ARCHIVES OF ELECTRICAL ENGINEERING VOL. X(Y), pp. 1-22 (2027)

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

A first-hitting-time framework for risk-based adaptive tripping and prognostics in gas-insulated switchgear

RUSLAN MAMMADOV, LALA BEKIROWA, ELVIN YUSUBOV*

Azerbaijan State Oil and Industry University, Department of Instrumentation Engineering
16/21 Azadliq Ave, Baku, Azerbaijan
*e-mail: mammrl1984@yahoo.com, lala.bakirova@asoiu.edu.az, *elvinyusifov05@gmail.com*

Early Access

Abstract: Gas-insulated switchgear (GIS) protects itself against dielectric breakdown by tripping when the SF_6 pressure falls below a fixed level. This logic cannot distinguish a slow leak from a rapid pressure loss, and a single threshold forces a compromise: set high, it de-energises the equipment unnecessarily during chronic leaks; set low, it cannot guarantee that the pressure remains above the critical level by the time the breaker completes its operation during a fast leak. This article proposes a trip rule in which the breaker operates once breach of critical level within the protection chain's reaction-time budget becomes near-certain under a conservative upper bound on the estimated leak rate. The method couples choked-orifice leak physics with an interacting multiple-model estimator of pressure and leak conductance, operating on temperature-compensated measurements; the probability of first hitting the critical level within the reaction budget is available in closed form through the inverse-Gaussian first-passage distribution. The proposed method is compared with four threshold-crossing baselines in a five-way ablation study across leak severities. The same running estimator additionally supplies a short-term breach-risk probability and a long-term remaining-useful-life estimate with an 80% prediction interval. Of the five methods, only the proposed one keeps the pressure above the critical level at breaker completion in every tested leak class while avoiding unnecessary tripping, and its remaining-time interval attains 80–100% empirical coverage of its nominal level. A first-hitting-time criterion thus provides uniformly safe operation and accurate prognostics from a single estimator, where threshold-crossing rules must trade one against the other.

Key words: adaptive protection, first-hitting-time, gas-insulated switchgear, interacting multiple model, prognostics, remaining useful life, SF_6 leak detection

1. Introduction

Gas-insulated switchgear (GIS) employs metal-encased grounding for high-voltage conductors and uses SF_6 gas for insulation purposes. The breakdown voltage of SF_6 gas is pressure dependent. For high voltage insulation systems, dielectric strength will depend on the state of the insulating medium [1]. Since pressure has a large impact on the dielectric strength of SF_6 gas, gas leaks will significantly lower the insulating strength of the gas, leading to the dangerous possibility of internal arcing, which can even lead to the rupture of the gas container. IEC 62271-203 [2] specifies requirements for devices that monitor SF_6 gas density, raising alarms

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

and lockout when the density drops, while NEMA US 80020-2022 [3] gives guidance on field leak detection. These devices will account for the pressure of the gas and the temperature will be considered as well. A constant threshold is inadequate to address events such as valve failure or seal blowout, rupture events that can traverse the entire safe operating range within seconds. Alternatively, consider the chronic leak from a gasket that is the result of slow degradation over time. If a constant trip level is set to a value that will allow a controlled response during normal operation of the system, it will be de-energised, resulting in frequent unnecessary de-energisations and their operational and financial costs. Conversely, if a constant trip level is set to a value that will minimise the shutdown frequency for chronic leak repairs, the gas system may not have adequate safety response to mitigate the impact of a fast leak.

This article describes and analyses the general control problem of a variety of trip decisions that are essentially ON/OFF decisions made within an operating range defined by $[x_{\text{low}}, x_{\text{high}}]$, governed by the monitored signal's rate of change. When within that range, and the signal is undergoing a rapid decline, the trip will occur at x_{high}. If the signal is undergoing a slow decline, the trip will occur at x_{low}. This provides a safety criterion that is explicit, quantifiable, and auditable, as opposed to the traditional safety margin that is empirically adjusted.

The main contributions are the following:

1. An adaptive trip decision based on a first-hitting-time (FHT) risk criterion is proposed. The estimated drift is first inflated to a $(1-\varepsilon)$ upper-confidence bound on the leak rate; the system then trips at the first instant at which breach of the floor within a predefined reaction-time budget T_{react} becomes near-certain (probability $\geq 1-\varepsilon$) under that conservative bound. The local trajectory is modelled as a drifted Wiener process, for which the hitting probability is available in closed form through the inverse-Gaussian distribution and is validated against direct stochastic simulation of the underlying process.
2. A physics-based, interacting multiple models (IMM) estimator is presented. This estimator is constructed using modes (healthy, orifice-leak, rupture-class) along the choked orifice leak equation, $\mathrm{d}P/\mathrm{d}t = -kP$, which, unlike ad hoc choices, leads to the leak conductance estimate and mode probabilities. These justify the trip decision and help classify faults and estimate defects in real time. This integrates protection and condition monitoring into one estimator (Section 4.3).
3. In Section 4.5, we conduct a five-way ablation study to identify the impact of each design component for the rupture-class, orifice-leak, and chronic leak regimes. We examine the impacts of physics-based versus arbitrary thresholding, nonlinear versus linear rate modelling, recursive versus batch estimation, single versus multi-model coverage, and risk versus ad hoc margins. We also examine a temperature-compensated stress test, a continuous-rate Monte Carlo validation (Section 6), and a sensitivity analysis of ε and T_{react}.
4. Utilising the (P, k) estimator in a rough hourly trend-monitoring path (further elaborated upon in Sections 4.6 and 6.5), we generate two complementary prognostic outcomes. The first is a short-term estimate of breach risk, termed P(breach x_{low}), resulting from the first-hitting-time Eq. (6), coupled with drift inflation (8), set at a seven-day and a thirty-day time frame. Both of these time frames are still considered short in relation to the time

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

constant of the leak. The second output is a long-term estimate, which is framed within an 80% uncertainty band. To approximate this, we apply the same delta method that is used in the trip criterion, to the exact exponential hitting-time function, as (6) and (8) are not applicable at this multi-month time frame. Consequently, this single estimator family offers the capability to generate a rapid protection decision, as well as make a near-term maintenance advisory and a long-term maintenance planning estimation.

2. Related work

Static and rate-aware protection schemes

Monitoring SF_6 gas leaks is predominantly considered temperature-corrected SF_6 monitoring with the setting of alarm thresholds, lockouts, and rate-of-change controls to set secondary trip point alarms [2, 3]. The reliability of fixed-threshold indicators would themselves be conditional, and recent evidence in the medium voltage range suggests that detection should be considered a probabilistic rather than a deterministic quantity [4]. The same would apply to the condition monitoring of equipment, which has also been considered a cost-based prioritisation problem [5]. An alternative SF_6 monitoring approach considers partial discharges and the arcing decomposition of SF_6 gas, to identify fault conditions [6]. Studies of arc modelling of the interrupting devices deal with the same medium from the switching duty [7]. This article focuses on trip decisions related to SF_6 gas leak (large-scale) and, therefore, does not analyse SF_6 gas decomposition to define fault conditions.

In Section 3, we further develop the repositioning threshold by the rate of change within a two-sided certified operating band and a defined probabilistic criterion. The repositioning-by-rate mechanism should not be confused with adaptive protection, which changes with the network topology, a phenomenon which is still actively researched in the context of smart grids and microgrids [8].

First-hitting-time methods in prognostics

The first hitting time (FHT) distribution of a drifted Wiener process is often found in the context of describing systems with degradation and their remaining useful life (RUL) [9, 10]. When researchers investigate the degradation of a system, they typically suggest that the system is considered to have failed once a certain limit is crossed. These researches apply the FHT distribution in order to conduct an offline, retrospective analysis of the RUL. The main aspect of our research is that we apply the same mathematical formulation in the context of an online, forward looking decision rule, and assessed against a hard dead-time budget. Furthermore, we suggest that the same running estimator provides other prognostic and similar outputs to [9, 10]. These outputs include the short-term violation risk, created by a short-horizon FHT, and a point estimate of the RUL in the month time frame created using the delta method on the hitting time, which in this context is considered to be a protection and fault classification mechanism, as suggested in Section 4.3.

