Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Analytical observability and decision margins for automatic transfer in active distribution networks

Rui Zhou

767874631@qq.com

Huaian Power Supply Branch of State Grid Jiangsu Electric Power Co., Ltd.

Jian Ji

Huaian Power Supply Branch of State Grid Jiangsu Electric Power Co., Ltd.

Guoming Yan

Huaian Power Supply Branch of State Grid Jiangsu Electric Power Co., Ltd.

Jin Zhu

Huaian Power Supply Branch of State Grid Jiangsu Electric Power Co., Ltd.

Peng Li

Huaian Power Supply Branch of State Grid Jiangsu Electric Power Co., Ltd.

Zhengxiang Ding

Huaian Power Supply Branch of State Grid Jiangsu Electric Power Co., Ltd.

Article

Keywords:

Posted Date: October 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11115196/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Analytical observability and decision margins for automatic transfer in active distribution networks

Rui Zhou*, Jian Ji , Guoming Yan , Jin Zhu , Peng Li , and Zhengxiang Ding

Huaian Power Supply Branch of State Grid Jiangsu Electric Power Co., Ltd., Huaian 223001, China
*Correspondence: Rui Zhou (767874631@qq.com).

ABSTRACT

Distributed generation can sustain bus voltage after source separation, while measurement faults can indicate voltage loss without an outage. An analytical framework links voltage observability, hierarchical evidence admission and transfer-command stability. The established voltage non-detection-zone boundary is specialised to automatic transfer, and closed-form conditions describe candidate admission and command perturbation margins. Independent integration and deterministic decision grids verify these conditions. A separately implemented averaged-network benchmark evaluates the sampled-signal pipeline with calibrated memberships, measurement faults and impaired communication. Calibration, development and final evaluation use 162, 108 and 324 disjoint physical trajectories, respectively. Across 1,296 evaluation variants, the full method produced 5 false actions in 1,008 non-action cases, compared with 144 for conventional voltage/current logic. However, 101 of 288 required actions remained unissued within 650 ms after disturbance. Median latency among successful required requests was 57 ms, compared with 319 ms for conventional logic; these are logical request times, not restoration times. Removing quality controls reduced missed requests but increased false actions. The analysis and benchmark distinguish observable voltage loss, admissible evidence and stable commands, while identifying the communication conditions that limit automatic-transfer availability.

1 Introduction

Automatic transfer restores supply by isolating an unavailable working source and connecting a standby source. In active distribution networks, inverter-interfaced distributed generators (IIDGs) can sustain local voltage after source separation, delaying the undervoltage confirmation used by conventional transfer logic. Distributed generation changes power-flow direction and protection behaviour[1–4], while interconnection requirements distinguish disturbance ride-through from unintentional islanding[5]. The resulting decision problem is not simply whether voltage is low. A supported island can retain voltage after utility supply is lost, whereas a potential-transformer (PT) secondary-circuit fault can indicate voltage loss while the primary circuit remains energised. Under light load, low incomer current does not resolve this ambiguity.

Research on restoration and protection has addressed different parts of this problem. Storage-assisted automatic transfer coordinates energy support with changeover[6], while reconfiguration studies examine network configuration under uncertain loads at a broader operational scale[7]. Adaptive protection studies examine changing operating conditions[8], and microgrid protection research analyses the limitations imposed by converter-interfaced generation[9]. Converter current limiting during unbalanced disturbances further affects the signals supplied to protection[10]. More recently, supervised learning has combined fault detection and classification while accounting for fault-resistance uncertainty[11]. These approaches address restoration, fault discrimination or converter behaviour, but a transfer device still needs to distinguish source loss from measurement failure and resolve incompatible action requests before closing a standby incomer.

Islanding detection provides several ways to obtain evidence of source separation. Reviews distinguish passive, active and communication-based techniques and compare their non-detection zones, perturbation requirements and implementation burdens[12,13]. Among active methods, maximum-power-point-tracking injection introduces a detection perturbation through photovoltaic control[14]. A two-level method combines a voltage-change trigger with a subsequent converter perturbation[15]. Negative-sequence current injection extends detection to grid-following and grid-forming inverters[16], while reinforcement-learning-driven active frequency drift adapts its perturbation parameters to operating conditions[17]. These methods improve the evidence available at the inverter, but depend on converter participation. Their detection outputs do not by themselves establish that a transfer command is admissible or that electrical closing permissions are satisfied.

Data-driven detection uses complementary information from measured signals. A wavelet-scalogram and histogram-of-oriented-gradients framework constructs fused features for islanding classification[18]. Another approach combines unscented Kalman filtering with a deep neural network[19]. Together with fault-resistance-aware classification[11], these studies illustrate how signal processing and learned discrimination can address operating variability. Their reported performance concerns the

tested signal populations and models. For the present transfer application, an additional requirement is an explicit response when an observation is invalid, evidence conflicts, or a candidate lies too close to an action boundary. This motivates analysing decision conditions separately from classifier performance.

The voltage limitation is already established in the non-detection zone (NDZ) of overvoltage/undervoltage protection (OVP/UVP). Zeineldin and Kirtley derived NDZ expressions for voltage- and frequency-dependent loads and different generator controls[20]. Setting-map methods jointly examine pickup and delay across islanding and non-islanding events[21]; subsequent work incorporates undervoltage blocking[22]. Consequently, the boundary $\lambda_{\mathrm{crit}} = U_{\mathrm{set}}^{\alpha}$ used here is a one-sided, fixed-frequency specialisation of known power-balance theory, not a new NDZ law. Its role is to identify the conditions under which voltage evidence becomes unavailable and another source of transfer evidence is required.

Combining that evidence requires attention to both conflict and temporal validity. Reliability- and uncertainty-based adjustment of basic probability assignments can precede combination under conflicting evidence[23]. Temporal fusion methods instead account for variation within a sequence of observations[24]. For AC microgrid frequency regulation, cyber-resilient control combines false-data-injection detection with observer-based correction and collaborative mitigation[25]. These contributions motivate explicit checks on evidence quality, but they do not justify treating every input as interchangeable. An aligned breaker transition and current loss describe an event; voltage-envelope characteristics express graded support. Missing remote voltage must not be substituted by zero, which has the different physical meaning of a dead line. Rejection methods allow a system to withhold an insufficiently supported decision[26]. Transfer arbitration additionally requires deterministic constraints on the final command.

To address these connected requirements, this paper develops an analytical framework linking undervoltage observability, hierarchical candidate admission and command stability. Thirteen features are organised into four first-match layers, with calibrated memberships representing graded support. Recent fuzzy-rule explanations illustrate the interpretability of readable rules[27]; the present calibration and arbitration rules are specified independently. A two-stage voltage envelope identifies when local undervoltage confirmation is reachable. Closed-form thresholds relate evidence support to communication quality, and arbitration conditions quantify command exclusivity and perturbation margins. Independent numerical integration and deterministic grids verify these relationships. A separate averaged-network benchmark then tests the complete sampled-signal implementation against conventional logic and four component ablations. This design distinguishes algebraic guarantees from computed benchmark behaviour without assuming that a unique command is necessarily a correct or timely command.

2 Analytical framework and residual-voltage model

The analysis begins with the measurement paths available to a transfer device and a voltage envelope that represents temporary IIDG support. Power balance establishes the observability boundary, while an independent event integrator checks the envelope and threshold-crossing calculations. Electrical closing permissions are then distinguished from the logical evidence examined in subsequent sections.

2.1 Measurement architecture and study settings

Figure 1 shows the measurement and communication architecture of the utility project under study. Substations B and C supply substation A through two incomers. The working incomer 1DL is normally closed and standby incomer 2DL is open. Breaker 3DL provides the bus-section connection used in the transfer configuration. Local acquisition supplies voltage, current, active power and breaker position. Remote link units transmit line-voltage and breaker-state information over fibre.

The architecture of the utility project under study motivates a 10-kV example with a 7-MW nominal load and four 4-MW IIDGs. Table 1 lists the parameters used to instantiate the analytical study. Installed capacity divided by nominal load is 2.286, but instantaneous penetration is defined separately as $\lambda = P_{\mathrm{DG}}/P_{\mathrm{L0}}$. This distinction prevents installed capacity from being confused with the generation operating point. The drawing defines measurement paths rather than a complete electrical network model; converter controls and feeder impedances are outside the envelope representation.

2.2 Two-stage envelope and its domain of validity

The load follows the established exponential representation reviewed by El-Fergany[28]. Its active-power relationship and the capped equilibrium used here are

$$P_{\mathrm{L}}(U) = P_{\mathrm{L0}} U^{\alpha}, \qquad U_{\infty} = \min\{1, \lambda^{1/\alpha}\}. \tag{1}$$

Here, U is per-unit voltage and $\alpha > 0$ is the load exponent. Below the voltage cap, U_{∞} follows from active-power balance. At the cap, the expression assumes regulation or curtailment accommodates excess generation; it does not model that control explicitly.

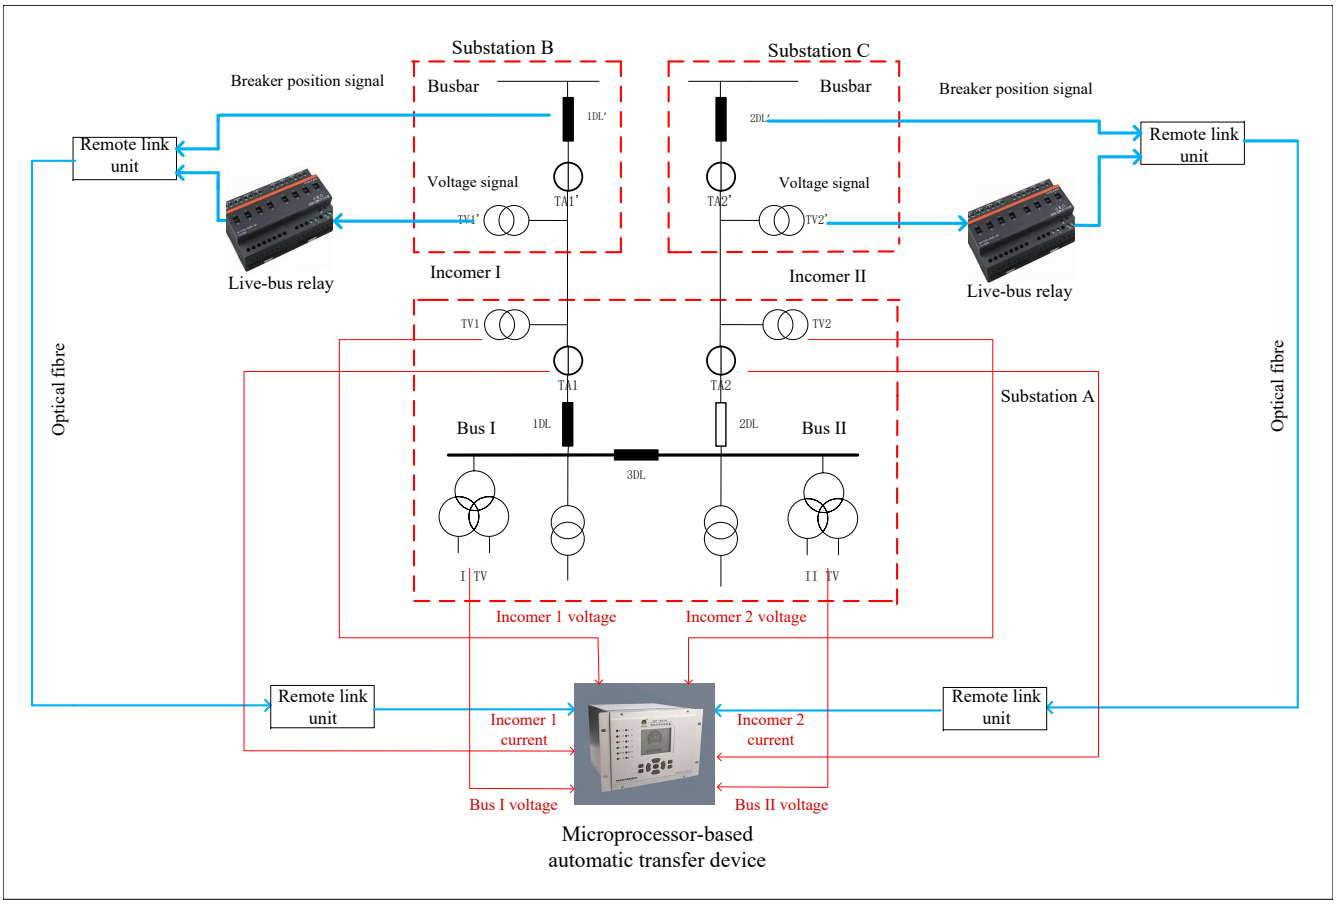


Figure 1. Measurement and communication architecture of the automatic-transfer system. Substations B and C feed the two incomers at substation A. Blue paths transmit remote breaker and voltage information; red paths deliver local voltage and current measurements. The topology, electrical symbols and device photographs are retained from the project Visio drawing.

