them apart. [D]

1. Hard-branch estimate: the density of the most probable integer hypothesis, reported for every decision.
2. Marginalized estimate: $\hat{N}_t$, reported for every decision.
3. Density-risk gate: $\hat{N}_t$, reported only when $\bar{q}_t \leq \alpha$.
4. Count gate with the same estimate: $\hat{N}_t$, reported only for the decisions with the highest count confidence, or inside an integer aperture, at the availability of comparison 3; count confidence is the posterior probability of the most probable integer hypothesis.

Comparing 1 and 2 measures what the estimate gains. Comparing 3 and 4 measures what a criterion on the density adds over a criterion on the count, with the estimate and the availability held fixed. Comparing 2 and 3 shows the error rate that withholding reaches at reduced availability; a count gate can reach a similar rate, so this comparison alone leaves the value of the density criterion open. Hard branch selection gated by count confidence mixes both effects and serves as the integer-recovery baseline. Methods are compared at matched availability, the same fraction of decisions with a reported density, with identical inputs and latency. [D]

4.4 Computation

The candidate route used in Section 5 takes as K the eight heaviest components and as $\hat{N}_t$ the best of four candidates, the mixture mean and the means of the three heaviest components, judged by $\bar{q}_t$. [D] A numerical reference calculation, a grid search over the full mixture with golden-section refinement, gives a reference value for the smallest risk, without a certificate of global optimality. On 10^4 decisions per case, the bound and the full-mixture risk of the same estimate agreed to within 3×10^{-16} , the level of rounding, and the candi-

date estimate's risk exceeded the reference value by at most 4.2 % for the prototype (median 2.1 %) and matched it to within 0.05 % for the ITER pair. At a threshold inside the prototype's range of $\bar{q}_t$, $\alpha = 2 \times 10^{-5}$, the bound withholds a density that the reference value would accept in 4.0 % of the checked decisions. [N] An implementation offered for evaluation states its route and reports $\bar{q}_t$ as a bound. [C]

4.5 Relatives and what is proposed

Integer least squares with lattice decorrelation, the LAMBDA method, solves the coupled integer problem, and its success-rate theory gives the multi-dimensional margin of Section 2.3 [11][19]. Integer-equivariant estimators shift by the same integer vector as the real-valued solution they start from; best integer-equivariant estimation, the member of this class with the smallest mean squared error, weights all integer candidates and is, for Gaussian errors, the probability-weighted sum over all integers, which the marginalized estimate resembles [13]. Integer-aperture estimation accepts an integer solution only inside an aperture chosen for a controlled fail rate [14]; both belong to a unified theory of ambiguity resolution [12]. Selective prediction supplies the risk–coverage curve on which methods are compared at matched availability [27]. [R]

The element proposed here is the placement of the acceptance criterion on a tolerance for the density that a TIP-type diagnostic supplies to plasma control, together with the four comparisons and the evaluation protocol. In the configurations of Section 5 this criterion withheld the same decisions as count confidence, so its distinct value, beyond the marginalized estimate it gates, rests on measured data, which this note leaves for evaluation. [C]

5. A MODEL-CONDITIONAL ILLUSTRATION

5.1 Assumptions

The illustration uses Gaussian priors with a density spread $\sigma_N = 0.05$ fringe, comparable to the prototype polarimeter's noise (Section 3.3), and path spreads σ_L from 1 to 10 μm , or no path prior, with prior errors independent of the measurement errors; independent phase noise of 1.9° on each color, which maps to 0.0078 fringe, $1.6 \times 10^{18} \text{ m}^{-2}$, in density and 0.04 μm in path, with a correlation of -0.79 , for the prototype pair; a search set with density steps up to 3 fringes and path steps up to 60 μm , 105 hypotheses for the prototype pair and 127 for the ITER pair; a tolerance $\epsilon = 0.05$ fringe; and 2×10^6 simulated decisions per case. The true integer state lies inside the search set by construction; the simulation gives all hypotheses of the search set equal prior weights and the out-of-set hypothesis zero prior weight, so the weights normalize over the search set alone. The prototype quotes 1.9° for its vibration-compensated interferometer [6], so the per-color covariance is an assumption until measured. [A]

5.2 Branch-induced displacement and total error

For the prototype pair, with hard branch selection and the density taken from the phases of the selected branch:

Path prior σ_L	d for (1, 2)	Wrong count	Branch-induced density displacement above 0.05 fringe	Total density error above 0.05 fringe
none	0.37	0.85	0.36	0.32
10 μm	1.10	0.58	5.8×10^{-3}	8.3×10^{-3}
5 μm	2.11	0.29	none, below 1.5×10^{-6}	7.8×10^{-5}
2 μm	5.21	9.2×10^{-3}	none, below 1.5×10^{-6}	none, below 1.5×10^{-6}
1 μm	10.4	one in 2×10^6	none, below 1.5×10^{-6}	none, below 1.5×10^{-6}