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

Multiple-model estimation techniques for fault diagnosis

Interacting multiple-model (IMM) estimation was initially proposed in [11] and later expanded in [12]. Multi-model Kalman/EKF banks and IMM-based fault diagnosis were developed in [13] and [14], respectively, and are concerned with sensor fault diagnosis of gas turbines as well as the diagnosis of faults in electromechanical systems. In the mentioned works, fault diagnosis is considered a classification and estimation problem that is loosely coupled to the control task.

The principal contribution of the current work is the design of a mode bank with leakage mechanisms, and the integration of a closed-form protective trip control law with the mode bank to allow the same estimator to serve both protective control and fault diagnosis functions.

Early Access

3. Problem formulation

Let $x(t)$ denote the monitored, temperature-compensated pressure signal (Section 4.2) and $[x_{\text{low}}, x_{\text{high}}]$ the certified operating band. A conventional protection scheme trips when $x(t) \leq x_{\text{low}}$, independent of the rate of change of $x(t)$. We instead seek a trip rule of the form

$$x(t) \leq x_{\text{trip}}(t), \quad x_{\text{trip}}(t) \in [x_{\text{low}}, \quad x_{\text{high}}], \tag{1}$$

where $x_{\text{trip}}(t)$ is driven by the estimated rate of change so that it sits near x_{high} when $x(t)$ is falling quickly and relaxes toward x_{low} when it is decreasing slowly or not at all, and where the rule as a whole is governed by an explicit, auditable, model-based criterion: the trip fires at the first instant at which, under a conservative $(1-\varepsilon)$ upper bound on the estimated leak rate, the probability of breaching x_{low} within a fixed reaction-time budget T_{react} (the sum of estimator lag, debounce/persistence time, and breaker mechanical operating time; Section 5) reaches $1-\varepsilon$. The safety this affords is a model-based design property, assessed empirically in Section 6, not a formally demonstrated bound (Section 4.4).

4. Methodology

4.1. Choked-orifice leak physics

For flow through a small orifice under choked (critical) conditions, mass flow rate is proportional to upstream pressure. Combined with the ideal-gas relation at locally constant temperature and enclosure volume, this gives

$$\frac{\mathrm{d}P(t)}{\mathrm{d}t} = -kP(t), \quad k \geq 0, \tag{2}$$

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

so that, for constant k, $P(t) = P_0 \exp\big(-k(t - t_0)\big)$. The leak-conductance parameter k (units s^{-1}) is zero absent a leak and increases with effective orifice area, so an online estimate of k is simultaneously the quantity needed for trip-time prediction and a natural defect-size / severity indicator.

4.2. Temperature compensation

SF_6 pressure varies with ambient temperature independent of any leak; at constant mass and volume, the ideal-gas relation gives $P \propto T$, i.e. a fractional pressure change of dT/T per unit temperature change – at $T \approx 293$ K this is ≈ 0.34% per °C, consistent with ≈ 0.3–0.4% per °C figure commonly cited for SF_6 near typical GIS operating conditions. All processing in this article operates on the temperature-compensated signal

$$\frac{P_{20}(t) = P(t) \cdot T_{\text{ref}}}{T_{\text{amb}}(t)}, \quad T_{\text{ref}} = 293.15 \text{ K}, \tag{3}$$

which removes this effect at the source rather than attempting to distinguish it downstream from a genuine slow leak (Section 6.2).

4.3. Physics-informed interacting multiple-model (IMM) estimator

We track the state $x = [P,\ k]^T$ with the (exactly integrated, for constant k over one sample) nonlinear process model

$$P_{i+1} = P_i \exp(-k_i T_s), \quad k_{i+1} = k_i + w_k, \quad w_k \sim N(0, q_k) \tag{4}$$

and scalar measurement $z_i = P_i + v_i$, $v_i \sim N(0, \sigma_P^2)$. Because the process model (4) is bilinear in (P, k), each mode is tracked with an extended Kalman filter (EKF) linearising about the current estimate, with Jacobian

$$F_i = \begin{bmatrix} \exp(-k_i T_s) & -P_i T_s \exp(-k_i T_s) \\ 0 & 1 \end{bmatrix}. \tag{5}$$

A sigma-point (unscented) alternative to Jacobian linearisation [15] and recursive Bayesian alternatives such as particle filtering [16] have both been examined for state estimation in power-system settings. The sigma-point route was not pursued here, since the scalar process nonlinearity in (4) is mild over one 0.1 s sample and the EKF's lower per-step computational cost is preferable for a 10 Hz protection-grade real-time implementation.

Rather than committing to one value of the process-noise intensity q_k – which forces a compromise between stable, low-variance tracking (good for slow/moderate leaks) and fast convergence to a large k (needed for rupture-class events) – we run three parallel mode-matched filters distinguished only by q_k, derived directly from the physical regimes of Section 4.1 (Table 1).

Table 1. IMM mode set derived from the choked-orifice leak physics

Mode	Physical interpretation	q_k
H (healthy)	no leak, $k \approx 0$	$(10^{-9})^2$

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

Mode	Physical interpretation	q_k
L (orifice leak)	fixed small defect, slowly varying k	$(10^{-5})^2$
BUP (burst)	abrupt/growing defect, k can jump	$(5 \times 10^{-3})^2$

The three mode-matched filters are combined with the standard IMM recursion [13, 14]: mixing each mode's prior with a Markov transition matrix **Π**, mode-matched EKF prediction/update, likelihood-weighted mode-probability update, and a probability-weighted combination of the mode states and covariances at each step. The combined mode probabilities $[\mu^{\mathrm{H}}, \mu^{\mathrm{L}}, \mu^{\mathrm{BUP}}]$ are an online estimate of P(healthy), P(orifice leak), P(burst), giving fault classification as a by-product of the same recursion that produces the state estimate used for the trip decision.

4.4. First-hitting-time risk criterion

Let $a(t) = P(t) - x_{\mathrm{low}} > 0$ be the estimated gap to the floor and model its short-horizon evolution as a Wiener process with drift, $\mathrm{d}Y = -\mu\mathrm{d}t + \sigma\mathrm{d}W$, $Y(0) = a$. The first-passage time $T_{\mathrm{hit}} = \inf\{t : Y(t) = 0\}$ follows an inverse-Gaussian (Wald) distribution with closed-form cumulative distribution function (CDF)

$$\Psi(a, \mu, \sigma, \tau) = \Phi\left((\mu\tau - a)/(\sigma\sqrt{\tau})\right) + \exp(2\mu a/\sigma^2)\Phi\left(-(\mu\tau + a)/(\sigma\sqrt{\tau})\right), \qquad (6)$$

where Φ is the standard normal CDF (evaluated in log-space via log Φ for numerical stability, since the two terms can individually overflow/underflow while their sum remains well-scaled).

We validated Eq. (6) against 2×10^5 Euler–Maruyama sample paths of the underlying stochastic differential equation ($a = 10$, $\mu = 2$, $\sigma = 0.5$, $\mathrm{d}t = 0.01$): at $\tau = 5$ the simulated hitting probability was 0.5132 against 0.5222 from (6) – a gap of roughly 8 × the naive Monte Carlo standard error (≈ 0.0011 at $n = 2 \times 10^5$), consistent with known Euler–Maruyama first-passage discretisation bias rather than a defect in the closed form; both are numerically negligible (~Φ(−16)) for $\tau = 1$ and saturate to 1 by $\tau = 10$, matching the deterministic hitting time $a/\mu = 5$.