Table 1. Parameters of the illustrative network and analytical decision study. Ride-through values and decision gates are design assumptions.

Group	Parameter	Setting
Network	Rated voltage; load	10 kV; 7 MW
Network	IIDG installed capacity	4×4 MW; capacity/load ratio 2.286
Voltage setting	Nominal no-voltage threshold	0.25 p.u.
Voltage envelope	Load exponent; time constants	$\alpha = 1.8$; $\tau_1 = 120$ ms; $\tau_2 = 20$ ms
IIDG disconnection	LVRT threshold; hold; execution	0.88 p.u.; 300 ms; 60 ms
Calibration rule	Core quantiles	0.10 and 0.90
Membership	Outer expansion; weight exponent	$\kappa = 1.5$; $\gamma = 2.0$
Decision gates	Credibility; maximum uncertainty	0.60; 0.30
Arbitration	Action and block ratios; latch	$k_a = 1.30$; $k_b = 0.80$; 40 ms
Delayed strategy	Permitted adaptive delay	40–600 ms

During stage I, connected IIDGs support the local voltage. A first-order envelope approximates the transition from the initial value of 1 p.u.:

$$U_1(t) = U_\infty + (1 - U_\infty)\exp(-t/\tau_1). \tag{2}$$

The first crossing of the study low-voltage ride-through (LVRT) threshold is

$$t_{\mathrm{L}} = \begin{cases} -\tau_1 \ln\left(\dfrac{U_{\mathrm{LVRT}} - U_\infty}{1 - U_\infty}\right), & U_\infty < U_{\mathrm{LVRT}}, \\ \infty, & U_\infty \geq U_{\mathrm{LVRT}}. \end{cases} \tag{3}$$

When t_{L} is finite, the model sets $t_{\mathrm{trip}} = t_{\mathrm{L}} + 0.300 + 0.060$ s. The complete envelope is

$$U(t) = \begin{cases} U_1(t), & t < t_{\mathrm{trip}}, \\ U_1(t_{\mathrm{trip}})\exp[-(t - t_{\mathrm{trip}})/\tau_2], & t \geq t_{\mathrm{trip}}. \end{cases} \tag{4}$$

For $t_{\mathrm{trip}} = \infty$, only stage I is present. A finite collapse above the LVRT equilibrium threshold therefore requires a separate anti-islanding or disconnection mechanism. Its delay cannot be inferred from Eqs. (2)–(4).

2.3 Undervoltage observability and its relation to the NDZ

The no-voltage criterion is reachable during stage I only when $U_\infty < U_{\mathrm{set}}$. For a setting below 1 p.u., the resulting boundary is

$$\lambda < \lambda_{\mathrm{crit}} = U_{\mathrm{set}}^{\alpha}. \tag{5}$$

Equation (5) expresses the lower-voltage boundary of the established voltage-dependent-load NDZ in generation-to-load coordinates[20]. To make the correspondence explicit, consider the uncapped power-balance model at fixed frequency and define $\Delta p = (P_{\mathrm{L0}} - P_{\mathrm{DG}})/P_{\mathrm{DG}} = \lambda^{-1} - 1$. A voltage no-trip band $U_{\min} \leq U \leq U_{\max}$ then becomes

$$U_{\min}^{\alpha} \leq \lambda \leq U_{\max}^{\alpha} \quad \Longleftrightarrow \quad U_{\max}^{-\alpha} - 1 \leq \Delta p \leq U_{\min}^{-\alpha} - 1. \tag{6}$$

This coordinate mapping is a specialisation, not the full voltage–frequency NDZ. The upper-voltage branch is shown only for comparison: the capped envelope in Eq. (1) does not describe overvoltage. The transfer no-voltage setting of 0.25 p.u. and the model LVRT threshold of 0.88 p.u. serve different functions and are not interchangeable or universal protection settings. The analysis uses the lower boundary to determine when local voltage can contribute to a transfer decision, then connects that condition to admission and arbitration.

Equality is excluded because an exponential approach to the threshold does not cross it in finite time. Below the boundary, the stage-I crossing time is

$$t_U = -\tau_1 \ln\left(\frac{U_{\mathrm{set}} - U_\infty}{1 - U_\infty}\right). \tag{7}$$

This expression applies only if crossing occurs before IIDG disconnection. Above the boundary, another mechanism must remove voltage support before undervoltage confirmation becomes possible.

At $U_{\mathrm{set}} = 0.25$ and $\alpha = 1.8$, direct calculation gives $\lambda_{\mathrm{crit}} = 0.08247$. The corresponding generation is 0.577 MW, or 3.61% of the installed 16-MW capacity. Thus, the loss of stage-I observability can arise well below the installed-capacity ratio usually associated with high penetration. The boundary shifts to 0.05519 and 0.11450 at settings of 0.20 and 0.30 p.u., respectively.

Figure 2a evaluates the stated envelope without imposing an arbitrary disconnection time. The $\lambda = 0.05$ and 0.30 cases reach the LVRT threshold and enter stage II. The $\lambda = 0.80$ and 1.20 cases remain supported because their equilibria exceed 0.88 p.u. This difference is important: an undervoltage element cannot detect every supported island simply by reducing its confirmation delay.

The sensitivity curves in Fig. 2c and their dependence on load exponent follow directly from

$$\frac{\partial \lambda_{\mathrm{crit}}}{\partial U_{\mathrm{set}}} = \alpha U_{\mathrm{set}}^{\alpha-1} > 0, \qquad \frac{\partial \lambda_{\mathrm{crit}}}{\partial \alpha} = U_{\mathrm{set}}^{\alpha} \ln U_{\mathrm{set}} < 0. \tag{8}$$

The derivatives at the study point are approximately 0.594 and -0.114, respectively. Raising the setting expands the observable region, whereas increasing the load exponent contracts it. Neither change removes the dependence on island power balance.

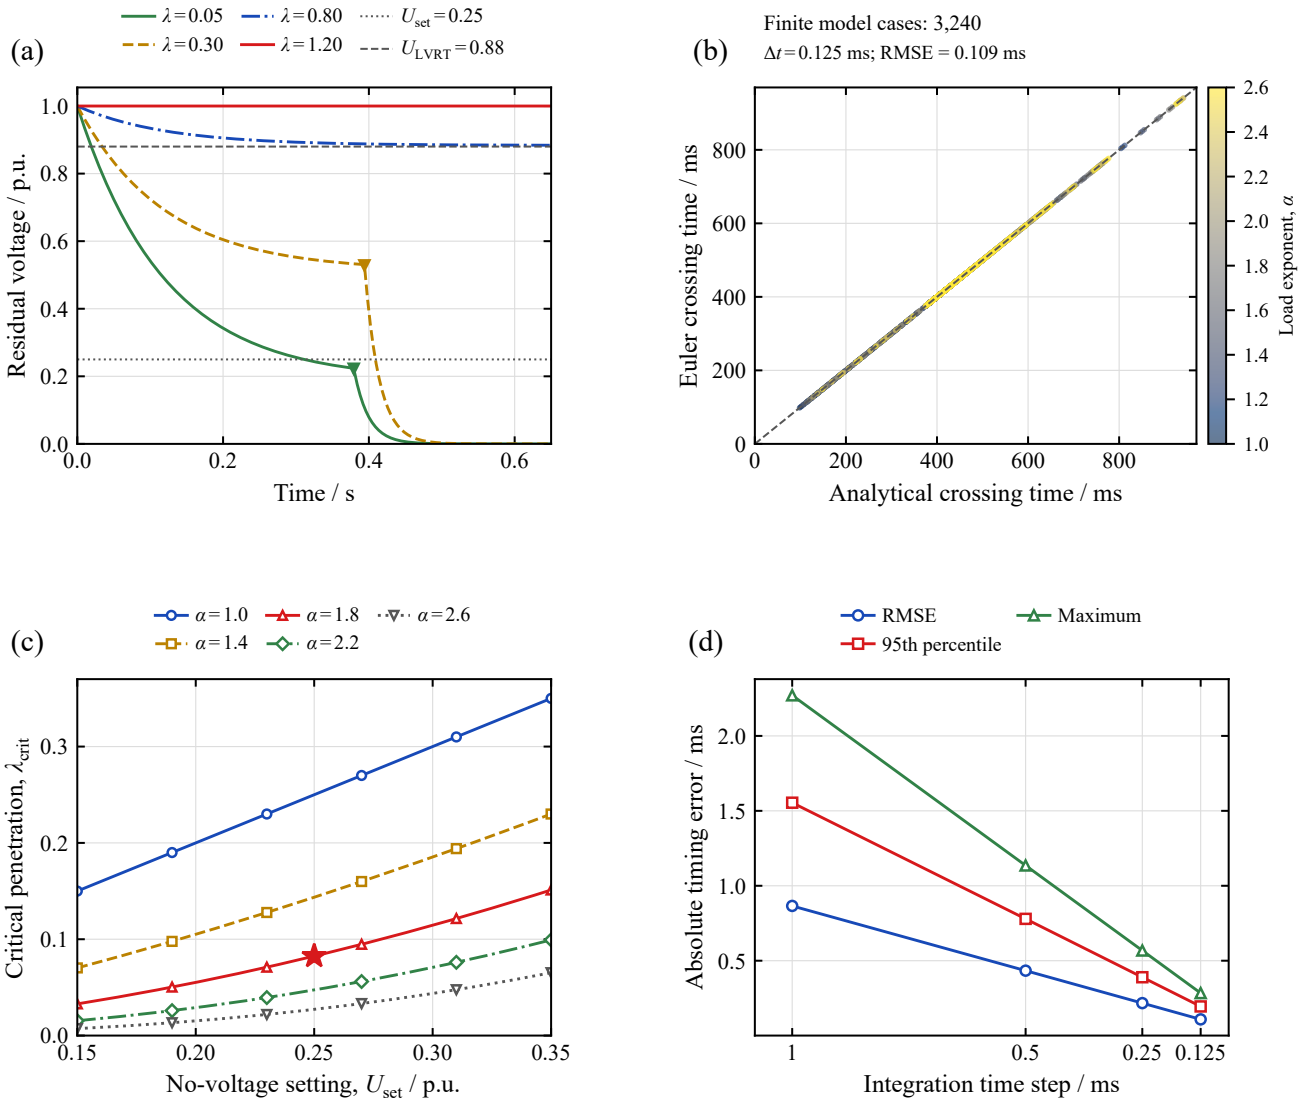


Figure 2. Two-stage residual-voltage model, numerical verification and parameter sensitivity. (a) Envelopes from Eqs. (2)–(4); triangles mark finite IIDG-disconnection times. (b) Analytical versus forward-Euler crossing times for 3,240 finite-crossing cases from a deterministic 4,941-combination grid. Colour denotes load exponent; the dashed line is the identity, not a fitted regression. (c) Critical penetration against voltage setting for five load exponents; the star marks the study point. (d) Timing-error convergence at four integration steps. Percentiles describe the parameter grid, not confidence intervals. Non-crossing cases remain in the accompanying case-level data.