[N] The displacement column counts only the density step of the selected branch; the total column adds the measurement error. Bounds for zero events are one-sided 95 % bounds. The total rate at 5 μm , 7.8×10^{-5} with a 95 % interval of $6.6\text{--}9.1 \times 10^{-5}$, comes mostly from selecting twice the near-null vector, in 137 of its 155 events, a displacement of 0.0374

fringe that measurement noise pushes across the tolerance. Three regimes appear. Without a path prior the near-null family spreads the density in steps of 0.0187 fringe, and the density inherits the full spread of the density prior. A coarse path prior of 5–10 μm leaves the count unresolved along the near-null direction and keeps the branch-induced displacement within tolerance. Exact count recovery, and with it the vibration state, needs a path prior near 2 μm or better. [M][N]

For the ITER pair every competitor in this search set displaces the density by at least 0.050 fringe, (3,7) by 0.050047 fringe, just beyond the tolerance. A wrong count therefore always produces a branch-induced displacement beyond tolerance, while the total error, which adds measurement noise, can fall inside it. Without a path prior, 62 % of decisions select a wrong count and 34 % exceed the tolerance in total; for path priors of 10, 5 and 2 μm the wrong-count rates are 0.18, 0.069 and 6.2×10^{-3} and the total-error rates 0.13, 0.069 and 6.2×10^{-3} . Where the two rates agree, they agree as a numerical result of the configuration: the dominant confusion at narrow path priors, (1,2), displaces the density by 0.162 fringe, far beyond the reach of the measurement noise. [N]

5.3 Four comparisons

Entries give the rate of total density errors above 0.05 fringe among the reported densities, with the number of events in parentheses. Every entry draws on the same 2×10^6 simulated decisions per configuration, so the columns form paired comparisons, and at each availability the gates report the same number of decisions, at 100 % every decision.

Configuration and availability	Hard-branch estimate, count gate	Marginalized estimate, count gate	Marginalized estimate, density-risk gate	Mean bound $\bar{q}_t$
Prototype, 5 μm path prior, 100 %	6.6×10^{-5} (132)	1.7×10^{-5} (33)	1.7×10^{-5} (33)	1.7×10^{-5}
Prototype, 5 μm path prior, 81.9 %, $\alpha = 2 \times 10^{-5}$	3.8×10^{-5} (63)	1.5×10^{-5} (24)	1.5×10^{-5} (24)	1.6×10^{-5}
ITER, 2 μm path prior, 100 %	6.2×10^{-3} (12,334)	6.2×10^{-3} (12,334)	6.2×10^{-3} (12,334)	6.2×10^{-3}
ITER, 2 μm path prior, 86.1 %, $\alpha = 10^{-3}$	6.3×10^{-5} (108)	6.3×10^{-5} (108)	6.3×10^{-5} (108)	7.4×10^{-5}

ITER, 2 μm path prior, 71.0 %, $\alpha = 10^{-4}$	1.2×10^{-5} (17)	1.2×10^{-5} (17)	1.2×10^{-5} (17)	1.4×10^{-5}
---	---------------------------	---------------------------	---------------------------	----------------------

[N] Exact Poisson 95 % intervals: 132 events, $5.5\text{--}7.8 \times 10^{-5}$; 63 events, $3.0\text{--}4.9 \times 10^{-5}$; 33 events, $1.1\text{--}2.3 \times 10^{-5}$; 24 events, $0.94\text{--}2.2 \times 10^{-5}$; 12,334 events, $6.1\text{--}6.3 \times 10^{-3}$; 108 events, $5.1\text{--}7.6 \times 10^{-5}$; 17 events, $0.70\text{--}1.9 \times 10^{-5}$. The hard-branch estimate takes the posterior density of the selected branch; taking the density directly from the phases gives the 7.8×10^{-5} of Section 5.2, on independent draws and within overlapping intervals. [N]

The estimate, comparisons 1 and 2. For the prototype the marginalized estimate cuts the error rate fourfold with every density available: on the same decisions, 112 errors of the hard-branch estimate disappear and 13 new ones appear. A rule that waits for a resolved count would have withheld every prototype density, because the count confidence stayed below 0.99 in all 2×10^6 simulated decisions. At 81.9 % availability the two count-gated columns differ, 63 against 24 events, and this difference also belongs to the estimate, since both use the same gate. For the ITER pair the estimate changes nothing: both estimates fail on the same 12,334 decisions. [N]