Two distinct uncertainties

The drift $\mu = kP$ is itself an estimate with error variance $\sigma_\mu{}^2$, obtained from the IMM's combined covariance via the delta method,

$$\sigma_\mu^2 = \hat{k}^2\mathrm{Var}(\hat{P}) + \hat{P}^2\mathrm{Var}(\hat{k}) + 2\hat{k}\hat{P}\mathrm{Cov}(\hat{P}, \hat{k}). \qquad (7)$$

This parameter (epistemic) uncertainty is distinct from the process (aleatoric) diffusion σ in (6), which represents genuine short-horizon wander of the trajectory about its mean path (sensor noise and unmodelled dynamics) and is treated as a fixed, small, sensor-characterised constant $\sigma = \sigma_{\mathrm{proc}}$. Conflating the two into a single margin – as a naive $k\sigma$ heuristic does – is counterproductive: inflating the diffusion term σ in (6) to stand in for parameter uncertainty can, for a fixed target probability, lower the required threshold in the dominant-term approximation

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

$a \approx \mu\tau - z_\varepsilon\sigma \cdot \sqrt{\tau}$, i.e. move it the wrong way, because larger σ also makes noise-driven hitting more likely on its own. We instead handle parameter uncertainty by inflating the drift used in (6) to a conservative upper bound,

$$\mu_c = \max\left(\hat{\mu} + z_\varepsilon\sigma_\mu, 0\right), \quad z_\varepsilon = \Phi^{-1}(1 - \varepsilon), \tag{8}$$

and let σ_{proc} alone populate (6). The trip rule is then

$$\text{trip when} \quad \Psi\left(a(t), \mu_c(t), \sigma_{\text{proc}}, T_{\text{react}}\right) \geq 1 - \varepsilon. \tag{9}$$

The tolerance ε plays a double role. First, $z_\varepsilon = \Phi^{-1}(1 - \varepsilon)$ sets the conservatism of the drift bound (8): the smaller ε, the more pessimistic the assumed leak rate, and the earlier (9) fires. Second, $1 - \varepsilon$ is the certainty threshold in (9): the rule waits until breach within T_{react} is near-certain under that pessimistic drift, which is what keeps the early-trip lead small for slow leaks (Section 6.1). The two effects act in opposite directions, and Section 6.4 shows empirically that the conservatism of (8) dominates, so that tightening ε monotonically increases the safety margin. We emphasise that the composite rule does not carry a formally demonstrated guarantee that the true pressure remains above x_{low} at breaker completion; that property is a model-based design target under the local Wiener approximation and the delta-method covariance of this section, and is assessed empirically across leak severities in Section 6.

Early Access

4.5. Baseline and ablation methods

To attribute performance to specific design decisions rather than to the proposed method as a monolithic whole, we compare five architectures:

Method A (static gain-scheduled)

$$x_{\text{trip}} = x_{\text{low}} + \left(x_{\text{high}} - x_{\text{low}}\right) \cdot \text{sat}\left(\left|\hat{\mathrm{d}P/\mathrm{d}t}\right| / \mathrm{d}P_{\text{max}}/\mathrm{d}t\right),$$

with the rate estimate a short boxcar-smoothed finite difference and the normalisation rate a fixed constant chosen, as a realistic design choice, to the moderate-leak characteristic rate rather than the true worst-case rupture rate – representing the common practice of sizing a gain to an expected rather than worst-case event.

Method B (deterministic least-squares (LS) extrapolation)

Dual-window (short/long) LS slope estimate feeding the linear time-to-limit (TTL) approximation

$$x_{\text{trip}} = \text{clip}\left(x_{\text{low}} + \Delta P + \left|\hat{\frac{\mathrm{d}P}{\mathrm{d}t}}\right| T_{\text{react}}, x_{\text{low}}, x_{\text{high}}\right),$$

with $\Delta P a k_0 \sigma_{\text{slope}}$ margin. Fixed-coefficient, fixed worst-case execution time; the most readily certifiable variant, but subject to window-fill latency.

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

Method C (linear Kalman filter)

The same linear TTL law as Method B, but with the rate estimate from a recursive constant-rate Kalman filter with innovation-based adaptive process noise – recursive and low-latency, but does not model the true exponential decay.

Method D (single-mode EKF)

The nonlinear (*P*, *k*) EKF of Section 4.3 run alone (no mode bank), feeding the exact exponential TTL law with a $k_0\sigma$ margin rather than the FHT criterion – isolates the benefit of correct nonlinear modelling from the benefits of

(i) the multi-mode bank and
(ii) the risk-based criterion.

Method E (proposed)

The physics-informed IMM of Section 4.3 feeding the first-hitting-time criterion (9). Methods A–D apply an *M*-of-*M* debounce (persistence) filter to their threshold-comparison decision; Method E's risk criterion (9) is evaluated directly each sample and is not additionally debounced, since it already integrates estimation uncertainty over the reaction-time horizon T_{react} itself. All methods remain subject to a hard, rate-independent back-stop trip at $x(t) \le x_{\mathrm{low}}$.

Early Access

4.6. Prognostics: short-term breach risk and long-term remaining-time estimate

The same running (*P*, *k*) estimator can additionally answer two slower questions that serve maintenance planning rather than protection:

(i) what is the probability that pressure breaches the lockout floor x_{low} within the next few days to weeks (a near-term alert), and
(ii) how many months of service remain before it does so at the current estimated leak rate (a long-term planning figure). We deliberately answer these with two different formulas, each valid on its own timescale, rather than one formula stretched across both.

4.6.1. Short-term breach-risk probability

For a near-term horizon τ_{short} of a few days to a month, the local Wiener-with-drift approximation behind Eq. (6) remains legitimate, because τ_{short} is still small relative to the multi-month-to-year time constant of the leaks this framework targets – exactly the validity condition required for the trip criterion – that the reaction-time budget be small relative to the leak's own time constant – just evaluated at a different (still short) τ. We therefore reuse Eqs. (6) and (8) verbatim, with the same conservative drift inflation used for the protective trip; unlike the trip rule (9), the result is reported directly as a probability (with a 50% alert cutoff used for classification in Section 6.5.2) rather than compared against $1 - \varepsilon$,

$$P(\text{breach within } \tau_{\mathrm{short}}) = \Psi\big(a(t), \mu_c(t), \sigma_{\mathrm{proc,slow}}, \tau_{\mathrm{short}}\big), \tag{10}$$

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

with $a(t) = \hat{P}(t) - x_{\text{low}}$ and $\mu_c(t)$ the same Eq. (8) quantity computed from the running estimator's current ($\hat{P}$, $\hat{k}$) and joint covariance via Eq. (7). The only change from the protective trip's use of (6) is the horizon ($\tau_{\text{short}} \in \{7, 30\}$ days here) and the diffusion constant: the fixed $\sigma_{\text{proc}} = 0.15\ \text{kPa}/\sqrt{s}$ (a tuned design constant for the sub-10-s protection path, deliberately smaller than the raw discretisation value $\sigma_P/\sqrt{T_s} \approx 0.95$ kPa/√s, since the drift-inflation term (8) already absorbs estimation uncertainty; Section 5) does not represent a meaningful wander rate over days to weeks, so we use instead $\sigma_{\text{proc,slow}} = \sigma_P/\sqrt{T_s}$ with the hourly sample period T_s = 3600 s – the diffusion coefficient whose discretisation reproduces the same per-sample measurement-noise variance (σ_P = 0.3 kPa), giving $\sigma_{\text{proc,slow}} = 0.005\ \text{kPa}/\sqrt{s}$, a small, sensor-characterised constant chosen analogously to σ_{proc}. Reusing (6)/(8) this way stops being appropriate once the leak itself is fast enough that even a 7–30 day window is not short relative to its own remaining life; Section 6.5 shows this boundary concretely for the 1-month scenario, whose entire natural crossing time is only 30.4 days.

4.6.2. Long-term remaining-time estimate

Equation (6) cannot simply be re-evaluated at a month-scale horizon τ: its short-horizon validity condition – that the horizon be small relative to the leak's own time constant – is violated by construction there, and the drift inflation (8) is likewise a short-horizon device . We therefore use the same delta-method covariance propagation as Eq. (7), but apply it to the exact hitting time of the exponential decay law (2) rather than to the local-Wiener drift. For current estimates $\hat{P} > x_{\text{low}}$ and $\hat{k} > 0$, the trajectory $\hat{P}\exp(-\hat{k}t)$ reaches x_{low} after

$$T(\hat{P}, \hat{k}) = \ln\left(\hat{P}/x_{\text{low}}\right)/\hat{k}, \tag{11}$$

the point remaining-time (RUL-style) estimate if the estimated leak conductance persists. With $\partial T/\partial \hat{P} = 1/\hat{k}\ \hat{P}$ and $\partial T/\partial \hat{k} = -T/\hat{k}$, first-order propagation of the estimator's joint covariance – exactly the technique of (7), applied to (11) instead of μ – gives

$$\text{Var}(T) = \text{Var}\left(\hat{P}\right)/\left(\hat{k}\hat{P}\right)^2 + \left(T/\hat{k}\right)^2\text{Var}\left(\hat{k}\right) - \left(2T/\left(\hat{k}^2\hat{P}\right)\right)\text{Cov}\left(\hat{P}, \hat{k}\right), \tag{12}$$

and we report the point estimate (11), converted to months, together with the 80% central interval $[\hat{T} - z_{0.9}\sigma_T, \hat{T} + z_{0.9}\sigma_T]$ with $\sigma_T = \sqrt{\text{Var}(T)}$, $z_{0.9} = \Phi^{-1}(0.90) \approx 1.28$ (a $p10$–$p90$ band), the lower bound clipped at zero.