2.4 Numerical verification of the voltage model

To check the envelope and its switching events independently, the analysis compared analytical crossing times with forward-Euler integration over 4,941 parameter combinations. The deterministic design varies load exponent, instantaneous penetration, voltage setting and stage-I time constant; The grid uses nine equally spaced load exponents from 1.0 to 2.6, 61 penetrations from 0.01 to 1.20, settings of 0.20, 0.25 and 0.30 p.u., and stage-I time constants of 0.08, 0.12 and 0.18 s. The remaining parameters follow Table 1. Of these combinations, 3,240 cross the no-voltage threshold within 2 s and 1,701 do not. Crossing status agrees at all four integration steps, including cases that do not enter stage II.

At the finest step of 0.125 ms, the crossing-time root-mean-square error is 0.109 ms and the maximum absolute error is 0.284 ms (Fig. 2b,d). Halving the step from 1 to 0.125 ms produces observed convergence orders of 0.997–0.999, consistent with first-order integration. The case-level data retain signed residuals and the complete crossing/non-crossing population. The integration advances $\dot{U} = (U_\infty - U)/\tau_1$ before disconnection and $\dot{U} = -U/\tau_2$ afterwards. Linear interpolation locates the first LVRT crossing, schedules disconnection 0.360 s later, and locates the no-voltage crossing. A step containing disconnection is split between the two stages; the integrator does not call analytical event-time expressions. If disconnection precedes the setting crossing, the reference time is $t_{\text{trip}} + \tau_2 \ln[U_1(t_{\text{trip}})/U_{\text{set}}]$. Unreachable crossings are stored as missing times, not as the 2-s horizon. Error statistics exclude these cases; $\log_2(E_h/E_{h/2})$ defines the observed RMSE convergence order. This comparison verifies numerical consistency of the stated voltage model; it does not measure controller decision accuracy or validate the model against physical observations.

2.5 Scope of electrical closing permission

Classification and credibility assessment produce a logical request. Electrical closing additionally requires working-source isolation, fault clearance, standby-source availability and valid breaker feedback. Where an energised island could be connected to another source, dead-bus or synchronism permission and the applicable IIDG isolation sequence must remain effective. The present analysis conditions on these permissions and does not predict a complete restoration time.

3 Hierarchical evidence and coordinated decision method

Voltage observability determines whether a local measurement can confirm source loss, but does not select a transfer command. The following method first checks observation validity, then assigns a causal candidate through ordered evidence layers. Calibrated support, communication-dependent uncertainty and arbitration gates determine whether that candidate can request action or block, or must remain under review.

3.1 Synchronous observations and measurement validity

Sampled-value protection studies examine packet loss and loss of synchronisation[29]. Such tests motivate explicit temporal-validity checks, rather than establish the numerical settings adopted here. A validity mask v_j identifies whether feature j is usable. Stale or discontinuous data receive $v_j = 0$; their physical values are not replaced by zero. For the utility project under study, the communication-quality rule is

$$q_c = \exp(-k_d \bar{d}/d_0)(1 - p_{\text{loss}})\mathbf{1}[a \le a_{\max}], \tag{9}$$

where $\bar{d}$ is mean delay, d_0 its reference, a data age and p_{loss} packet-loss fraction. The delay sensitivity is $k_d = 1$. The analysis varies q_c directly over $[0, 1]$. Mapping measured network delays to this scalar requires application-specific choices of d_0, $a_{\max}$ and the averaging window.

The thirteen features comprise bus-voltage r.m.s. (F1), voltage gradient (F2), residual-envelope time constant (F3), incomer-current r.m.s. (F4), current step (F5), and negative-sequence voltage ratio (F6). They also include power direction (F7), switch/current timing consistency (F8), pseudo-voltage-loss evidence (F9), IIDG output variability (F10), remote voltage (F11), remote breaker state (F12), and communication quality (F13). The following reference protocol defines causal feature availability; it is not a recovered hardware calibration or an experimentally validated feature pipeline. At the study sampling rate of 4 kHz and nominal frequency of 50 Hz, an 80-sample trailing window supplies voltage and current r.m.s. values every 1 ms. F2 is the least-squares slope of the latest 20 voltage-envelope values, with units p.u./s and window-end timestamps. Invalid samples invalidate the feature rather than being imputed.

For F3, a candidate event at t_0 starts acquisition of 120 envelope values at $t_0 + 20, \ldots, t_0 + 139$ ms, excluding mixed pre-/post-event r.m.s. windows. For each $\tau = 5, 6, \ldots, 500$ ms, constrained least squares fits $U_i = c + A\exp(-s_i/\tau)$ with $0 \le c \le 1$, $A \ge 0$ and $s_i = t_i - (t_0 + 20\ \text{ms})$. The minimum-residual fit is retained. Reference rejection conditions are an additional disconnection within the window, peak-to-peak variation below 0.01 p.u., residual RMSE above 0.02 p.u., or a time-constant optimum at a search endpoint. These declared implementation choices require deployment-specific assessment. F3 is unavailable before approximately 139 ms and cannot justify an earlier command; the switch/current branch need not wait for it.

For F6, an 80-sample rectangular-window fundamental DFT supplies phase phasors. With $a = \exp(\mathrm{j}2\pi/3)$, positive and negative sequences are $(V_a + aV_b + a^2V_c)/3$ and $(V_a + a^2V_b + aV_c)/3$. Their magnitude ratio is valid only when the positive sequence is at least the reference floor of 0.05 p.u. This nominal-frequency protocol requires frequency tracking if off-nominal operation introduces substantial leakage. Membership calibration remains separate from feature extraction.

The current-step feature compares like quantities rather than an instantaneous peak with an r.m.s. value:

$$I_{\mathrm{ref}}(k) = \max_{1 \le n \le N_w} I_{\mathrm{rms}}(k-n), \qquad F_5(k) = [I_{\mathrm{ref}}(k) - I_{\mathrm{rms}}(k)]_+. \tag{10}$$

Here, N_w spans the preceding 20-ms cycle and $[x]_+ = \max(0,x)$. Engineering normalisation is applied after extraction. The signed-power feature uses $F_7 = \mathrm{clip}(0.5 + 0.5\bar{P}/P_N, 0, 1)$, preserving reverse-flow depth without a discontinuous sign switch. Pseudo voltage loss requires low measured voltage, a closed incomer and absence of a substantial current step. Persistently low current alone cannot establish or refute a PT secondary-circuit fault.

3.2 First-match routing

Figure 3 separates the method into three connected modules. Offline calibration supplies fixed membership functions, evidence weights and remote-dependency indicators (panel a). Online identification applies this configuration to valid observations (panel b). Strategy coordination then converts candidate support into an action, block or review command subject to credibility and uncertainty gates (panel c). Calibration is not repeated as an iterative optimiser in the online execution cycle.

Within the online module, layer 1 identifies an inadvertent incomer trip from confirmed breaker opening and absent current; F8 records transition alignment. Layer 2 evaluates current-loss causes using local voltage and valid post-event remote topology evidence. Layer 3 tests voltage-channel inconsistency when measured voltage falls without a current step. Layer 4 distinguishes reverse power, light load and normal operation if earlier layers remain undecided.

Once a layer supplies a decisive candidate, later layers are not allowed to overwrite its physical interpretation. Credibility assessment and final arbitration still apply to that candidate. Communication degradation precedes these layers as an observation state, rather than competing with electrical events as another exclusive class. It can therefore coexist with genuine source loss.

3.3 Calibration, exclusivity and branch support

The routing layers combine discrete event evidence with continuously varying measurements. Membership functions place these measurements on a common support scale, while scenario-specific weights favour features that discriminate the target condition. Calibration is therefore an offline requirement, distinct from the numerical analysis of admission and arbitration performed here. For scenario m and feature j, the chosen trapezoidal membership $\mu_{mj}(x)$ has core breakpoints $b = Q_{0.10}$ and $c = Q_{0.90}$. Its outer breakpoints are

$$a = \max\{0, b - \kappa(c-b)\}, \qquad d = \min\{1, c + \kappa(c-b)\}, \quad \kappa = 1.5. \tag{11}$$

The membership is unity on $[b,c]$, zero outside $[a,d]$, and linear on the transition intervals. Zero-width transitions use the corresponding limiting step, avoiding division by zero. Binary events retain their logical interpretation rather than requiring an artificial continuous transition.

The empirical exclusivity score compares support within the target scenario with the strongest competing population:

$$J_{mj} = \widehat{\Pr}\{\mu_{mj}(X_j) > 0.5 \mid Y = m\} - \max_{\ell \ne m} \widehat{\Pr}\{\mu_{mj}(X_j) > 0.5 \mid Y = \ell\}. \tag{12}$$

When positive exclusivity exists, feature weights follow

$$w_{mj} = \frac{[J_{mj}]_+^{\gamma}}{\sum_{k=1}^{N} [J_{mk}]_+^{\gamma}}, \qquad \gamma = 2. \tag{13}$$

Otherwise, separation between target and competing scenario means provides the fallback:

$$D_{mj} = \min_{\ell \ne m} \frac{|\bar{x}_{mj} - \bar{x}_{\ell j}|}{\sqrt{(s_{mj}^2 + s_{\ell j}^2)/2}}, \qquad w_{mj} = \frac{D_{mj}^2}{\sum_k D_{mk}^2}. \tag{14}$$

Here, s^2 denotes calibration-sample variance. The expression requires finite pooled variances and a positive normalising denominator. If neither positive exclusivity nor finite separation establishes usable support, the result is review with a calibration diagnostic. Deterministic binary events retain explicit logical weights rather than infinite numerical weights. Fully collapsed