The gate, comparisons 3 and 4. With the marginalized estimate held fixed, the density-risk gate and the count gate accepted identical sets of decisions at every tested availability for both pairs, and therefore the same error counts. Neither tested configuration establishes a gating advantage of the density criterion over count confidence. Withholding itself, comparisons 2 and 3, does little for the prototype, whose bound lies between 1.3×10^{-5} and 2.0×10^{-5} for every decision, so that withholding 18 % of decisions lowers the rate from 1.7 to 1.5×10^{-5} , and much for the ITER pair, where withholding 14 % lowers it about a hundredfold with either gate. A rule that waits for a count confidence of 0.99 keeps 94.3 % of the ITER densities available at 3.6×10^{-4} , 678 events with an interval of $3.3\text{--}3.9 \times 10^{-4}$, with either estimate. Each mean bound lies inside the corresponding 95 % interval in this simulation. Bound validity is assessed one-sided, as specified in F1. [N]

The identity of the gates has a structural reason in both configurations. For the prototype the posterior concentrates on the near-null family, whose weights form a discrete Gaussian with a shape fixed by one scalar, the distance of its center from the nearest lattice point; as that distance grows, the count confidence falls and the bound rises, so the two rank the decisions alike. For the ITER pair every competitor with appreciable weight displaces the density beyond the tolerance, so the bound follows the posterior mass outside the most probable hypothesis, which is one minus the count confidence. The gates can part where the density spread varies between decisions, or where competitors with appreciable weight displace the density by different amounts, some inside the tolerance and some beyond it.

[M] The demonstrated benefit is therefore one of estimation: marginalization preserves useful density despite unresolved counts. A separate advantage of the density-risk gate remains a hypothesis for measured data, F2b of Section 7.1. [C]

5.4 Search bounds and the operating domain

For the path, $L_{\max} = v_{\max} \Delta t_{\text{gap}}$; an assumed rate of optical path change of 30 mm/s over gaps of 1, 2, 4 and 8 ms gives 30, 60, 120 and 240 μm . [A] With a 5 μm path prior, the rates of both pairs stay unchanged across these bounds within the simulation noise, because distant hypotheses carry negligible weight. Without a path prior the bound itself acts as the path information: for the prototype the total error rate is 2 % at 30 μm , where the near-null family stops at twice its vector, and 32 % from 60 μm upwards; for the ITER pair the smallest competitor changes from (1,2) at 30 μm to (3,7) at 60 μm and to (10,23) beyond 106 μm , which raises the wrong-count rate from 0.11 to 0.91. [N] The contrast between the pairs is therefore an operating-domain result, valid for the stated bounds and priors, and every rate needs its search bounds stated beside it. [C]

5.5 Reach of the illustration

The calculations establish a model-conditional mechanism: within a stated search set, and with assumed priors and noise, an unresolved count can coexist with an accurate density, and a marginalized estimate captures that. Two mechanisms carry different weights for the two pairs, a better estimate for the prototype and withholding for the ITER pair, and in neither configuration did the density-risk gate withhold differently from a count gate applied to the same estimate. Everything that measured data decide remains open: the real tails of the residuals, the correlations between priors, the per-color noise covariance, the frequency of out-of-set events, the calibration of the reported probabilities, and the latency budget. [D][C]

6. EVALUATION PROTOCOL

6.1 Data, replay and ground truth

The protocol uses raw interferometer–polarimeter records with documented fringe jumps during pellets, ELMs, disruptions and disruption mitigation, obtained through collaboration. Replay acts where count errors arise: on detector waveforms before demodulation, or on complex I/Q signals before phase extraction and unwrapping, by imposing amplitude fades, phase noise, fast phase ramps and interruptions with reacquisition; or on the tracker's count state, by corrupting it directly. A multiple of 2π added to an already wrapped phase leaves it unchanged, so replay at that level tests nothing. [D][C]

Implementation tests come first, and they examine two outputs separately. The uncorrected inversion, the density and path computed from the lifted phases with the counts as given, must return the lattice residuals of a (1,2) count-state injection on the prototype pair to numerical precision: a density offset of 0.0187 fringe and an optical path offset of 10.39 μm . The estimator's density must stay unchanged by the same injection, since its posterior over absolute counts stays the same while the search set covers both count states and the priors carry no information from the tracker's count state. Both checks test the inversion and the bookkeeping; because the residuals follow algebraically from the same inversion, they leave the physical two-color model untested. Physical tests need independent truth: bench optics without plasma, where the density is zero and an independent displacement sensor records the mirror displacement, or natural events identified by evidence independent of the two-color inversion. [D]

Ground truth for natural events comes from evidence independent of the estimator's inputs, such as Thomson scattering, a laser-scattering measurement of local electron density, other chords through a profile model, closure at the pulse end, or an offline reconstruction from signals outside the estimator's inputs, and its labels carry an acknowledged uncertainty that propagates into the metrics. The polarimeter and the deployed corrector enter as inputs or baselines. Training, calibration and test sets are split by complete discharge. [D][C]

6.2 Baselines and matched availability

Baselines: the deployed correction algorithm of the data source; continuity unwrapping [18]; integer least squares or state-space estimation [7][11]; integer least squares with integer-aperture acceptance at a fixed fail rate, and the best integer-equivariant estimate [13] [14]; recent machine-learning methods [8][10]. Every method receives the same sensor in-

puts and the same latency budget, and methods are compared at matched availability on risk–coverage curves [27]. [D][R]