The near-term alert (10) reuses the short-horizon stochastic model (6)/(8), where the aleatoric wander of the trajectory over days is what dominates and the horizon is short enough that the Wiener approximation still holds; the long-term estimate (11)–(12) instead uses the exact deterministic hitting-time functional of the physical decay law (2), where the dominant uncertainty is epistemic – how well (P, k) are currently known – propagated by the same delta-method machinery as (7) but applied to a different functional, appropriate to a horizon where (6)'s own approximation no longer holds. Both run on the same course hourly trend-monitoring path (Section 5, Study 3).

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

5. Experimental setup

Unless noted otherwise: x_{low} = 500 kPa, x_{high} = 650 kPa, sample period T_s = 0.1 s (10 Hz protection-grade sampling), sensor noise σ_P = 0.3 kPa (nominal) or 0.6 kPa (degraded sensor, Study 2), protective-trip diffusion constant σ_{proc} = 0.15 kPa/√s, reaction budget T_{react} = 3 s, debounce M = 5 samples (0.5 s) for the threshold-comparison baselines A–D (Method E's risk criterion is evaluated undebounced; Section 4.5), ablation-baseline margin multiplier k_σ = 3.5, proposed-method risk tolerance $\varepsilon = 10^{-4}$, IMM transition matrix

$$\mathbf{\Pi} = [0.995\ 0.004\ 0.001;\ 0.010\ 0.985\ 0.005;\ 0.005\ 0.010\ 0.985],$$

(rows/columns ordered H, L, BUP) with prior mode probabilities [0.90, 0.08, 0.02].

Study 1 (ablation)

Three discrete leak time constants – τ = 15 s (rupture-class), 600 s (moderate), 36 000 s (chronic) – plus a no-leak/thermal-noise scenario, evaluated for all five methods via Monte Carlo (200/150/40/20 repetitions respectively; Method E used 200/150/15/10 repetitions for the slow/false-trip cases specifically, to bound the added cost of its three-filter mode bank). The resulting asymmetry means the worst-case-minimum comparisons for the chronic and no-leak classes in Table 2 carry less statistical power for Method E than for the baselines (Section 6.1).

Study 1b (temperature compensation)

A separate 48-hour, T_s = 60 s trend-monitoring-path simulation with a ± 15°C diurnal ambient cycle and zero true leak, comparing rate estimates computed on raw P versus P_{20}.

Study 2 (expanded validation)

Method E vs. Method A across

$$\tau \in \{10, 24.3, 59, 143, 348, 846, 2057, 5000\}\ \text{s} \times \sigma_P \in \{0.3, 0.6\}\ \text{kPa},$$

with the leak onset time itself randomised uniformly in [0, 60] s (rather than fixed at t = 0), 35 Monte Carlo repetitions per condition, reporting the mean safety margin at breaker completion with 95% confidence intervals and the fraction of runs with a negative (unsafe) margin.

Sensitivity analysis

Method E only, moderate-leak scenario (τ = 600 s), sweeping the risk tolerance $\varepsilon \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$ (holding $T_{\text{react}} = 3$ s) and, separately, $T_{\text{react}} \in \{1, 2, 3, 5, 8\}$ s (holding $\varepsilon = 10^{-4}$), with 30 repetitions per value for the margin/lead metrics and, for the corresponding false-trip check, 10 (ε sweep) and 5 (T_{react} sweep) independent 30–60 min no-leak traces per value.

Study 3 (prognostic validation)

The remaining-time estimate of Section 4.6 is evaluated on an hourly (T_s = 3600 s) trend-monitoring path – coarser than both the 10 Hz protection path and the 60 s temperature-compensation path, as appropriate for a months-scale horizon – with the same σ_P = 0.3 kPa sensor noise, over three leak scenarios whose true natural band-crossing times $\ln(x_{\text{high}}/x_{\text{low}})/\lambda$ equal 1, 3,

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

and 6 months ($\lambda = 9.98 \times 10^{-8}$, 3.33×10^{-8}, and $1.66 \times 10^{-8}\ s^{-1}$ respectively), with 20 Monte Carlo noise realisations each, simulated from leak onset to floor crossing. The (P, k) EKF of Section 4.3 is re-tuned for the hourly timescale ($q_k = 10^{-20}$, $q_p = (0.02)^2$; empirically verified to converge to the true k within roughly one to two weeks of hourly data and to remain stable thereafter). At monitoring-window checkpoints of 3, 7, 14, 30, and 60 days of accumulated data (those shorter than each scenario's crossing time) we record the predicted remaining time and its 80% interval, the true remaining time, the prediction error, and whether the interval covers the truth.

6. Results

6.1. Ablation study across leak severities (Study 1)

Table 2 reports, for each of the three leak classes and all five methods, the mean trip time, the safety margin at breaker completion $P_{\text{true}}(t_{\text{trip}} + T_{\text{react}}) - x_{\text{low}}$ (mean, std, and worst-case minimum over the Monte Carlo repetitions), and the early-trip lead as a percentage of the leak's natural band-crossing time $\ln(x_{\text{high}}/x_{\text{low}})/\lambda$. This lead is unwarranted conservatism for the moderate and chronic classes, where, on average, all five methods are safe. For the rupture class, a large lead is the intended safety margin rather than waste, since the fixed-window baselines do not respond in time. No method produced a false trip in the no-leak/thermal-noise scenario (Section 5; 0/20 runs for A–D, 0/10 for E, combined ≥ 120 simulated hours). The last column gives the root-mean-square error (RMSE) of each method's rate estimate.

Table 2. Study 1: five-way ablation across rupture-class, moderate, and chronic leaks. Margin is $P_{\text{true}}(t_{trip} + T_{\text{react}}) - x_{low}$ in kPa (negative = unsafe); Lead % is early-trip time as a percentage of the natural band-crossing time $\ln(x_{\text{high}}/x_{\text{low}})/\lambda$ – unnecessary conservatism for moderate/chronic leaks, but the intended safety margin for the rupture class (Section 6.1)

Leak class	Method	Trip t (s)	Margin mean ± std	Margin min	Lead %	Rate RMSE
Rupture (τ = 15 s)	A: static gain-scheduled	1.50	−18.47 ± 0.00	−18.47	61.9	9.163
	B: LS-extrapolation (linear)	1.40	−15.25 ± 0.00	−15.25	64.4	8.875
	C: linear Kalman (linear)	0.80	4.49 ± 0.53	1.18	79.6	4.395
	D: single EKF + ad hoc margin	0.89	1.65 ± 1.38	−2.15	77.5	2.480
	E: IMM + FHT (proposed)	0.20	**24.99 ± 0.68**	**21.64**	94.8	2.773
Moderate (τ = 600 s)	A	5.12	141.27 ± 3.27	131.11	96.7	0.299
	B	151.62	2.34 ± 0.83	0.68	3.7	0.204

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

Leak class	Method	Trip t (s)	Margin mean ± std	Margin min	Lead %	Rate RMSE
	C	153.28	0.95 ± 1.34	−0.23	2.6	0.088
	D	153.98	**0.37 ± 0.22**	−0.23	2.2	0.125
	E	152.54	1.57 ± 0.89	**0.77**	3.1	0.111
Chronic (τ = 36 000 s)	A	6912.6	36.42 ± 5.46	28.32	26.8	0.283
	B	9235.4	2.88 ± 0.08	2.67	2.2	0.180
	C	8851.6	8.28 ± 3.39	3.43	6.3	0.079
	D	9397.1	**0.63 ± 0.08**	**0.48**	0.5	0.022
	E	8883.5	7.84 ± 5.01	3.25	5.9	0.074

Bold marks the best (safest worst-case margin, or tightest mean margin where all methods are already safe) entry per class

Figures 1–3 visualise these results. Three findings stand out.