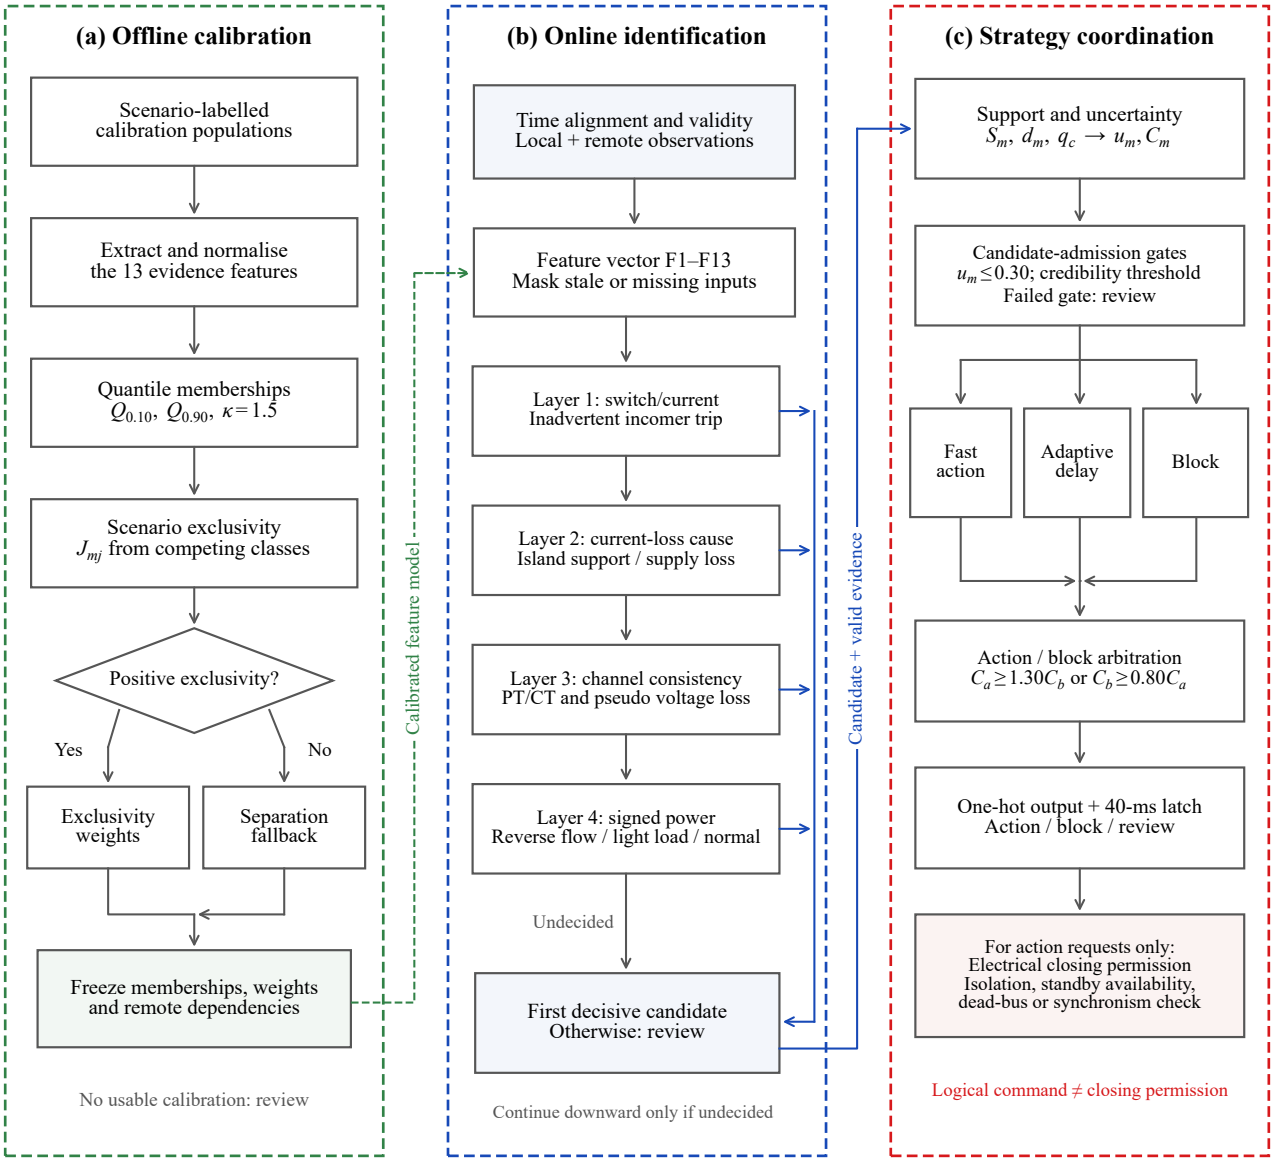


Figure 3. Architecture of hierarchical evidence identification and coordinated transfer decisions. (a) Offline calibration of memberships and evidence weights, with a separation-based fallback when no positive exclusivity is available. (b) Validity-masked observations enter four first-match layers; each decisive layer has a separate exit to candidate assessment. PT and CT denote potential and current transformers, respectively; the channel checks include PT secondary-circuit faults. (c) Support and uncertainty determine candidate admission, strategy selection and action/block arbitration. The output latch enforces a single command. Electrical closing permissions apply only to action requests and cannot be bypassed by high credibility. The dashed arrow transfers calibrated configuration; solid arrows carry online evidence or decisions.

membership intervals require an explicit point-membership or event interpretation. The analytical evaluations condition on normalised weights and candidate supports. They do not estimate a calibrated classifier's accuracy.

Let r_j indicate remote dependence and $Z_m = \sum_j w_{mj} q_c^{r_j} v_j$. For $Z_m > 0$, the adjusted weights and support are

$$\widetilde{w}_{mj} = \frac{w_{mj} q_c^{r_j} v_j}{Z_m}, \qquad S_m = \sum_{j=1}^{N} \widetilde{w}_{mj} \mu_{mj}(x_j). \tag{15}$$

For local features, the factor $q_c^{r_j}$ is defined as one, including at $q_c = 0$. When $Z_m = 0$, the branch lacks usable evidence and returns review. Normalisation is within a scenario, not across scenario supports, so a decisive event is not diluted by the number of competing labels.

3.4 Credibility and communication-dependent uncertainty

The remote-dependency fraction is $d_m = \sum_j w_{mj} r_j$. With $S_{-m} = \max_{\ell \neq m} S_\ell$, the conflict measure adopted for the utility project under study is

$$K_m = \mathrm{clip}\left(1 - \frac{S_m - S_{-m}}{S_m + S_{-m}}, 0, 1\right). \tag{16}$$

Only comparable candidates admitted by the current route enter this calculation. If both supports vanish, the result is review rather than an evaluated ratio. For an electrical candidate, uncertainty and credibility are

$$u_m = \mathrm{clip}\{0.6 d_m K_m (1 - q_c) + 0.5(1 - S_m), 0, 0.9\}, \qquad C_m = (1 - u_m) S_m. \tag{17}$$

The first term couples communication loss to conflict and remote dependence. The second penalises insufficient support. These coefficients define a reference decision rule, rather than fitted probabilities. Their consequences are analysed below and audited on a separate development set.

A candidate is accepted when $u_m \leq 0.30$ and $C_m \geq 0.60 + 0.10 d_m D$, where D indicates degraded operation. A purely local branch has $d_m = 0$. While its local inputs remain valid, its support, uncertainty and credibility gate are invariant to remote-link quality. This property does not imply invariance of every remote-assisted strategy or of physical breaker execution.

3.5 Delay adaptation and fail-safe arbitration

Confirmed candidates request fast action, delayed action or block. Delayed strategies interpolate between fast and conservative settings using credibility κ_m and normalised IIDG variability ν_m:

$$T_m = \mathrm{clip}\left\{T_{\min} + (T_{\max} - T_{\min}) \frac{(1 - \kappa_m)^2 + 0.5 \nu_m}{1.5},\ T_{\min}, T_{\max}\right\}, \tag{18}$$

with $T_{\min} = 40$ ms and $T_{\max} = 600$ ms. The analysed delay rule uses $\kappa_m = C_m$ normally and $\kappa_m = 0.7 C_m$ in degraded operation. This delay-specific adjustment is distinct from the invariance of locally supported classification. The limits are permitted bounds, not claims that every operating point reaches both endpoints. These bounds describe the strategy delay, excluding breaker mechanics and electrical closing permissions.

Let C_a and C_b be the strongest eligible action and blocking credibilities. Taking maxima keeps both quantities on the same scale regardless of the number of available strategies. After the relevant uncertainty gate, the output is

$$o = \begin{cases} \text{action}, & C_a \geq 0.60 \text{ and } C_a \geq 1.30 C_b, \\ \text{block}, & C_b \geq 0.60 \text{ and } C_b \geq 0.80 C_a, \\ \text{review}, & \text{otherwise.} \end{cases} \tag{19}$$

Degraded remote-dependent candidates must also pass their tightened gate. The 0.80 factor multiplies action credibility in the blocking rule. Reversing that relationship changes the decision regions and is not an equivalent rearrangement.

For positive credibilities, simultaneous action and block would require $C_a \geq 1.30 \times 0.80 C_a = 1.04 C_a$, which is impossible. The absolute gates exclude the origin. Increasing C_a at fixed C_b cannot leave the action region, while increasing C_b cannot enter it. One-hot encoding and a 40-ms latch provide a single stable logical output; electrical closing interlocks remain separate.

3.6 Numerical implementation and AI-assisted preparation

The evaluations were implemented in Python 3.12 with NumPy 1.26.4 and Matplotlib 3.9.2. The voltage grid and event-interpolation rule are specified in the voltage-model methods; the admission and arbitration analyses below specify their complete grids and inclusive boundary conventions. Closed-form event times and forward-Euler integration use separate calculation paths. Source data retain non-crossing model cases and all nominal arbitration outputs.

An AI-assisted workflow supported English-language editing, LaTeX preparation, figure code, numerical implementation and consistency checks. The executable models, parameter manifests and resulting numerical records are supplied for inspection and reproduction. The authors retain responsibility for the scientific interpretation and submitted content.

3.7 Independent network and sampled-signal benchmark

The benchmark evaluates the complete observation-to-request pipeline using a separately implemented positive-sequence network. It neither calls the analytical envelope nor supplies scenario labels to the controller. A 10-kV voltage base and 7-MVA power base define per-unit quantities. The working source connects through $z = R + \mathrm{j}X$ to an aggregate IIDG and a voltage-dependent load. The standby source is represented by the closing-permission boundary, not by an executed closing event.

Let $g \in \{0,1\}$ denote working-source connectivity, E its voltage, V the local bus phasor, and I_d the converter current. Kirchhoff's current law is

$$g\frac{E-V}{z} + I_d = (P_0 - \mathrm{j}Q_0)|V|^{\alpha-1}\exp(\mathrm{j}\arg V). \tag{20}$$

Active and reactive demand share the exponent $\alpha > 1$ in this benchmark. On an island, the solution is

$$|V| = \left(\frac{|I_d|}{\sqrt{P_0^2+Q_0^2}}\right)^{1/(\alpha-1)}, \qquad \arg V = \arg I_d + \tan^{-1}(Q_0/P_0). \tag{21}$$

Connected operation uses a damped fixed-point load-flow solution, with the current-balance residual checked at every step. This quasi-static network retains feeder impedance and reactive demand, but excludes electromagnetic and load-frequency transients.

An averaged converter response supplies the dynamic state,

$$\tau_i \dot{I}_d = I_d^\star - I_d, \qquad I_d^\star = \mathrm{sat}_{I_{\max}}\left[\frac{P_d - \mathrm{j}Q_d}{\max(|V|,0.10)}e^{\mathrm{j}\theta}\right]. \tag{22}$$

The radial saturation preserves current angle while enforcing its magnitude limit. The phase tracker uses

$$\dot{\theta} = \mathrm{clip}\{80\arg(Ve^{-\mathrm{j}\theta}), -10\pi, 10\pi\}. \tag{23}$$

This bounded first-order tracker is not a detailed hardware PLL. The current state advances with an exponential first-order update at 0.25 ms, using the preceding target. A continuous 0.360-s interval below 0.88 p.u. sets the converter target to zero. No universal anti-islanding response is assumed when this timer does not operate. The network may consequently produce an initial voltage step and subsequent recovery, unlike the monotone analytical envelope.

Each trajectory varies $R \in [0.015, 0.050]$, $X \in [0.040, 0.120]$, $|E| \in [0.98, 1.02]$, and $\alpha \in [1.25, 2.40]$. Ordinary loads use $P_0 \in [0.65, 1.20]$ and $Q_0/P_0 \in [0.12, 0.40]$; light loads use $P_0 \in [0.05, 0.14]$. Converter response times span 20–100 ms. The current limit is a factor from 1.10–1.30 times $\max(P_d, 0.20)$. These are declared benchmark ranges, not recovered device settings or identified parameter distributions.