6.3 Metrics

Exact count recovery, density accuracy at a stated tolerance and vibration-state accuracy, reported separately per event class; density error split into branch-induced displacement and continuous error; the four comparisons of Section 4.3; calibration of reported probabilities, shown as reliability diagrams, and validity of the bound $\bar{q}_t$, reported apart; out-of-set and abstention rates; recovery time after a jump; real-time latency; discharge-level risk under the union bound; zero-event results with one-sided 95 % bounds; search bounds stated with every rate. [D]

6.4 Requested from partner groups

Detector waveforms or complex I/Q signals of the interferometer and polarimeter, with time stamps and tracker states; equilibrium reconstructions for $\bar{B}_{\parallel}$ along each chord; Thomson scattering and other chords for independent truth; vibration or path sensor records where they exist; wavelength calibration and window temperature logs; event annotations for pellets, ELMs, disruptions and mitigation; and bench access for zero-density optical tests. Offline evaluation on archived records needs no change to an operating machine. [C]

7. VALIDATION LADDER, FALSIFIERS AND INVALIDATORS

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

Stage	Purpose	Primary outputs	Failure condition
0, this note	Model-conditional mechanism	Lattices, margins, parity accounting, synthetic comparisons, Appendix B	Mismatch between text and reproduction script
1	Implementation tests and raw-signal replay	Count-state injection checks; replay with controlled injections; the reference calculation over the full search set beside the bound $\bar{q}_t$	Lattice residuals missed in the uncorrected inversion; F1 fires on replay
1b	Bench optics	Zero-density tests with independent path truth	F3 fires
2	Prototype records, off-line	The four comparisons at matched availability against all baselines	F1 or F2 fires; retirement criterion met
3	Real-time shadow operation	Latency within the control budget; calibration stable across campaigns	Calibration drift; latency over budget

7.1 Model falsifiers

F1, calibration and bound validity. Two properties are tested separately in each event class. Calibration: the reported probabilities of the integer hypotheses and of the out-of-set hypothesis, and the density risk over the search set from the reference calculation, match the observed frequencies within their statistical uncertainty. Bound validity, assessed one-sided: among released densities, overall and in each bin of $\bar{q}_t$, the observed frequency of errors beyond ϵ stays at or below the mean reported bound within its statistical uncertainty. A conservative bound can stay valid while the probabilities are miscalibrated, so the two verdicts are reported apart: systematic miscalibration falsifies the estimator's error model, and an error frequency above the bound falsifies the availability rule as configured. [C]

F2, advantage at matched availability. With identical inputs and latency, and in the event classes where unresolved counts occur, two claims are tested separately. F2a, esti-

mate: the marginalized estimate gives risk–coverage curves for density accuracy below those of the hard-branch estimate under the same gate, of the deployed algorithm and of the machine-learning baselines. F2b, gate: with the marginalized estimate held fixed, the density-risk gate $\bar{q}_t \leq \alpha$ gives risk–coverage curves below those of gating by count confidence or by integer-aperture acceptance. Section 5.3 supports F2a in one synthetic configuration and finds the two gates identical in both, so F2b is an untested hypothesis for measured data. A failure of either part is reported as a null result. [C]

F3, two-color model. In controlled optical experiments with independent truth, the measured responses follow the two-color model. On the bench the density truth is zero, and an independent displacement sensor records the mirror displacement Δx , which changes the optical path by $\Delta L = 2 \Delta x$ in double pass. The mirror moves while both trackers lose and reacquire the signal, and two outputs are tested separately. The uncorrected inversion must show the lattice residuals: where the reacquired counts are off by (1,2) from the true counts, which the sensor and the zero density fix, a density offset of 0.0187 fringe and an optical path that departs from $2 \Delta x$ by 10.39 μm , equivalent to 5.20 μm of mirror displacement, for the prototype pair, within the combined uncertainty from phase noise, wavelength calibration and chromatic drift. The density residual scales with $\lambda_1 - 2\lambda_2$, so one nanometer in λ_2 changes it by 1.3 %. The final estimate must stay within ϵ of zero density at the rate its reported bound $\bar{q}_t$ allows, with a path state that agrees with $2 \Delta x$ within its reported uncertainty. Natural events count toward F3 only when evidence independent of the two-color inversion identifies them, because a path step and a density step inferred from the same inversion confirm each other by construction. Residuals that depart from the lattice beyond the combined uncertainty falsify the two-color model for that system, and final densities outside the tolerance more often than the reported bound allows count against the estimator under F1; the verdict rests on the uncertainty budget. [C][N]