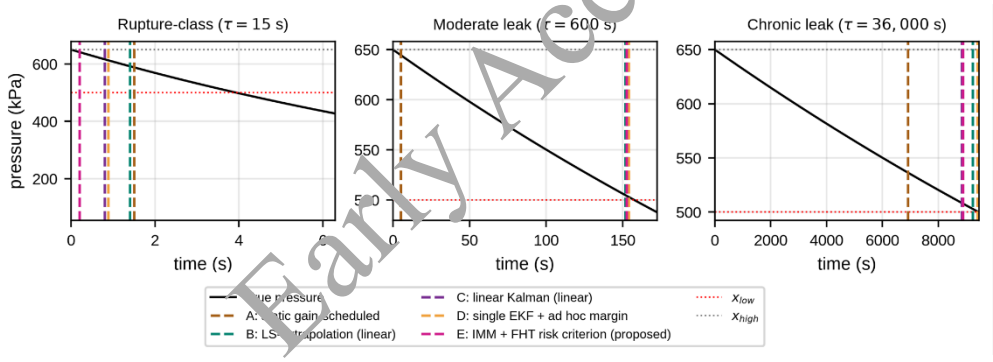


Fig. 1. True pressure trajectories and mean trip times for all five methods, for the rupture-class, moderate, and chronic leak scenarios of Study 1

Method E is the only method with a positive worst-case margin in every tested class. For the rupture-class leak, the fixed-window batch estimators (A, B) cannot react in time at all: their window-fill latency alone exceeds the ≈ 3.9 s natural crossing time of a τ = 15 s leak, giving deeply negative margins (−18.5, − 5.2 kPa). The recursive linear Kalman filter (C) and single-mode EKF (D) do much better by avoiding window-fill latency, but D's ad hoc $k_\sigma\sigma$ margin is occasionally unsafe in the worst observed case (−2.15 kPa) because the EKF's estimate of k has not yet converged from its zero-leak prior when the event is this fast. Method E resolves this: its burst mode is specifically designed (via a large process-noise q_k) to converge within about one second, and the formal first-hitting-time criterion at $\varepsilon = 10^{-4}$ gives a large positive margin (24.99 ± 0.68, worst case 21.64 kPa) with no observed unsafe run across 200 repetitions.

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

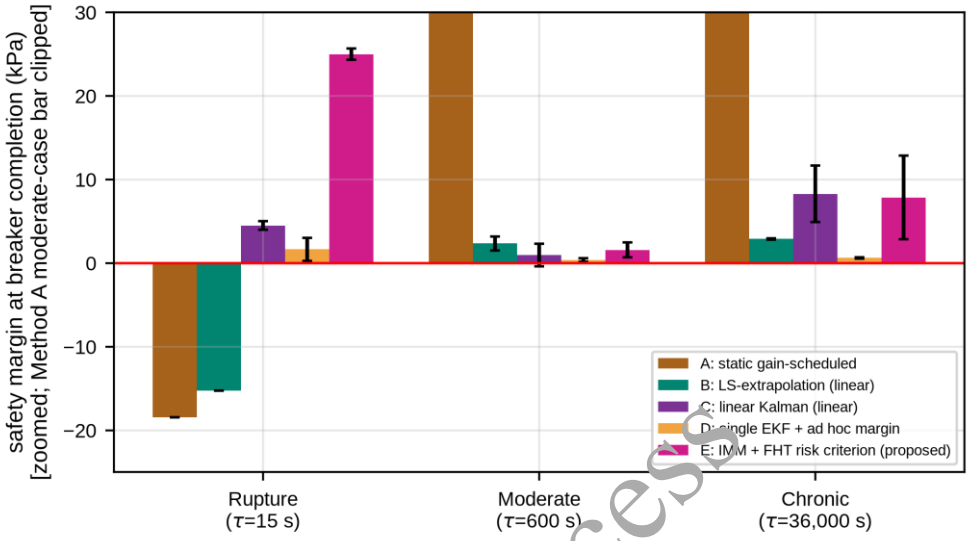


Fig. 2. Safety margin at breaker completion by method and leak class (zoomed; Method A's moderate-case bar, 141 ± 3 kPa, is clipped for legibility – see Table 2). Method E is the only method with a positive margin in every class

No single ablation baseline is uniformly tightest. For moderate and chronic leaks – where all five methods are already safe on average – the single-mode EKF (D) achieves the tightest mean margin and lowest wasted lead time (as low as 0.5% of the natural crossing time for the chronic case), edging out the proposed Method E (5.9%). This is a limitation: D's ad hoc margin is tuned to the same nominal confidence level as E's ε but, unlike E, does not correctly propagate the estimator's own drift uncertainty through the full nonlinear, time-varying risk calculation of Section 4.4 – so it is tighter on average at the cost of the occasional worst-case violation noted above. Method E trades a modest amount of average tightness for a safety guarantee that holds uniformly across all tested severities.

The arbitrary-gain baseline (A) fails in both directions. Its normalisation rate, sized to a plausible "expected" leak rather than the worst case, saturates almost immediately for the rupture-class leak (late, unsafe trip) while over-reacting to the moderate and chronic leaks, discarding 96.7% and 26.8% of their respective natural service life. This is the concrete failure mode motivating Sections 4.1–4.4: a single non-physical gain cannot simultaneously be conservative enough for fast events and permissive enough for slow ones.

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

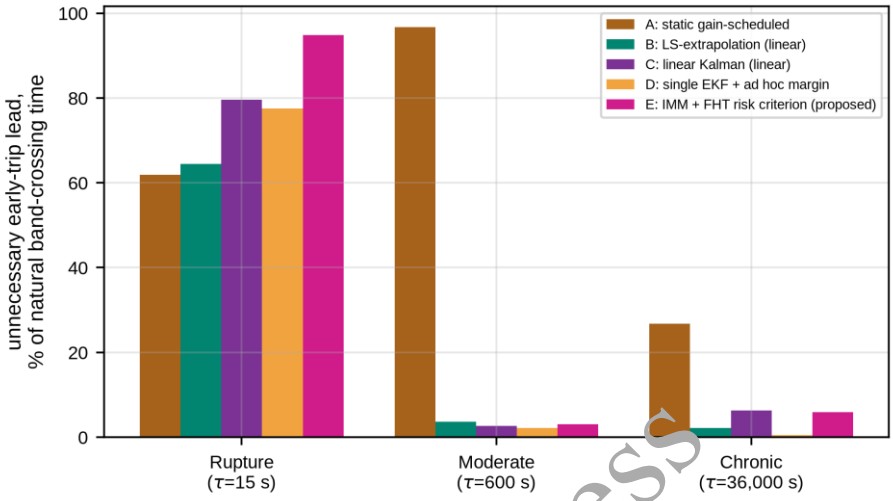


Fig. 3. Unnecessary early-trip lead time, as a percentage of each leak's natural band-crossing time, by method and leak class

6.2. Temperature compensation

Figure 4 shows the separate Study 1b trace: $a \pm 5$°C diurnal ambient cycle with zero true leak, on a coarser T_s = 60 s trend-monitoring path with a 4-hour rate-estimation window.

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

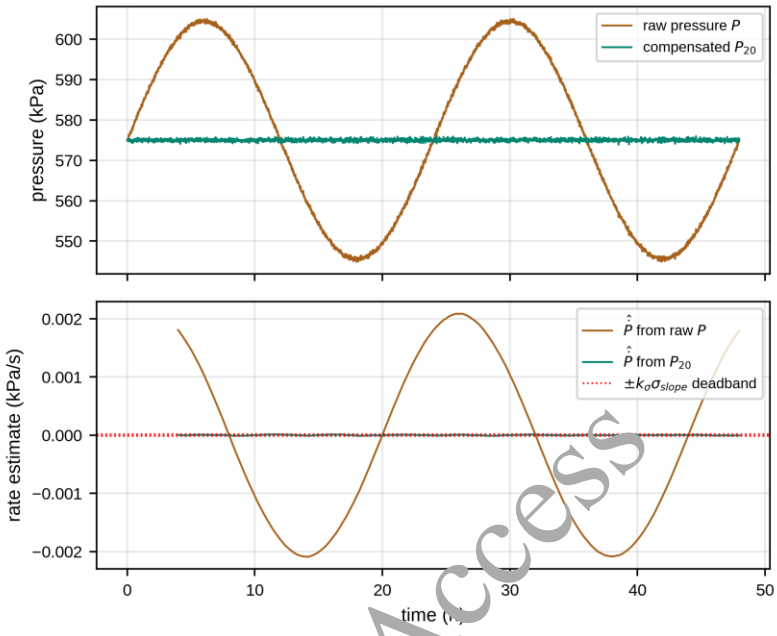


Fig. 4. Raw vs. temperature-compensated (P_{20}) pressure and rate estimate under a ±15°C diurnal cycle with zero true leak