Nine scenarios comprise normal operation, PT loss, light load, incomer trip, supported island, source loss, reverse power, light-load PT loss, and CT loss. Their generation-to-load ranges are, respectively, 0.08–0.80, 0.05–0.65, 0–0.25, 0.25–1.05, 0.60–1.05, 0–0.15, 1.15–1.70, 0–0.25, and 0.05–0.65. Disturbances occur at 0.250 s within 0.900-s records. Incomer trip opens the local breaker; supported island opens the remote source breaker. Source loss opens an intervening supply path while remote-source voltage remains present. Normal, light-load and reverse-power trajectories are non-switching controls. The reference task requests action for incomer trip or source loss, and withholding of transfer for the other scenarios. Supported-island cases therefore represent a conservative request policy, not proof that every island should remain connected indefinitely.

Three-phase voltage and current samples are reconstructed from the fundamental phasors at 4 kHz. Independent Gaussian sample noise has a base standard deviation of 0.002 p.u. PT or CT faults scale all three affected channels by a factor from 0.01–0.10 after the event. The independent local IIDG power channel supplies F10; its additional meter error is not varied. Breaker-position reporting delays vary from 2–20 ms. All compared methods receive identical samples and reported breaker states for a given case.

3.8 Executable feature configuration and calibration split

The implementation uses the trailing windows specified above and updates decisions every millisecond. Voltage and current normalisers are 1.2 and 2.0 p.u. F2 maps as $\mathrm{clip}(0.5+\dot{U}/100,0,1)$, and F3 maps as $\mathrm{clip}(\widehat{\tau}/0.5,0,1)$, with $\widehat{\tau}$ in seconds. F5 uses a 1.0-p.u. normaliser; F6 uses a 0.30 negative-sequence ratio. F7 uses a 1.0-p.u. active-power base. F11 divides remote voltage by 1.2; F12 is the received breaker state. F13 controls information quality and receives zero class-discrimination weight.

A causal event record bridges short current transitions and the longer envelope fit. A current drop above 0.04 p.u. starts a peak hold lasting at least 300 ms and continuing while current remains below 0.15 p.u. Remote-assisted current-loss routing requires a held drop of at least 0.08 p.u. Confirmed local breaker opening and current below 0.15 p.u. instead form a deterministic local candidate with unit support and zero remote dependence. This branch does not require a large pre-event current or wait for F3. F8 records switch/current timing consistency, with unit value within 20 ms and a linear decrease to zero at 100 ms. F9 requires measured voltage below 0.25 p.u., a closed local breaker and no active current-loss record. Retained history prevents a completed current step from being misinterpreted as a later PT-channel failure.

F10 divides the absolute deviation from its preceding 100-ms mean by $\max(|\overline{P}_d|,0.05)$, then normalises by 0.5. F3 acquisition starts at the first observed current drop, substantial voltage drop, or breaker transition. Its earliest availability is 139 ms after that detected event, rather than after the unobserved physical switching time. After a fit attempt, a further 0.10-p.u. decrease over 20 ms can initiate another episode at least 160 ms after the preceding onset. The previous estimate is invalidated until the new fit passes its acceptance checks. Invalid F3 and F6 estimates remain masked. Current-loss routing with unavailable remote evidence returns review instead of entering the normal-operation branch. Suspected CT loss is tested when current disappears while local and remote voltage remain present.

Calibration, development and final evaluation contain 162, 108 and 324 complete physical trajectories, respectively. Each scenario contributes 18, 12 and 36 trajectories. Calibration uses six fixed snapshots at 60, 140, 220, 300, 550 and 640 ms after the event, giving 972 labelled feature vectors. The later snapshots include converter-disconnection decays when they occur. Invalid entries are excluded feature by feature. Quantile breakpoints and exclusivity weights are then frozen; the supplied calibration file contains every breakpoint, weight and valid-sample count. The variance denominator of the separation fallback has a numerical floor of 10^{-8} in squared normalised-feature units. Missing separation receives zero weight, and a zero normaliser gives review.

Physical-parameter generators use seeds 190100, 190200 and 190700 for the three splits. Separate per-trajectory streams generate sample noise and communication events. Calibration and development alternate noise multipliers one and two; final evaluation balances multipliers one, two, four and eight within each scenario. No trajectory or time window crosses a split. Earlier benchmark inspection motivated the implementation corrections above; the new final seed was evaluated only after these corrections were frozen. Development-only one-factor checks vary the current-step routing threshold, credibility gate, action ratio and communication coefficient. The nominal coefficients are retained as a reference configuration, not claimed to be optimal. The supplied revision protocol records this development history; it is not an externally preregistered experiment.

3.9 Communication perturbations, baselines and outcome definitions

Packets are sent every 10 ms and retain source timestamps. Nominal delays are uniformly distributed over 2–6 ms. Delayed/lossy links use 30–80 ms delays and an independent 20% packet-loss probability. Interrupted links stop delivering packets 20 ms before the disturbance. Silent-error links deliver fresh packets but invert the remote breaker state and corresponding voltage status after the disturbance. This last case tests an incorrect payload that a freshness check cannot detect.

The receiver holds the newest received source timestamp, rejects data older than 60 ms, and evaluates quality over a 200-ms window. Remote-assisted routing also requires a source timestamp no earlier than the detected current-loss episode. A transport-fresh packet can otherwise describe pre-event topology. The reference delay d_0 is 20 ms. A missing packet contributes to estimated loss only after its 60-ms deadline has elapsed. These conventions prevent future packet arrivals or losses from entering the current decision. An open remote breaker retains the supported-island withholding policy even after local voltage decays. All four communication conditions reuse each complete trajectory, producing 1,296 related evaluation variants, rather than 1,296 independent physical trials.

Six implementations share the same input records. The conventional baseline requires voltage below 0.25 p.u. and current below 0.15 p.u. continuously for 300 ms. Flat routing replaces the ordered candidate mask with simultaneous competition among all nine classes. The no-quality ablation accepts held remote values and disables quality discounting, stale-data rejection and post-event timestamp checks. The local-only ablation masks both remote electrical features. Uniform weighting assigns equal weight to twelve features and excludes F13. All other calculations remain unchanged. This comparison tests the supplied conventional logic and internal components, not superiority over every published transfer method.

Fast action requires 40 ms of continuous eligibility. Other action requests use Eq. (18); a common 40-ms latch applies to action and block outputs in all implementations. Logical latency is measured from the physical disturbance to the first issued action request. It includes sampling, feature availability, received-data delay and the eligibility timer, but excludes breaker

mechanics and closing permissions. A false action is any post-event action in a non-action scenario. A missed action is absence of a request within the 650-ms post-event observation window in an action-required scenario. This finite-window outcome does not distinguish a later response from permanent refusal. Review is a terminal review state without an earlier action. Latency statistics include successful required actions only, and must therefore be interpreted with missed-action counts. No confidence interval treats communication variants as independent trials.

4 Communication and arbitration analysis

The decision equations permit separate analysis of candidate admission and final command selection. The admission boundary identifies the support required under degraded communication. The arbitration analysis then determines whether eligible commands can conflict and how much credibility perturbation a selected command can tolerate. These results condition on fixed candidate identities and valid upstream observations.

4.1 Closed-form admission boundary

The communication rule admits a useful analytical reduction. Define the non-negative communication penalty as $B_m = 0.6 d_m K_m (1 - q_c)$, and let $c_D = 0.60 + 0.10 d_m D$. For an admissible candidate, the uncertainty limit is below the clipping ceiling. Thus, clipping is inactive at the admission boundary and

$$C_m = (0.5 - B_m) S_m + 0.5 S_m^2, \qquad u_m = B_m + 0.5(1 - S_m). \tag{24}$$

Solving the credibility and uncertainty inequalities gives the minimum required support:

$$S_{\min}(B_m, c_D) = \max\left\{ -(0.5 - B_m) + \sqrt{(0.5 - B_m)^2 + 2c_D},\ 0.4 + 2B_m \right\}. \tag{25}$$

Admission is feasible only if $S_{\min} \le 1$. This expression provides an inspectable link between communication dependence, evidence support and the conservative review response.

For a local branch, $B_m = 0$ and $c_D = 0.60$, giving $S_{\min} = 0.70416$. The credibility gate is then more restrictive than the uncertainty gate. At nominal operation, the two requirements become equal at $B_m = 0.22857$, with support $S_m = 0.85714$. Above that penalty, uncertainty determines the admission boundary. When $B_m > 0.30$, even unit support violates the uncertainty limit and the branch must remain unadmitted.

At fixed support $S_m = 0.90$ and conflict $K_m = 0.50$, local credibility remains 0.855 throughout the communication sweep. For complete remote dependence, credibility ranges from 0.585 to 0.855 as quality rises from zero to one. The first point also exceeds the uncertainty limit. This difference follows from the decision equations and does not estimate a sensor failure probability.

The analysis checked Eq. (25) against the original clipped rule over 183,618 parameter combinations. The design uses 101 support levels, three remote dependencies, three conflicts, 101 communication qualities and two degradation states. Away from exact floating-point boundary ties, both formulations return identical admission decisions. Both support and quality take equally spaced values on $[0, 1]$; dependence and conflict each take 0, 0.5 and 1, and $D \in \{0, 1\}$. Zero support is unadmitted. For positive support, the selected conflicts correspond to absent competing support, competing support $S_m/3$, and equally strong or stronger competition, respectively. Exact inclusive decisions use integer indices $s = 100 S_m$, $q = 100 q_c$, $d = 2 d_m$ and $k = 2 K_m$. Define $U_{\text{num}} = \min\{1800, 3dk(100 - q) + 10(100 - s)\}$. Admission requires $U_{\text{num}} \le 600$ and $(2000 - U_{\text{num}}) s \ge 10000(12 + dD)$. All 183,618 rows are retained in Supplementary Data 1. The root-form comparison excludes 1,145 cases within 10^{-12} of an original gate; it agrees on the other 182,473 cases. Exact decisions remain supplied for all boundary cases.

4.2 Command exclusivity and credibility thresholds

The nominal arbitration grid contains 501 levels on each credibility axis, giving 251,001 combinations. Direct evaluation returns 61,954 action points, 94,351 block points and 94,696 review points, with no simultaneous action and block. The axes are $C_a = i/500$ and $C_b = j/500$ for integers $i, j = 0, \ldots, 500$. Exact inclusive comparisons use $i \ge 300$, $10i \ge 13j$ for action, and $j \ge 300$, $5j \ge 4i$ for block. These counts describe the geometry of the specified rule. They are not event probabilities or device reliability estimates.

The narrower ratio-only margin, $C_a < 1.30 C_b$ and $C_b < 0.80 C_a$, contains 3,785 points. Low-credibility cases enlarge the full review region when the absolute gates are applied. This distinction matters when reporting conservative operation: counting only the ratio margin omits most review outputs in the nominal grid.

For general positive absolute gate c_0, action ratio k_a and block ratio k_b, the condition $k_a k_b > 1$ is sufficient for disjoint commands. Equality does not provide this guarantee because both inclusive inequalities can hold on their common boundary. Figure 4 compares this transition at fixed $k_b = 0.80$ and varies the absolute gate separately. Absolute-gate sensitivity uses

$c_0 = 0.1, 0.2, \ldots, 0.9$; action-ratio sensitivity uses $k_a \in \{1.05, 1.10, 1.20, 1.25, 1.30, 1.40, 1.50\}$. Both retain the same credibility grid. At $k_a = 1.25$, 26 points satisfy both inclusive commands, demonstrating why equality is insufficient. Raising the absolute gate can only remove admitted commands and expand review; it does not alter the ratio condition for exclusivity.