F4, controlled aliasing. Phase ramps of known size $\Delta\phi$ per tracking step, injected into the raw signals, produce the apparent change $\Delta\phi_{\text{apparent}} = \text{wrap}(\Delta\phi)$: ramps between π and 2π per step appear reversed, and a ramp of 2π per step appears as stillness. Departures on controlled injections point to tracker behavior outside the nearest-branch model. On natural data, reversals also arise from fades, filtering and faulty reacquisition, so natural reversals test the event classifier, while the aliasing law is tested on injections. [C]

7.2 Experimental invalidators

Each of the following invalidates the affected comparison and calls for its repetition: ground truth derived from the estimator’s own inputs; training, calibration and test sets split within discharges; replay applied to wrapped phases; baselines given different inputs, latency or coverage; gates compared with different density estimates; rates reported with-

out their search bounds; zero-event results reported as zero rates; natural events identified through the same two-color inversion that they test; and probabilities on measured data normalized over a search set without an out-of-set hypothesis. [D]

7.3 Retirement criterion

If, on prototype records at matched availability with identical inputs and latency, the marginalized estimate fails to lower density errors relative to the hard-branch estimate under the same gate in the event classes where unresolved counts occur, the decision rule retires. If the estimate passes and F2b fails, the density-risk gate retires, and the marginalized estimate is gated by count confidence or integer-aperture acceptance instead. The accounting discipline remains useful in any case: count, density and vibration accuracy reported separately; displacement and continuous error kept apart; search bounds stated; zero-event results bounded; every continuous state and every prior counted in the observability accounting. [D]

8. SCOPE AND LIMITS

What the approach is. A decision rule and an evaluation protocol for density measurements from two-color interferometer–polarimeter systems: a posterior over integer hypotheses, a marginalized density estimate, a density risk with its conservative bound, an availability rule, and an accounting discipline. Its first testable domain is offline evaluation on records of the DIII-D prototype of ITER’s TIP. [D]

Where its claims end. The approach adds a decision on top of established measurement physics, and its claims end at the density measurement: plasma behavior, confinement, burn, materials and magnets lie outside it, and so does any statement about fusion performance. Its numbers are model-conditional; reliability on measured data, real-time feasibility at TIP rates, and benefit to density control remain to be demonstrated. DIP and other topologies call for their own lattices and models. [D]

Where it would pay, if the baselines are beaten. Discharge phases in which fringe jumps cluster, such as disruption mitigation and high-density fueling, where the marginalized estimate, released under a valid bound, keeps the control loop supplied while a rule that waits for a resolved count would blank it; and long pulses, in which discharge-level risk accumulates. [C]

Inherited limits. Gaussian priors and noise, with prior errors independent of measurement errors; an assumed per-color covariance; a true integer state inside the search set by construction, with zero prior weight on the out-of-set hypothesis; a cold-plasma model in the decision simulation, with the temperature corrections of Section 3.4 evaluated separately; a single chord without a profile model; and a candidate-based estimate with the conservative bound $\bar{q}_r$, checked against a numerical reference calculation. [D]

9. CONCLUSION

EN-003 moves one decision in the density-measurement chain from the integer count to the density itself. Its quantitative content is a chain with stated assumptions: the slip lattice of the instrument, the joint margin in the metric of density and path priors, the parity accounting with nuisance states and priors counted, and a density risk with a conservative bound on which availability is decided. In a Gaussian model the chain shows where the benefit lies: for the DIII-D prototype the marginalized estimate preserves useful density despite unresolved counts, for the ITER pair withholding lowers the error rate, and in both configurations the density-risk gate withheld the same decisions as a count-confidence gate applied to the same estimate.

Whether the estimation benefit survives measured data, and whether a density criterion gates better than a count criterion anywhere, are the note's open questions; the protocol, falsifiers and retirement criterion above state how to answer them. A clean loss retires the decision rule and leaves the accounting discipline intact. EN-003 is offered as the parent framework requires: precise enough to fail on a named instrument, at a stated tolerance, against a stated baseline.