Early Access

Raw (uncompensated) pressure produces a smooth, sustained apparent depressurisation rate peaking at 2.09×10^{-3} kPa/s – about 100 × the estimator's own noise-driven deadband (2×10^{-5} kPa/s) – which would persistently and incorrectly elevate an uncompensated trip threshold well above x_{low} for roughly half of every day, entirely from thermal cycling. Applying the P_{20} compensation of Eq. (3) reduces the peak spurious rate to 1×10^{-5} kPa/s, below the deadband, confirming that temperature compensation is a necessary pre-processing step for any leak-detection window operating on an hours-or-longer timescale, and is not automatically handled by the fast (10 Hz) protection-path window lengths used in Study 1 (which are short enough, at ≤ 30 s, to be structurally insensitive to a 24-hour thermal cycle regardless of compensation).

6.3. Expanded validation (Study 2)

Figure 5 sweeps the leak time constant continuously over $\tau \in [10, 5\,000]$ s (roughly two and a half decades) at two noise levels, with randomised leak onset, comparing Method E against the static-gain baseline (Method A); Table 2 already covers the τ = 15, 600, and 36 000 s points in more ablation detail, so Study 2 focuses on the continuum between and beyond them.

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

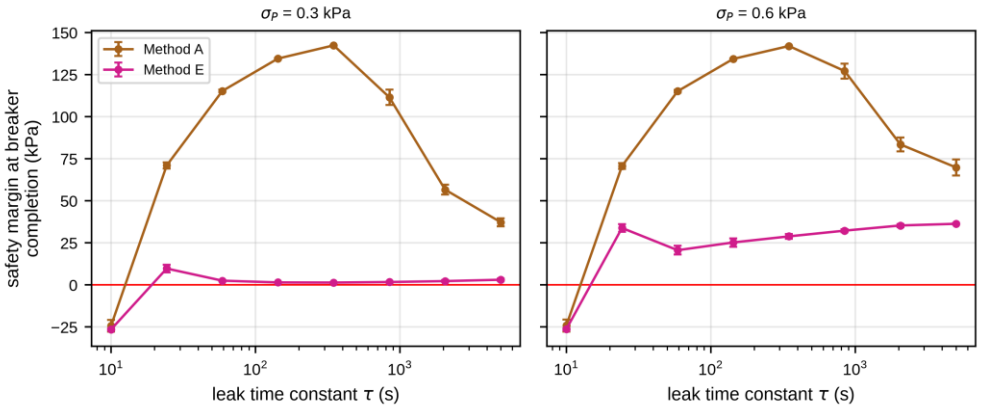


Fig. 5. Study 2: safety margin at breaker completion vs. leak time constant τ (log scale), for two sensor noise levels, with 95% confidence intervals over 35 repetitions per point (randomised leak onset). Method E's margin stays small and stable across more than two decades of leak severity; Method A's swings non-monotonically and far larger in magnitude. Both methods show a negative margin at τ = 10 s, a physical limit (Section 6.3)

Early Access

At τ = 10 s, both methods show a negative margin in 100% of repetitions (−24.3 ± 3.5 and −26.4 ± 0.4 kPa for A and E respectively at σ_p = 0.3). At τ = 10 s the natural band-crossing time is $\ln(x_{\text{high}}/x_{\text{low}})\cdot\tau \approx 2.6$ s, which is shorter than the T_{react} = 3 s reaction-time budget itself – even a hypothetical instantaneous, perfect detector cannot guarantee a positive margin once the physical transit time is shorter than the mechanical/electronic protection chain's own operating time. This is a fundamental physical limit, not a deficiency of the estimation or decision architecture, and is consistent with GIS engineering practice of relying on mechanical pressure-relief devices, not electronic protection, against sub-second catastrophic events.

For every $\tau \geq 24.3$ s, Method E's margin is small and remarkably stable – between 1.4 and 3.1 kPa at σ_P = 0.3 across more than a 200 × range of leak severity – while Method A's margin swings non-monotonically between 37 and 142 kPa over the same range, reflecting its single fixed gain being well-matched to only a narrow band of leak rates near its normalisation point and poorly matched elsewhere. At the degraded noise level (σ_P = 0.6), Method E's margin increases and stabilises around 29–36 kPa for $\tau \geq 348$ s – an appropriately more conservative response to genuinely worse rate-estimation uncertainty via the $\mu_c = \hat{\mu} + z_\varepsilon \sigma_\mu$ inflation of Section 4.4 – while remaining far more consistent than Method A's response to the same noise increase. No negative-margin runs were observed for either method at any $\tau \geq 24.3$ s across 35 repetitions × 2 noise levels × 7 rate values (490 runs per method).

6.4. Sensitivity analysis

Figure 6 sweeps Method E's two free parameters independently on the moderate-leak scenario (τ = 600 s). Tightening the risk tolerance from $\varepsilon = 10^{-2}$ to 10^{-6} increases the mean margin monotonically from 0.87 to 2.01 kPa (worst-case minimum from 0.43 to 1.10 kPa) at a modest

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

lead-time cost (4.05 → 5.41 s), with no false trips observed at any ε (0/10 no-leak runs per value; ≈ 38 simulated hours in total) – a smooth, predictable, one-parameter safety/availability trade-off, consistent with the design goal stated in Section 1. Sweeping the reaction-time budget T_{react} from 1 to 8 s (at fixed $\varepsilon = 10^{-4}$) similarly increases the margin monotonically (0.28 → 5.39 kPa mean), but at T_{react} = 1 s the worst-case margin over 30 repetitions was negative (−0.15 kPa) despite zero false trips across its 5 no-leak runs- the one place in this study where the nominal $\varepsilon = 10^{-4}$ guarantee was not empirically achieved.

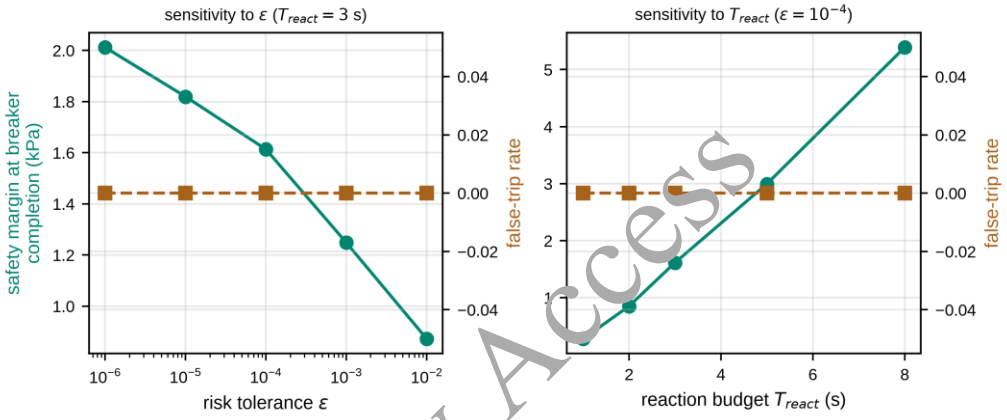


Fig. 6. Sensitivity of Method E's safety margin and false-trip rate to the risk tolerance ε (left) and reaction-time budget T_{react} (right), moderate-leak scenario

Early Access

We attribute this to T_{react} = 1 s being comparable to the IMM's own mode-convergence time-scale, violating the implicit assumption behind (6) that the dead-time window is long enough for the local Wiener-with-drift approximation and the delta-method uncertainty propagation to be accurate – a concrete boundary condition on the framework's applicability.

6.5. Prognostics validation: short-term risk and long-term remaining time (Study 3)

6.5.1. Long-term remaining-time estimate

Table 3 reports, for each leak scenario and monitoring-window checkpoint, the median predicted remaining time (11) against the true remaining time, the median 80% interval, the mean and worst single-run prediction error, and the empirical coverage of the nominal 80% interval over the 20 Monte Carlo repetitions.