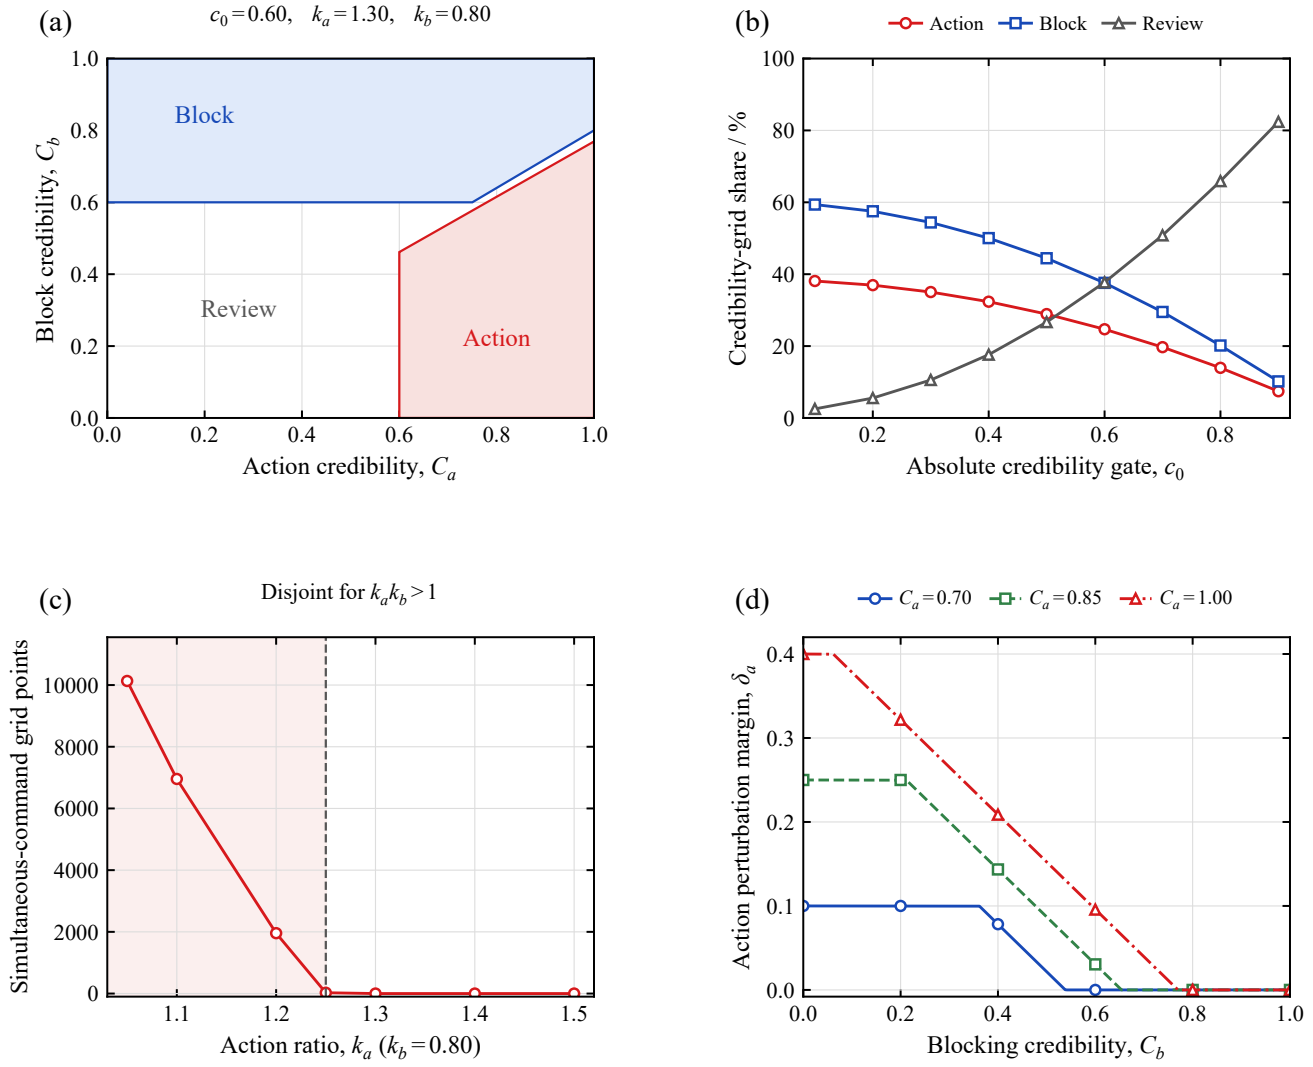


Figure 4. Arbitration geometry, parameter sensitivity and bounded perturbation margins. (a) Nominal action, block and review regions after candidate admission. (b) Fractions of a fixed 251,001-point credibility grid for nine absolute gates. (c) Simultaneous-command counts as k_a varies at $k_b = 0.80$ and $c_0 = 0.60$; the vertical line marks product one. (d) Symmetric credibility perturbation margins for three initial action credibilities; zero denotes no positive action certificate. Grid fractions are geometric descriptions, and no experimental probabilities are inferred.

4.3 Certificate for bounded credibility perturbations

Logical exclusivity alone does not establish that a selected command is stable to errors in its inputs. Consider an action point with perturbations $|\varepsilon_a|, |\varepsilon_b| \le \delta$. The least favourable changes reduce C_a and increase C_b. Sufficient and exact bounds for a symmetric box contained in the action half-planes are

$$\delta_a = \min\left\{C_a - c_0,\ \frac{C_a - k_a C_b}{1 + k_a}\right\}. \tag{26}$$

For $\delta_a > 0$, every perturbation smaller than δ_a preserves the action inequalities. The analogous blocking certificate is

$$\delta_b = \min\left\{C_b - c_0,\ \frac{C_b - k_b C_a}{1 + k_b}\right\}. \tag{27}$$

These certificates condition on the candidate set and uncertainty gates remaining unchanged. A loss of measurement validity may invalidate that premise and must be handled by the validity gate.

At $(C_a, C_b) = (0.85, 0.45)$, the action certificate is 0.11522 credibility units. The ratio inequality, rather than the absolute gate, limits this point. The numerical test evaluated the adverse-direction corner at 99% of the positive certificate for every positive-margin nominal action point. No violation occurred. The check covers 61,708 action points with margin greater than 10^{-10}. The continuous certificate follows because the adverse corner $(-\delta, +\delta)$ minimises both action inequalities. The absolute difference becomes $C_a - c_0 - \delta$, and the ratio difference becomes $C_a - k_a C_b - (1 + k_a)\delta$. Their simultaneous non-negativity gives Eq. (26); interchanging action and block gives Eq. (27). Thus the proof covers the continuous box, whereas the finite check verifies its implementation. Restricting perturbed credibilities to $[0, 1]$ only reduces the feasible box and preserves the sufficient bound.

5 Complete-pipeline benchmark results

5.1 Network consistency and representative signal trajectories

The independent network calculation retained all 594 physical trajectories. Its maximum current-balance residual was 5.48×10^{-15} p.u. in double precision. Halving the step from 0.250 to 0.125 ms across all 108 development trajectories gave a voltage RMSE of 3.08×10^{-6} p.u. The maximum difference was 1.00×10^{-4} p.u. These checks establish numerical consistency of this particular averaged model, not agreement with measured network behaviour.

Figure 5 displays the first pre-indexed evaluation trajectory for three scenarios under nominal communication. Current loss provides an earlier indication of incomer opening than the eventual voltage collapse. In the supported-island case, remote topology distinguishes a block request from source-loss action. Under light-load PT loss, measured voltage collapses while current remains small but does not disappear. Conventional voltage/current logic cannot resolve this distinction. The command strips retain all review transitions and the declared latch, rather than showing only a selected successful command.

5.2 Shared-input comparisons and component ablations

Table 2 reports all 7,776 method runs on the same 1,296 evaluation variants. The full method produced 5 false actions among 1,008 non-action variants, whereas conventional logic produced 144. However, the full method left 101 of 288 required requests unissued, compared with 180 for conventional logic. Its terminal review count was 260. These outcomes separate protection against an erroneous request from the ability to issue a required request within the observation horizon.

Median successful-request latency was 57 ms for the full method and 319 ms for the conventional baseline. These medians describe different successful subsets, not a paired estimate of restoration-time reduction. Among 46 variants with successful required requests under both methods, the median paired baseline-minus-full latency was 253 ms. Figure 6d displays the conditional distributions and their denominators. Missed requests remain in the outcome table and are not assigned an artificial latency equal to the record length. The 650-ms window censors later responses, particularly when voltage confirmation and the conventional 300-ms timer operate sequentially.

The ablations did not establish uniform dominance of the full method. Flat routing produced four false actions but missed 213 required requests. Uniform weights produced twelve false actions and missed 101 requests. Disabling quality controls reduced missed requests to 36, but produced ten false actions. The local-only ablation produced no false actions and missed 144 required requests. Thus, lower false-action counts did not imply better availability, and the full method was not best on every metric.

Table 2. Complete-pipeline evaluation on 324 physical trajectories under four communication conditions. The 1,296 variants contain 288 action-required and 1,008 non-action cases. Latency includes successful required actions only.

Method	False action	Missed action	Review	Median (ms)	P95 (ms)
Full method	5	101	260	57	541.3
Voltage/current	144	180	208	319	319.0
Flat routing	4	213	95	403	547.8
No quality control	10	36	147	58	560.0
Local only	0	144	480	56	59.0
Uniform weights	12	101	216	57	181.9

5.3 Communication limits and parameter audit

The full method missed 0 of 72 required requests under nominal communication. This count increased to 29 for delayed/lossy links, 36 for interrupted links and 36 for silent payload errors. All five false actions occurred with silent payload errors.