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] Van Zeeland, M. A., Akiyama, T., Becoulet, M., and Kim, C. (2023). ITER Toroidal Interferometer and Polarimeter (TIP) beam refraction in 3D density profiles. *Fusion Engineering and Design* 193, 113618. <https://www.sciencedirect.com/science/article/abs/pii/S0920379623002028>
- [3] Innocente, P., Mazon, D., Joffrin, E., and Riva, M. (2003). Real-time fringe correction algorithm for the JET interferometer. *Review of Scientific Instruments* 74, 3645. Preprint: <https://scipub.euro-fusion.org/wp-content/uploads/2014/11/EFDP02001.pdf>
- [4] Juhn, J.-W., et al. (2010). Fringe-jump corrected far infrared tangential interferometer/polarimeter. *Review of Scientific Instruments* 81, 10D540. https://nstx.pppl.gov/DragNDrop/Publications/Presentations/Publications/2010%20Papers/2010_locked/Juhn_RSI.pdf
- [5] Akiyama, T., Kawahata, K., Tanaka, K., Tokuzawa, T., Ito, Y., Okajima, S., Nakayama, K., Michael, C. A., Vyacheslavov, L. N., Sanin, A., Tsuji-Iio, S., and LHD Experiment Group (2010). Interferometer systems on LHD. *Fusion Science and Technology* 58, 352. <https://scientific-publications.ukaea.uk/wp-content/uploads/Published/FusionSTVOL58p352.pdf>
- [6] Van Zeeland, M. A., et al. (2018). Tests of a full-scale ITER toroidal interferometer and polarimeter (TIP) prototype on the DIII-D tokamak (invited). *Review of Scientific Instruments* 89, 10B102. <https://pubs.aip.org/aip/rsi/article/89/10/10B102/364372/>
- [7] Blanken, T. C., Felici, F., Rapson, C. J., de Baar, M. R., Heemels, W. P. M. H., the TCV team, and the ASDEX Upgrade team (2018). Control-oriented modeling of the plasma particle density in tokamaks and application to real-time density profile reconstruction. *Fusion Engineering and Design* 126, 87–103. <https://www.sciencedirect.com/science/article/pii/S0920379617309031>
- [8] Wang, Y. F., Hanada, K., Sakurai, D., Liu, H. Q., Lan, T., Gao, X., and Wu, X. H. (2023). Phase jump detection and correction based on the support vector machine. *Plasma Physics and Controlled Fusion* 65, 065001. <https://doi.org/10.1088/1361-6587/accaa2>
- [9] Yao, Y., Yang, Y., Ti, A., Song, Y., Zhang, J., Wang, Y., Zhang, Y., Liu, H., and Jie, Y. (2024). A fringe jump counting method for the phase measurement in the HCN laser interferometer on EAST and its FPGA-based implementation. *Plasma Science and Technology* 26, 045601. <https://doi.org/10.1088/2058-6272/ad0dec>

- [10] Lu, Z., Ling, B., Yao, Y., Li, J., Wang, S., Lian, H., Liu, H., and Jie, Y. (2026). A random forest-based fringe jump prediction and correction method for laser interferometers on EAST. *Fusion Engineering and Design* 230, 115873. <https://www.sciencedirect.com/science/article/abs/pii/S0920379626002607>
- [11] Teunissen, P. J. G. (1995). The least-squares ambiguity decorrelation adjustment: a method for fast GPS integer ambiguity estimation. *Journal of Geodesy* 70, 65–82. <https://link.springer.com/article/10.1007/BF00863419>
- [12] Teunissen, P. J. G. (2003). Towards a unified theory of GNSS ambiguity resolution. *Journal of Global Positioning Systems* 2(1), 1–12. <https://gnss.curtin.edu.au/wp-content/uploads/sites/21/2016/04/Teunissen2003Towards.pdf>
- [13] Teunissen, P. J. G. (2003). Theory of integer equivariant estimation with application to GNSS. *Journal of Geodesy* 77, 402–410. <https://doi.org/10.1007/s00190-003-0344-3>
- [14] Teunissen, P. J. G. (2003). Integer aperture GNSS ambiguity resolution. *Artificial Satellites* 38(3), 79–88.
- [15] 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>
- [16] 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>
- [17] Dupke, A. (2026). STAR Drive: Sampled Transport by Alias Ratcheting — A Lock-Window Pulse-Scheduling Architecture for Conventional Propulsion (Engineering Note EN-002, Version 1.0). Zenodo. <https://doi.org/10.5281/zenodo.21833505>
- [18] Itoh, K. (1982). Analysis of the phase unwrapping algorithm. *Applied Optics* 21, 2470. <https://opg.optica.org/ao/abstract.cfm?uri=ao-21-14-2470>
- [19] Teunissen, P. J. G. (1998). Success probability of integer GPS ambiguity rounding and bootstrapping. *Journal of Geodesy* 72, 606–612. <https://link.springer.com/article/10.1007/s001900050199>
- [20] Chow, E. Y., and Willsky, A. S. (1984). Analytical redundancy and the design of robust failure detection systems. *IEEE Transactions on Automatic Control* 29, 603–614. <https://doi.org/10.1109/TAC.1984.1103593>