Table 3. Study 3: prognostic remaining-time-to-critical-pressure validation. RUL is the remaining time until P reaches x_{low}, in months; error is (predicted − true)/true per run; coverage is the fraction of the 20 repetitions whose 80% interval contained the true remaining time

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

Scenario	Elapsed (d)	True RUL (mo)	Pred. RUL (mo)	80% interval (mo)	Mean err. (%)	Worst |err.| (%)	Coverage
1 month ($\lambda = 9.98 \times 10^{-8}$ s^{-1})	3	0.90	0.90	[0.89, 0.92]	+0.14	1.78	0.95
	7	0.77	0.77	[0.76, 0.78]	+0.04	0.47	1.00
	14	0.54	0.54	[0.53, 0.55]	−0.09	0.86	1.00
	30	0.014	0.014	[0.013, 0.015]	−1.35	8.95	0.80
3 months ($\lambda = 3.33 \times 10^{-8}$ s^{-1})	3	2.90	2.92	[2.76, 3.08]	+0.42	4.92	0.95
	7	2.77	2.77	[2.66, 2.88]	+0.12	1.24	1.00
	14	2.54	2.54	[2.44, 2.64]	−0.20	2.09	1.00
	30	2.01	2.01	[1.93, 2.09]	−0.33	1.94	1.00
	60	1.03	1.03	[0.98, 1.07]	−0.37	1.75	1.00
6 months ($\lambda = 1.66 \times 10^{-8}$ (s^{-1})	3	5.90	5.97	[5.32, 6.62]	+0.90	9.28	0.95
	7	5.77	5.78	[5.32, 6.24]	+0.25	2.46	1.00
	14	5.54	5.53	[5.10, 5.96]	−0.37	3.97	1.00
	30	5.01	4.98	[4.59, 5.37]	−0.62	3.59	1.00
	60	4.03	4.01	[3.69, 4.33]	−0.64	3.18	1.00

The exact-hitting-time prognostic of Section 4.6 locks onto the true remaining life within days of hourly monitoring in every scenario. With only 3 days of data the median prediction is already within about 1% of the truth (mean error +0.14% to +0.90% across the three scenarios), though individual noise realisations can still deviate by up to 9.3% (6-month scenario) and the interval is correspondingly wide ([5.32, 6.62] months around a 5.90-month truth). By 7–14 days the worst single-run error has contracted to 0.5–4%, and the 80% interval tightens steadily as data accumulate – for the 6-month scenario, from ±0.65 months at day 3 to ±0.32 months at day 60. Figure 7 shows this convergence for the 3-month scenario: after a brief initial transient while the EKF's $\hat{k}$ pulls away from its zero-leak prior, the median prediction tracks the declining true remaining time essentially exactly for the rest of the leak's life.

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

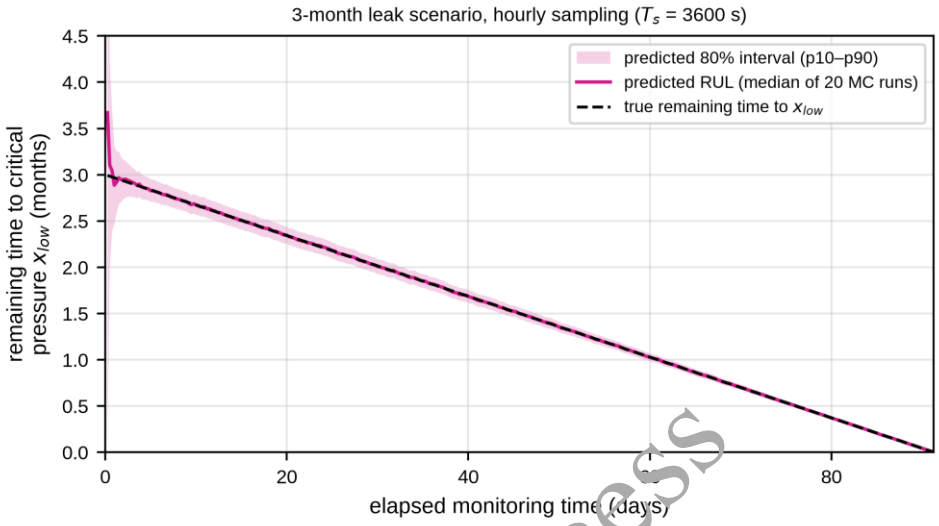


Fig. 7. Study 3, 3-month leak scenario: median predicted remaining time to x_{low} with its 80% interval (shaded; medians over 20 Monte Carlo repetitions) against the true remaining time, as a function of elapsed monitoring time on the hourly trend path. The prediction converges to the truth within days and its interval tightens as data accumulate

Early Access

Empirical coverage of the nominal 80% interval is 95–100% at nearly every checkpoint – conservative rather than optimistic, the safe direction for a maintenance-planning quantity. One checkpoint below nominal (coverage 0.80, day 30 of the 1-month scenario) happens when there is about 0.4 days of life left, and the interval has shrunk to approximately ±0.001 months (about 40 min). This is a scenario where a scheduled maintenance decision would have already been made. In this case, the protection path of Section 4.4 would be the operative safeguard. The small systematic pattern – slight overprediction at the shortest windows, where $\hat{k}$ still carries a remnant of its zero-leak prior, turning into a sub-1% underprediction bias at long windows – is well inside the reported interval at every checkpoint.

6.5.2. Short-term breach-risk probability

Table 4 evaluates the near-term alert (10) at $\tau_{\mathrm{short}} \in \{7, 30\}$ days, at the same checkpoints and over the same 20 Monte Carlo repetitions per scenario, against the true remaining time at each checkpoint (deterministic given the scenario, since only the sensor noise driving the estimator differs across repetitions). Using a 50% predicted-probability alert threshold as the classification cutoff, the near-term probability correctly separates "breach expected within τ_{short}" from "not expected" in 27 of the 28 scenario × checkpoint × horizon combinations evaluated (96.4% accuracy): every one of the 22 combinations predicted quiet (probability below 50%) was genuinely quiet, and 5 of the 6 combinations predicted in alert (probability at or above 50%) genuinely went on to breach within that horizon.

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

Table 4. Study 3: representative short-term breach-risk probabilities (selected checkpoints; full 28-combination set underlies the aggregate accuracy above). Predicted probability is the mean over 20 Monte Carlo repetitions of $\Psi(a, \mu_c, \sigma_{\mathrm{proc,slow}}, \tau_{\mathrm{short}})$; actual breach is whether the (deterministic) true remaining time at that checkpoint is within τ_{short}

Scenario	Elapsed (d)	True remaining (d)	τ_{short} (d)	Pred. *P*(breach)	Actual breach?
1 month	3	27.4	7	≈ 0	No
1 month	3	27.4	30	1.00	Yes
1 month	14	16.4	7	≈ 0	No
1 month	14	16.4	30	1.00	Yes
1 month	30	0.4	7	1.00	Yes
1 month	30	0.4	30	1.00	Yes
3 months	30	61.3	30	6.5×10^{-7}	No
3 months	60	31.3	7	≈ 0	No
3 months	60	31.3	30	0.79	No
6 months	60	122.6	7	≈ 0	No
6 months	60	122.6	30	9.0×10^{-16}	No

The predicted probability behaves exactly as the short-horizon Ψ formula should when the true remaining time is far from τ_{short}: essentially 0 whenever the leak has months of remaining life against a 7–30 day window (e.g. 9.0×10^{-16} for the 6-month scenario at day 60, 122.6 days still remaining), and essentially 1 once the true remaining time drops below the horizon (e.g. the 1-month scenario's day-30 checkpoint, 0.4 days remaining). The single miscall in the 28-combination set is the 3-month scenario at day 60 against the 30-day horizon: true remaining time is 31.3 days, i.e. just 1.3 days beyond the 30-day cutoff, yet the predicted probability is already 0.79 – a near-boundary false alarm, and arguably the conservative-side error a maintenance alert should prefer. Figure 8 visualises the full 28-combination set: predicted probability collapses sharply from ≈ 1 to ≈ 0 almost exactly at the τ_{short} boundary (vertical dotted line) in both panels, with the single boundary miss visible as the one off-colour point above the 50% line on the right. The 1-month scenario illustrates the boundary condition stated in Section 4.6.1: its entire natural crossing time is only 30.4 days, so the 30-day alert horizon is not considered a small fraction of remaining life at the start of monitoring but very nearly the whole of it.