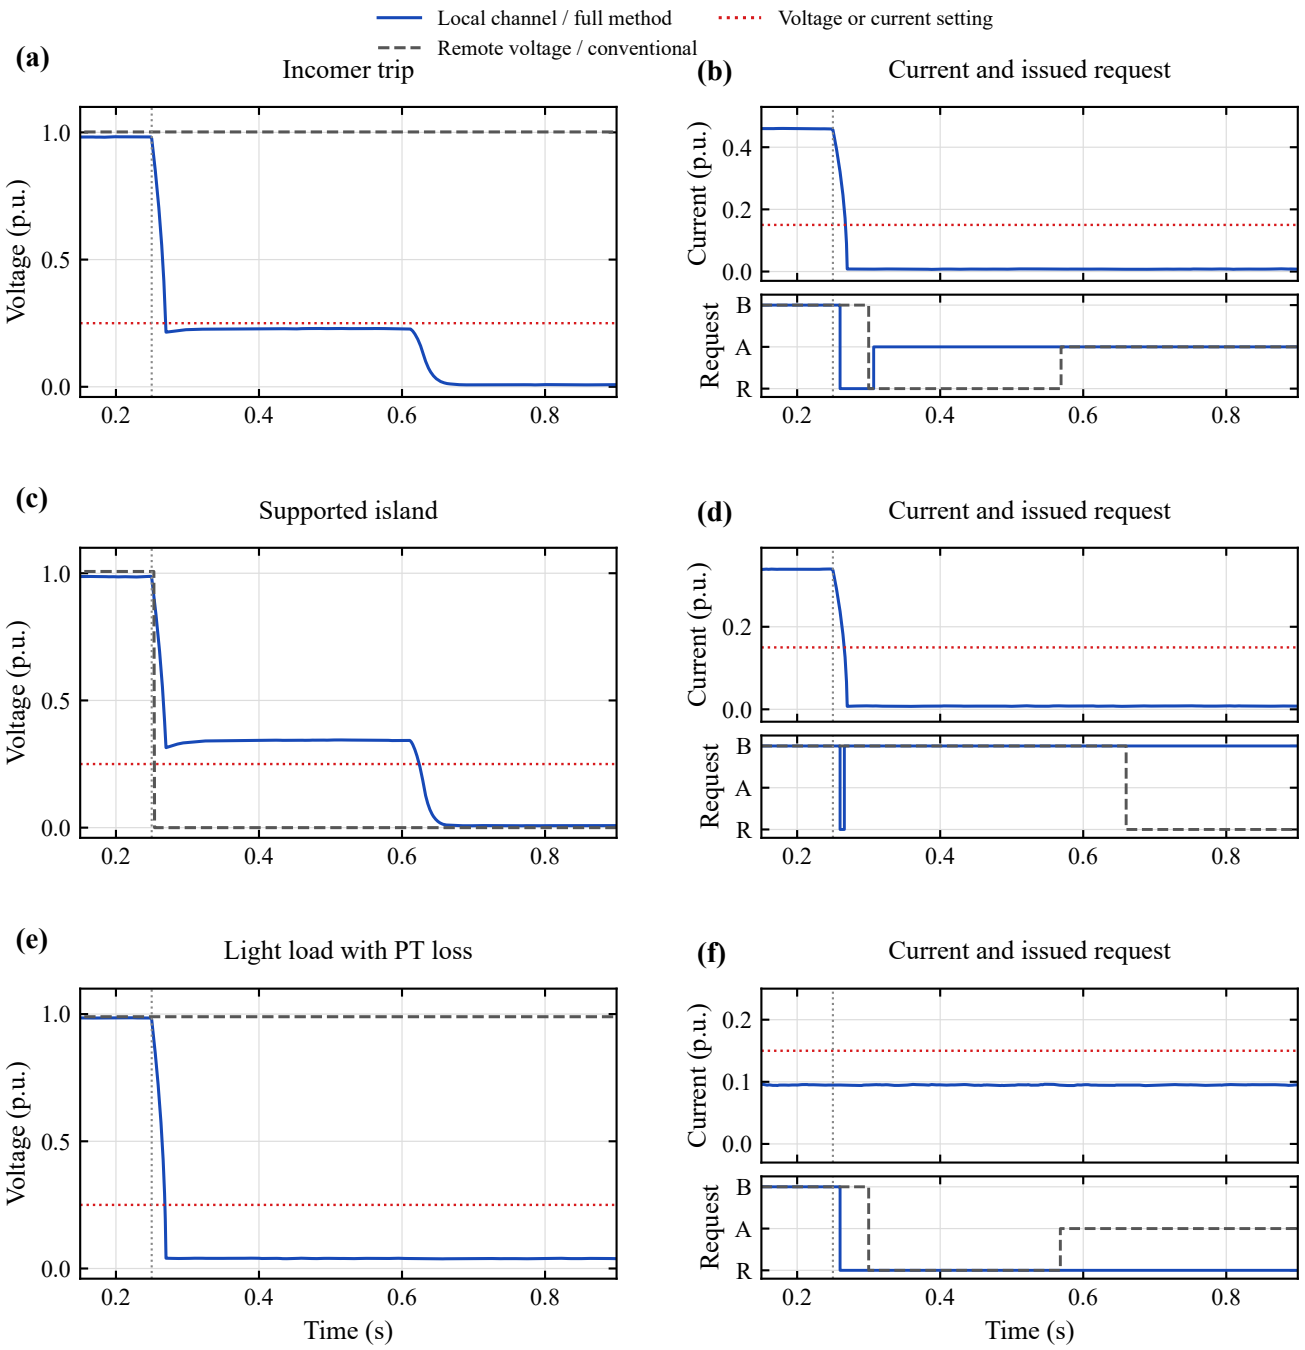


Figure 5. Sampled observations and issued logical requests in the independent network benchmark. Rows show incomer trip, supported island and light-load PT loss. Left panels show trailing local voltage r.m.s. and received remote voltage. Right panels show trailing incomer-current r.m.s. above full-method and conventional request traces. R, A and B denote review, action and block. Dotted vertical lines mark the 0.250-s physical disturbance; red dotted lines are the voltage or current setting. Each row is the first pre-indexed final-evaluation trajectory of its scenario, not a best-performing selection. Waveforms include the declared measurement noise. No physical closing or load restoration is executed.

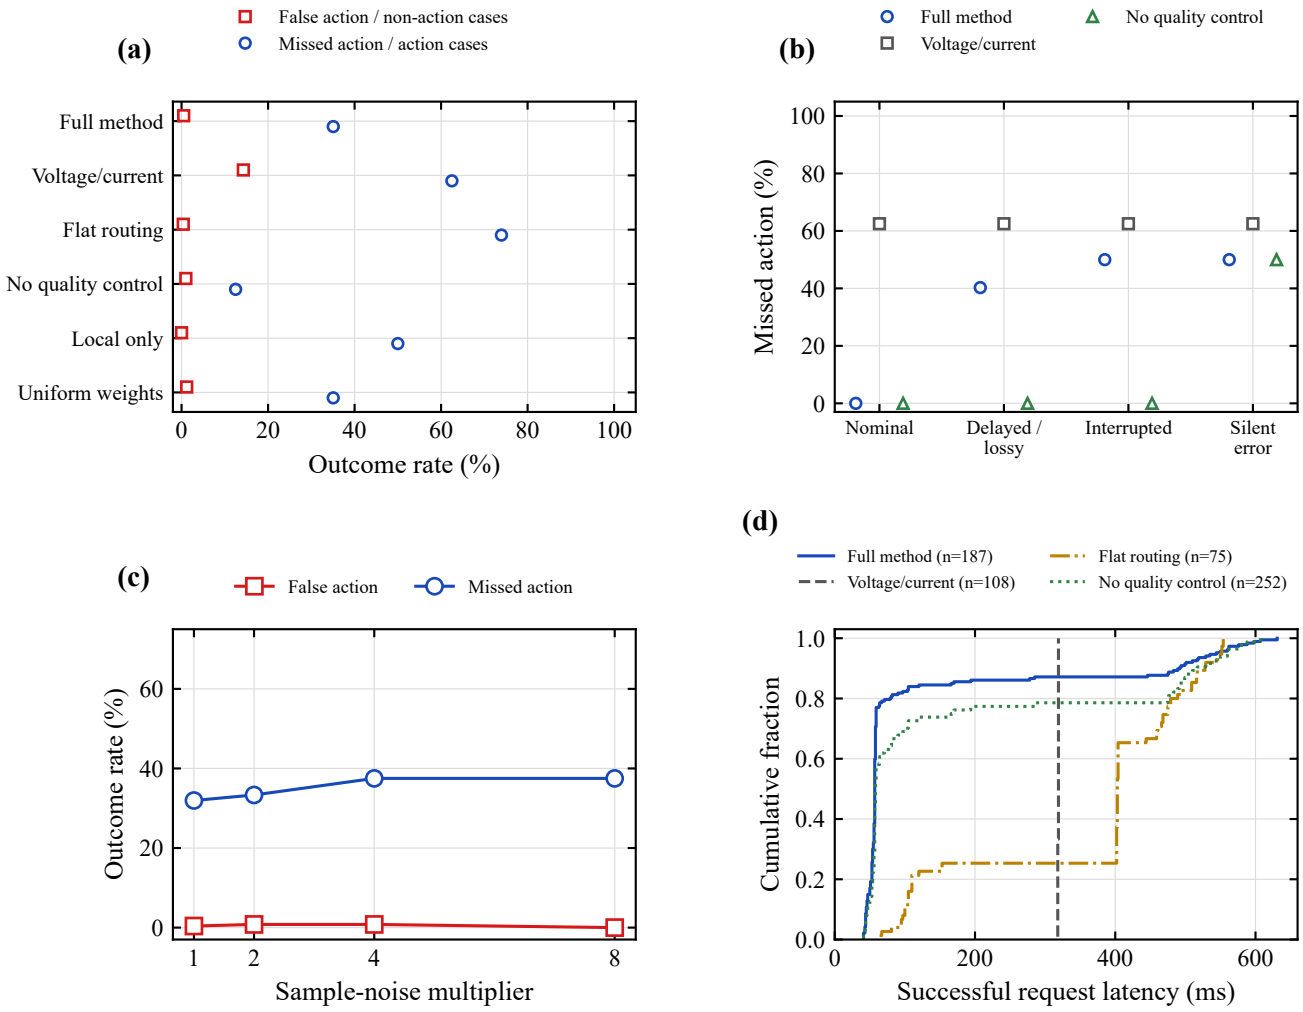


Figure 6. Complete-pipeline performance and its boundaries. (a) False-action rates use 1,008 non-action variants as denominator; missed-action rates use 288 action-required variants. (b) Missed-action rates by communication condition, with 72 required-action cases per condition. (c) Full-method outcomes by sample-noise multiplier, each containing 81 distinct physical trajectories and 324 communication variants; action and non-action denominators are 72 and 252. Noise strata contain different balanced trajectories and are not a paired noise sweep. (d) Empirical cumulative latency distributions among successful required requests only, with successful counts in the legend. All results are from the declared computational benchmark. Communication variants reuse physical trajectories, so these counts do not represent independent field trials or confidence intervals.

Freshness and packet-loss checks cannot certify payload correctness. A branch requiring remote evidence can also remain unavailable after packet expiry, even when its local observations are physically plausible. The local incomer-trip branch and remote-assisted source-loss branch therefore require separate assessment.

Across the 432 development variants, the one-factor parameter audit retained zero false actions and 34 missed requests. Terminal reviews ranged from 66 to 77, compared with 75 at the nominal configuration. Current-step and credibility settings affected review counts, while the tested action ratios and communication-conflict coefficients did not change them. This plateau does not identify an optimum. In particular, the communication-conflict coefficient has no effect when a uniquely routed candidate has no competing support and hence $K_m = 0$. Wider conflicting-evidence cases are needed to assess that term. Development and final-evaluation outcomes differed, reinforcing the need to retain the separate evaluation split.

6 Discussion

The contribution is a linked analysis of observability, candidate admission and command stability, rather than a new voltage NDZ or a claim that elementary inequalities are novel in isolation. Established power-balance theory identifies when an undervoltage observation is unavailable. First-match routing then gives breaker and current-transition evidence a defined place before graded credibility assessment. The admission and arbitration boundaries make explicit what additional conditions are needed before an observation can support a stable logical command.

The admission boundary explains how communication degradation can produce different responses across branches. A purely local branch retains its gate when valid local support is unchanged. A remote-dependent branch may require progressively stronger support, eventually reaching an infeasible requirement. This mechanism provides a quantitative basis for conservative review without treating an unavailable remote measurement as a zero-voltage observation.

The arbitration analysis distinguishes command exclusivity from command stability. The product condition controls whether action and block can coexist. The perturbation certificates instead quantify distance from the selected command's boundaries. A command can be unique yet have a very small margin. Reporting both properties therefore conveys more information than showing only a single-output truth table.

Improving undervoltage sensitivity by raising U_{set} moves the observability boundary but does not remove the supported-island blind region for settings below nominal voltage. Shortening the confirmation delay likewise cannot create a threshold crossing that the equilibrium prevents. Conversely, an excessively high pickup may admit non-islanding voltage depressions. The practical implication is to coordinate pickup and delay with blocking conditions and independent, valid remote evidence, rather than tune the no-voltage element alone. This interpretation complements setting-map and undervoltage-block coordination studies[21,22]; it is not a universal setting recommendation. A failed remote link must trigger the specified local fallback or review, not a fictitious remote dead-line indication.