- [21] Van Zeeland, M. A., Carlstrom, T. N., et al. (2017). Tests of a two-color interferometer and polarimeter for ITER density measurements. *Plasma Physics and Controlled Fusion* 59, 125005. <https://doi.org/10.1088/1361-6587/aa8c49>
- [22] Bertin Technologies (2022). Final design of the ITER Density Interferometer Polarimeter (DIP) delivered to ITER Organization. <https://www.bertin-technologies.com/bertin-technologies-delivers-to-iter-organization-the-final-design-of-the-density-interferometer-polarimeter-dip/>
- [23] Bagnato, F., Sirinelli, A., Vayakis, G., Watts, C., Sengupta, R., and Carrat, V. (2025). ITER Density Interferometer Polarimeter: theoretical estimation of temperature, pressure and humidity variation effect on phase measurement and correction strategy. *Journal of Instrumentation* 20, C03001. <https://doi.org/10.1088/1748-0221/20/03/C03001>
- [24] Mirnov, V. V., Brower, D. L., Den Hartog, D. J., Ding, W. X., Duff, J., and Parke, E. (2013). Electron kinetic effects on interferometry and polarimetry in high temperature fusion plasmas. *Nuclear Fusion* 53, 113005. https://wippl.wisc.edu/pub_files/journal/Mirnov1017.pdf
- [25] Iwasaki, K., Sugiyama, S., Ohtani, Y., and Sakamoto, Y. (2026). Consideration of density measurement using toroidal interferometer and polarimeter on JA DEMO. *Fusion Engineering and Design* 225, 115638. <https://www.sciencedirect.com/science/article/abs/pii/S0920379626000256>
- [26] Akiyama, T., Van Zeeland, M. A., Finkenthal, D. K., LeSher, M., Finden, R., Trost, P., Gattuso, A., Miranda, S., Groh, M., de Looz, M.-A., and Watts, C. (2025). Impact of window temperature variations on ITER toroidal interferometer and polarimeter (TIP) measurements. *Plasma Physics and Controlled Fusion* 67(11), 115007. <https://doi.org/10.1088/1361-6587/ae1411>
- [27] El-Yaniv, R., and Wiener, Y. (2010). On the foundations of noise-free selective classification. *Journal of Machine Learning Research* 11, 1605–1641. <https://jmlr.org/papers/v11/el-yaniv10a.html>

APPENDIX A. NOTATION AND TERMS

Symbol	Meaning
$\tilde{\phi}, \phi, k$	Wrapped phase; true phase; branch integer
$h_m, x_m, \hat{k}_m, \Delta mn, S_{mn}, q_m$	Harmonic index; lifted phase; selected branch; wrapped and lifted residuals; count error of channel m
$d, \delta, \Sigma, b, \sigma_r$	Mahalanobis distance; lattice vector between hypotheses; covariance of the decision residual; scalar margin; residual spread
$p, p_i, n, P_{\text{any}}$	Per-repair probability, equal or individual; repairs per discharge; probability of at least one wrong repair
$\mathbf{y}, H, \mathbf{u}, \mathbf{f}, \mathbf{e}, \Pi, \mathbf{r}$	Measurements; measurement matrix; continuous unknowns; faults; noise; parity matrix; parity residual
M, n_w, T_w, n_f, n_s	Scalar measurements; continuous unknowns; window length; states free per sample; states shared across the window
N, N_F, n_e, r_e	Line-integrated density; line density per fringe; electron density; classical electron radius
$\lambda, \lambda_1, \lambda_2$	Wavelength; first and second color
$\Delta L, \Delta x, \delta N, \delta L, \delta L_{12}$	Optical path change; mirror displacement, with $\Delta L = 2 \Delta x$ in double pass; density and path steps of a slip (k_1, k_2); chromatic path difference
$\psi, C_F, B\ , \bar{B}\ , N_{\text{pol}}$	Faraday angle; Faraday constant; parallel magnetic field; its density-weighted mean along the chord; polarimeter density
e, ϵ_0, m_e, c	Elementary charge; vacuum permittivity; electron mass; speed of light
T_e, τ_e	Electron temperature; $T_e/(m_e c^2)$
$S, N_{\text{max}}, L_{\text{max}}, \nu_{\text{max}}, \Delta t_{\text{gap}}$	Search set; its density and path bounds; maximum rate of optical path change; gap duration
$x, x_k, x_{\text{prior}}, \xi_m$	True state; solution under hypothesis k ; prior mean; measurement error of the solutions

$\Sigma_p, \Sigma_m, \Sigma_{pm}, G$	Prior and measurement covariances; cross-covariance of prior and measurement errors; gain
$r_k, \rho_k, w_k, o, \rho_o, w_o$	Residual; prior weight; posterior weight; out-of-set hypothesis; its prior and posterior weights
$m_k, m_{k,N}, C, s, \Phi$	Posterior mean under hypothesis k ; its density component; posterior covariance; density spread; standard normal distribution function
$q_t, \bar{q}_t, \eta_k, \eta_o, K$	Density risk; its conservative bound; probability of an error beyond tolerance under hypothesis k and under the out-of-set hypothesis; evaluated components
$\hat{N}t, \epsilon, \alpha, y1 : t$	Density estimate; tolerance; risk level for availability; measurements up to time t
σ_N, σ_L	Spreads of the density and path priors
$\Delta\phi, \Delta\phi_{\text{apparent}}$	Injected phase ramp per tracking step; its apparent value after wrapping

Terms

Best integer-equivariant estimation. The integer-equivariant estimator with the smallest mean squared error; it weights all integer candidates, and for Gaussian errors it is the probability-weighted sum over all integers.