The formula still classifies every 1-month-scenario checkpoint correctly here (Table 4), because at these particular checkpoints the true remaining time happens to be clearly on one side or the other of 30 days; the concern the validity condition raises is not that (6)/(8) fail at this horizon, but that as the horizon and the leak's own time constant converge, the short-horizon approximation's error margin is no longer small by construction, and a fast-progressing leak could produce a less clean-cut case than the ones sampled here.

Early Access

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

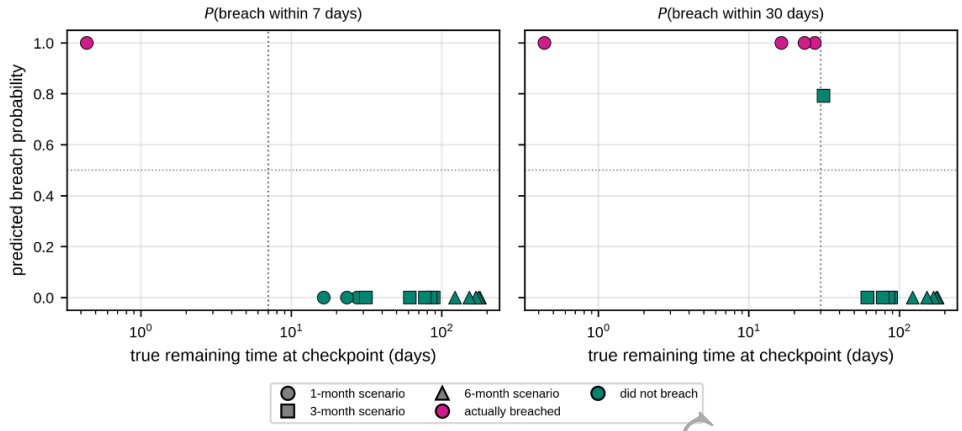


Fig. 8. Study 3: predicted near-term breach probability vs. true remaining time at the checkpoint, for the 7-day (left) and 30-day (right) alert horizons, across all three scenarios and checkpoints. Marker colour indicates whether a breach actually occurred within that horizon

Early Access

7. Conclusions

This research introduces a novel first-hitting-time risk metric for adaptive protective tripping characterised by an inverse-Gaussian closed-form solution and demonstrates the concept through explicit stochastic simulation. The design integrates both the protection decision and the real-time online fault classification using a physics-based, multimode estimator. Improvements in the method stem from particular, independently motivated design decisions, which were verified using a five-way ablation study. The method was further described by an extended Monte Carlo validation, which incorporated a sensitivity analysis of adjustable risk tolerance and reaction-time budgets. Using the same estimator, short-term breach risk was calculated, and the first hitting time metric was successfully applied to correctly classify 27 of 28 scenario/checkpoint/horizon combinations, with alerts triggered at the 50% threshold. The estimator also yielded a long-term, months-scale, remaining time prediction, obtained by applying the same delta-method uncertainty propagation to the exact exponential hitting time. The predictions were found to align with the ground truth within a few percent. The main limitations of this study are the absence of field experiments on real-world equipment, since the results are obtained from simulation, and the need for careful calibration of the method's design parameters; hardware-in-the-loop validation and a systematic calibration procedure are the subject of future work.

Appendix. Numerical evaluation of the first-hitting-time CDF

Equation (6) is evaluated by computing $\log \Phi(z_1)$ and $\log \Phi(z_2) + 2\mu a/\sigma^2$ (via a numerically stable $\log \Phi$ implementation) and combining as

This paper has been accepted for publication in the AEE journal. This is the version, which has not been fully edited and content may change prior to final publication.
Citation information: DOI 10.24425/aee.2027.8063

$$\exp(m)[\exp(\log \Phi(z_1) - m) + \exp(\log \Phi(z_2) + 2\mu a/\sigma^2 - m)],$$

with m the maximum of the two log-terms, avoiding overflow in $\exp(2\mu a/\sigma^2)$ when $\Phi(z_2)$ is simultaneously extremely small.

References

[1] Lech M., Węgierek P., *Effect of selected parameters on the electrical strength of a high-voltage vacuum insulation system*, Archives of Electrical Engineering, vol. 72, no. 3, pp. 597–611 (2023), DOI: 10.24425/aee.2023.146039.
[2] International Electrotechnical Commission, IEC 62271-203: *High-voltage switchgear and control-gear – Part 203: AC gas-insulated metal-enclosed switchgear for rated voltages above 52 kV*, IEC (2022).
[3] National Electrical Manufacturers Association, *NEMA US 80020-2022: Field Leak Detection of SF_6 Gas-Insulated Equipment*, NEMA (2022).
[4] Bula Bunjaku I., Alidemaj A., Gashi D., Latifi D., Bula E., *Distance-dependent detection reliability of fixed-threshold residual earth-fault indicators in medium-voltage feeders*, Archives of Electrical Engineering, vol. 75, no. 3, pp. 501–522 (2026), DOI: 10.24425/aee.2026.1525.
[5] Karimabadi A., Hajiabadi M.E., Kamyab E., *Prioritization approach for circuit breakers to equip with condition monitoring devices*, Archives of Electrical Engineering, vol. 69, no. 2, pp. 403–422 (2020), DOI: 10.24425/aee.2020.133034.
[6] Mahdi A.S., Abdul-Malek Z., Arshad R.N., *SF_6 decomposed component analysis for partial discharge diagnosis in GIS: a review*, IEEE Access, vol. 10, pp. 27270–27288 (2022), DOI: 10.1109/ACCESS.2022.3156926.
[7] Sawicki A., *Arc models for simulating processes in circuits with a SF_6 circuit breaker*, Archives of Electrical Engineering, vol. 68, no. 1, pp. 147–159 (2019), DOI: 10.24425/aee.2019.125986.
[8] Islam K., Kim D., Abu-Siada A., *A review on adaptive power system protection schemes for future smart and micro grids, challenges and opportunities*, Electric Power Systems Research, vol. 230, 110241 (2024), DOI: 10.1016/j.epsr.2024.110241.
[9] Zhang Z., Si X., Hu C., Lei Y., *Degradation data analysis and remaining useful life estimation: a review on Wiener-process-based methods*, European Journal of Operational Research, vol. 271, no. 3, pp. 775–796 (2018), DOI: 10.1016/j.ejor.2018.02.033.
[10] Si X.S., Wang W., Hu C.H., Zhou D.H., *Remaining useful life estimation – a review on the statistical data driven approaches*, European Journal of Operational Research, vol. 213, no. 1, pp. 1–14 (2011), DOI: 10.1016/j.ejor.2010.11.018.
[11] Blom H.A.P., Bar-Shalom Y., *The interacting multiple model algorithm for systems with Markovian switching coefficients*, IEEE Transactions on Automatic Control, vol. 33, no. 8, pp. 780–783 (1988), DOI: 10.1109/9.1299.
[12] Bar-Shalom Y., Li X.R., Kirubarajan T., *Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software*, John Wiley & Sons (2001).
[13] Pourbabaee B., Meskin N., Khorasani K., *Sensor fault detection, isolation, and identification using multiple-model-based hybrid Kalman filter for gas turbine engines*, IEEE Transactions on Control Systems Technology, vol. 24, no. 4, pp. 1184–1200 (2016), DOI: 10.1109/TCST.2015.2480003.
[14] Lee A.S., Wu Y., Gadsden S.A., AlShabi M., *Interacting multiple model estimators for fault detection in a magnetorheological damper*, Sensors, vol. 24, no. 1, 251 (2024), DOI: 10.3390/s24010251.
[15] Julier S.J., Uhlmann J.K., *Unscented filtering and nonlinear estimation*, Proceedings of the IEEE, vol. 92, no. 3, pp. 401–422 (2004), DOI: 10.1109/JPROC.2003.823141.
[16] Kozierski P., Lis M., Owczarkowski A., Horla D., *Particle filter in power system state estimation – bad measurement data and branch disconnection*, Archives of Electrical Engineering, vol. 64, no. 2, pp. 237–248 (2015), DOI: 10.1515/aee-2015-0020.

Early Access