The utility example supplies operating context, but the analytical boundaries are expressed in dimensionless variables. Equation (5) can be recomputed for any $0 < U_{\mathrm{set}} < 1$ and $\alpha > 0$ within the assumed load model; the support threshold and arbitration margins can likewise be recomputed for other admissible decision parameters. No feeder impedance appears in these expressions because network dynamics have been abstracted into the stated envelope and evidence inputs, not because their physical influence has been shown to vanish. Transferability therefore follows from the parameterised assumptions and must be checked when those assumptions change. Open code and complete grids make that check reproducible, but do not independently prove physical generality.

The independent network benchmark connects sampled waveforms to calibration, feature extraction, routing and command timers. Its results do not support an unconditional reliability claim. The full method reduced false actions relative to conventional logic, but still produced 5 false actions and missed 101 required requests. Disabling quality controls reduced missed requests to 36, while increasing false actions from 5 to 10. This comparison exposes a security–availability trade-off rather than uniform dominance of the complete design. Choosing an operating point requires an explicit cost for erroneous transfer, withheld requests and delay, followed by assessment on the intended network.

The calibration results also identify a mechanism that deserves caution. Small empirical exclusivity differences can concentrate weights on weakly discriminating channels. The revised event logic allowed F3 estimation during later converter-disconnection decays, rather than retaining only the initial step-and-recovery window. Some calibration entries remained invalid because the required decay was absent or the fit failed its stated acceptance conditions. Estimator availability does not establish discriminating value, and a thirteen-feature implementation does not demonstrate thirteen independent performance gains. Application-specific calibration should examine feature stability across trajectories and operating domains before deployment.

The averaged network includes feeder impedance, reactive demand, converter current saturation and phase tracking. It still omits switching harmonics, detailed PLL and anti-islanding controls, motor dynamics, unbalanced primary faults and actual closing operations. It is therefore a computational stress test of transfer-request logic, not hardware or field validation. The analytical envelope is not fitted to these network trajectories. Its monotone transient and regulated voltage cap remain assumptions whose accuracy requires separate waveform identification.

The perturbation certificate applies in credibility coordinates and does not directly bound raw sensor noise. Validity changes and candidate switching can leave the certificate's fixed-candidate domain. Source-current transition memory, feature waiting times and packet expiry must consequently be included when applying the bounds to a device. Electrical isolation, synchronism permissions and a validated restoration sequence remain separate requirements.

7 Conclusion

IIDG support can delay undervoltage confirmation after source separation, while measurement faults can produce an apparently similar voltage-loss indication. This study addresses the resulting transfer-decision problem by linking residual-voltage observability, hierarchical evidence admission and command stability. The framework orders thirteen features into four first-match layers, preserves invalid-data semantics, and separates logical requests from electrical closing permissions.

The voltage analysis specialises the established NDZ relationship to the transfer no-voltage setting. At $U_{\mathrm{set}} = 0.25$ and $\alpha = 1.8$, stage-I observability is lost at a generation-to-load ratio of 0.08247, corresponding to 0.577 MW in the illustrative network. Independent integration over 4,941 combinations reproduces all 3,240 finite crossings and 1,701 non-crossings. Step refinement gives first-order convergence and a crossing-time RMSE of 0.109 ms at 0.125 ms. These findings show why reducing the confirmation delay cannot resolve an equilibrium-supported blind region.

The candidate-admission analysis gives a local minimum support of 0.70416 and identifies a communication penalty above 0.30 as infeasible under the stated uncertainty gate. The 183,618-combination evaluation checks these boundaries with explicit equality conventions. The condition $k_a k_b > 1$ is sufficient for disjoint commands. The nominal arbitration grid contains 94,696 review points and no simultaneous action and block. A separate perturbation bound gives an action margin of 0.11522 at $(C_a, C_b) = (0.85, 0.45)$.

Complete-pipeline evaluation used 324 physical trajectories, four communication conditions and six shared-input implementations. The full method produced 5 false actions in 1,008 non-action variants, but left 101 of 288 required requests unissued. Median latency among its successful required requests was 57 ms. These outcomes show why command exclusivity, availability and speed must be assessed separately. The framework provides reproducible conditions and an executable benchmark for that assessment; it does not establish deployment readiness or measured device reliability.

Data availability

All numerical data generated and analysed in this study are included in Supplementary Data 1 and 2. Supplementary Data 1 contains the 4,941 voltage-model cases, 183,618-point admission grid and 251,001-point arbitration grid. Its network-benchmark folder contains parameter manifests, all 594 physical trajectories, frozen calibration parameters, final-evaluation feature frames and all 7,776 method outcomes. Supplementary Data 2 contains the sampled three-phase voltage/current records for these same trajectories. Development sensitivities, representative traces and verification summaries are also supplied in Supplementary Data 1. The READMEs define variables, units, split assignments and figure-to-file mappings. These records are generated computational data, not field measurements.

Code availability

The complete Python code for model integration, feature extraction, calibration, routing, communication, arbitration and figure generation is supplied with Supplementary Data 1. The package includes software requirements, reproducibility tests and commands for regenerating every result. The implementation uses Python 3.12, NumPy 1.26.4, SciPy 1.13.1 and Matplotlib 3.9.2. Analytical grids are deterministic; the network benchmark records fixed split-specific and per-trajectory random seeds. Times New Roman fonts are required for the supplied plotting style and are not redistributed.

References

1. Ackermann, T., Andersson, G. & Söder, L. Distributed generation: a definition. *Electr. Power Syst. Res.* **57**, 195–204 (2001).
2. Razavi, S.-E. *et al.* Impact of distributed generation on protection and voltage regulation of distribution systems: a review. *Renew. Sustain. Energy Rev.* **105**, 157–167 (2019).
3. Manditereza, P. T. & Bansal, R. Renewable distributed generation: the hidden challenges—a review from the protection perspective. *Renew. Sustain. Energy Rev.* **58**, 1457–1465 (2016).
4. Gadanayak, D. A. Protection algorithms of microgrids with inverter interfaced distributed generation units—a review. *Electr. Power Syst. Res.* **192**, 106986 (2021).

5. IEEE. IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces. IEEE Std 1547-2018 (2018).
6. Zheng, T., Ma, Y., You, Y. & Hu, S. A fast standby automatic switching scheme for photovoltaic power station considering energy storage cooperative control. *Trans. China Electrotech. Soc.* **41**, 961–974 (2026). In Chinese.
7. Bai, X., Mavrocostanti, Y., Strickland, D. & Harrap, C. Distribution network reconfiguration validation with uncertain loads—network configuration determination and application. *IET Gener. Transm. Distrib.* **10**, 2852–2860 (2016).
8. Barra, P. H. A., Coury, D. V. & Fernandes, R. A. S. A survey on adaptive protection of microgrids and distribution systems with distributed generators. *Renew. Sustain. Energy Rev.* **118**, 109524 (2020).
9. Hooshyar, A. & Iravani, R. Microgrid protection. *Proc. IEEE* **105**, 1332–1353 (2017).
10. Joshi, J. *et al.* Control strategy for current limitation and maximum capacity utilization of grid connected PV inverter under unbalanced grid conditions. *Sci. Rep.* **14**, 10118 (2024).
11. Barkhi, M., Pourhossein, J. & Hosseini, S. A. Integrating fault detection and classification in microgrids using supervised machine learning considering fault resistance uncertainty. *Sci. Rep.* **14**, 28466 (2024).
12. Khamis, A., Shareef, H., Bizkevelci, E. & Khatib, T. A review of islanding detection techniques for renewable distributed generation systems. *Renew. Sustain. Energy Rev.* **28**, 483–493 (2013).
13. Kulkarni, M. S. *et al.* Enhancing grid resiliency in distributed energy systems through a comprehensive review and comparative analysis of islanding detection methods. *Sci. Rep.* **14**, 12124 (2024).
14. Bakhshi-Jafarabadi, R., Sadeh, J. & Popov, M. Maximum power point tracking injection method for islanding detection of grid-connected photovoltaic systems in microgrid. *IEEE Trans. Power Deliv.* **36**, 168–179 (2021).
15. Bakhshi-Jafarabadi, R., Sadeh, J., Chavez, J. d. J. & Popov, M. Two-level islanding detection method for grid-connected photovoltaic system-based microgrid with small non-detection zone. *IEEE Trans. Smart Grid* **12**, 1063–1072 (2021).
16. Kim, H. J., Kim, M. S. & Lee, E. S. Enhanced islanding detection for GFM and GFL inverters using negative sequence current injection. *IEEE Trans. Ind. Electron.* **72**, 10119–10129 (2025).
17. Abo-Khalil, A. G. *et al.* Reinforcement learning driven adaptive active frequency drift for fast and reliable islanding detection. *Sci. Rep.* **16**, 876 (2026).
18. Pal, K. *et al.* Intelligent islanding detection framework for smart grids using wavelet scalograms and HOG feature fusion. *Sci. Rep.* **15**, 30351 (2025).
19. Chauhdary, S. T. *et al.* Microgrid anti islanding protection scheme based on deep neural network algorithm and unscented Kalman filtering. *Sci. Rep.* **15**, 28726 (2025).
20. Zeineldin, H. H. & Kirtley, J. L., Jr. Performance of the OVP/UVP and OFP/UFP method with voltage and frequency dependent loads. *IEEE Trans. Power Deliv.* **24**, 772–778 (2009).
21. Motter, D. & de Melo Vieira, J. C. The setting map methodology for adjusting the DG anti-islanding protection considering multiple events. *IEEE Trans. Power Deliv.* **33**, 2755–2764 (2018).
22. Motter, D. & Vieira, J. C. M. Improving the islanding detection performance of passive protection by using the undervoltage block function. *Electr. Power Syst. Res.* **184**, 106293 (2020).
23. Tang, Y. *et al.* A new basic probability assignment generation and combination method for conflict data fusion in the evidence theory. *Sci. Rep.* **13**, 8443 (2023).
24. Zhang, S., Gao, Y., Kou, Y. & Guo, Q. Temporal evidence fusion evaluation method considering time sequence variation trend. *Sci. Rep.* **15**, 25129 (2025).
25. Huang, T., Wu, D. & Ilić, M. Cyber-resilient automatic generation control for systems of AC microgrids. *IEEE Trans. Smart Grid* **15**, 886–898 (2024).
26. Hendrickx, K., Perini, L., Van der Plas, D., Meert, W. & Davis, J. Machine learning with a reject option: a survey. *Mach. Learn.* **113**, 3073–3110 (2024).
27. Zhu, X., Wang, D., Pedrycz, W. & Li, Z. Fuzzy rule-based local surrogate models for black-box model explanation. *IEEE Trans. Fuzzy Syst.* **31**, 2056–2064 (2023).
28. El-Fergany, A. A. Reviews, challenges, and insights on computational methods for network reconfigurations in smart electricity distribution networks. *Arch. Comput. Methods Eng.* **31**, 1233–1253 (2024).
29. Sharma, J. & Sidhu, T. S. A low voltage microgrid protection scheme using digital instrument transformers. *IET Renew. Power Gener.* **18**, 1007–1018 (2024).

Author contributions

Rui Zhou: Methodology, Writing—original draft. Jian Ji: Methodology, Writing—review and editing. Guoming Yan: Data curation, Writing—review and editing. Jin Zhu: Methodology, Data curation. Peng Li: Data curation, Writing—review and editing. Zhengxiang Ding: Methodology, Writing—review and editing. All authors approved the manuscript for submission.

Funding

This work was supported by the Science and Technology Project of State Grid Jiangsu Electric Power Co., Ltd. (grant no. J2026058), concerning reliability improvement of automatic-transfer devices under high-penetration distributed generation.

Competing interests

The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryData1.zip
- SupplementaryData2.zip