Branch. The integer that turns a wrapped phase into the true phase.

Count confidence. Posterior probability of the most probable integer hypothesis.

Count gate. Release of the density only for the decisions with the highest count confidence, or inside an integer aperture.

DEMO. A demonstration fusion power plant designed to follow ITER; JA DEMO is the Japanese design.

Density risk. Posterior probability q_t that the density error exceeds the tolerance; $\bar{q}_t$ is its conservative bound.

Density-risk gate. Release of the density only when $\bar{q}_t \leq \alpha$.

Disruption, disruption mitigation. A sudden loss of plasma confinement; mitigation injects material, for example shattered pellets, to soften its consequences and briefly raises the density to very high values.

ELM. Edge-localized mode, a periodic relaxation of the steep edge pressure gradient in high-confinement plasmas.

Fringe, fringe jump. One full cycle, 2π , of interferometer phase; a whole-fringe error in the accumulated count.

Hard branch selection. Choosing the most probable integer hypothesis and reporting the density it implies.

I/Q signals. In-phase and quadrature components of a demodulated detector signal; the angle of the complex signal they form is the wrapped phase.

Integer-aperture estimation. Accepting an integer solution only inside an aperture region chosen for a controlled fail rate.

Integer-equivariant estimation. Estimation whose result shifts by the same integer vector as the real-valued solution it starts from.

Lifted phase. A wrapped phase plus 2π times its selected branch.

Marginalized estimate. Density estimate that weights all integer hypotheses by their posterior probabilities, chosen to make the bound $\bar{q}_t$ small.

Matched availability. Comparison of methods at the same fraction of decisions with a reported density.

Near-null direction. Lattice vector of count errors that moves the density little and the inferred path strongly.

Optical path. Path length of the probing beam weighted by the refractive index; a mirror displacement Δx along a double-pass beam changes it by $2 \Delta x$.

Out-of-set hypothesis. Explicit hypothesis that the true integer state lies outside the search set.

Parity space. Space of residuals free of the continuous unknowns; its dimension is the number of measurements minus the number of unknowns.

Pellet injection. Injection of frozen fuel pellets for fueling or ELM pacing, which raises the density abruptly.

Polarimeter. Diagnostic measuring the Faraday rotation of a probing beam, proportional to the density weighted by the parallel magnetic field.

Pseudo-measurement. A prior on a state written as an extra row of the measurement model.

QCL. Quantum-cascade laser, the 5.22 μm source of the DIII-D prototype's second color.

Risk-coverage curve. Error rate plotted against the fraction of decisions accepted.

Search set. Finite set of integer hypotheses bounded by the largest density and path changes plausible across a gap.

Thomson scattering. Laser-scattering diagnostic of local electron density and temperature, independent of the interferometer.

TIP, DIP. ITER's Toroidal Interferometer and Polarimeter, the primary density-control diagnostic, and its Density Interferometer Polarimeter, a complementary dispersion interferometer.

Tokamak. Toroidal magnetic-confinement device; DIII-D and ITER are tokamaks.

Uncorrected inversion. Density and path computed from the lifted phases with the counts as given.

Wrap. Map of a phase onto its principal interval, $(-\pi, \pi)$.

APPENDIX B. REPRODUCTION SCRIPT AND MANIFEST

The script reproduces every [N] value of this note; the release package distributes it as `en003_reproduce.py`, with pinned dependencies in `requirements.txt` and the output of the reported run in `reference_output.txt`. It fixes all random seeds, one seed family per section, uses the same model throughout, and checks the candidate route against the numerical reference calculation on a subsample. Runtime is about four minutes on two processor cores. Random streams of numpy generators can differ between numpy versions, and floating-point rounding between platforms, so stochastic values may move within their stated intervals; the reported run used CPython 3.11.15 with numpy 2.4.4 and scipy 1.17.1. [N]

Code availability. The release package is archived with this note at doi.org/10.5281/zenodo.23105860. Its `README.pdf` gives the verification commands, and `SHA256SUMS` lists the checksum of every other file; in the verified environment the script reproduces `reference_output.txt` line by line. [N]

Manifest. Section 2.3: discharge-level risk and zero-event bound. Section 3.1: line density per fringe. Section 3.2: slip lattices for both pairs, the density-only boundary and the (3,7) displacement. Sections 3.3 and 3.4: polarimeter resolution in density units, temperature corrections, chromatic leakage, wavelength sensitivity of the near-null step. Sections 2.4 and 3.5: parity dimensions and the noise-weighted signature angle. Section 5.1: noise mapping. Section 5.2: decision table, Poisson interval and failure breakdown at 5 μm , ITER rates. Sections 4.4 and 5.3: four comparisons with exact Poisson intervals, paired error counts, the identity of the accepted sets of both gates, and the reference check. Section 5.4: search-bound sensitivity. [N]