Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira, Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria
© 2026 ESREL 2026 Organizers. *Published by* Research Publishing, Singapore.
doi: 10.3850/ESREL2026061419_esrel26-p29173-cd

Probabilistic Prognostics for Power Electronic Devices: A Comparative Study of Monte Carlo and Latin Hypercube Sampling for Uncertainty-Aware RUL Estimation and Cost-Optimal Maintenance

Hiteshree Suresh Sakhare

Subir Chowdhury School of Quality and Reliability, IIT Kharagpur, India. E-mail: hiteshreesakhare@kgpian.iitkgp.ac.in

Heeralal Gargama

Subir Chowdhury School of Quality and Reliability, IIT Kharagpur, India. E-mail: heeralalgargama@hijli.iitkgp.ac.in

Abstract

Prognostics and Health Management (PHM) of power electronic devices is essential for ensuring reliability and enabling cost-effective maintenance in safety-critical applications such as electric vehicles (EVs) and energy conversion systems. In insulated-gate bipolar transistors (IGBTs), collector–emitter voltage (V_{CE}) degradation serves as a critical health indicator, as its progressive increase precedes catastrophic failure. This paper presents a comprehensive probabilistic prognostics framework that integrates statistical degradation modeling, uncertainty-aware Remaining Useful Life (RUL) estimation, probability-of-failure (PoF) assessment, hazard rate analysis, and cost-optimal maintenance decision-making. The failure threshold is defined as the maximum observed degradation voltage obtained from accelerated multi-device experimental datasets. Degradation rates are modeled using a normal distribution, and uncertainty propagation is performed using both Monte Carlo (MC) sampling and Latin Hypercube Sampling (LHS) techniques.

A detailed comparative analysis evaluates convergence behavior, computational efficiency, and statistical consistency between MC (5000 samples) and LHS (500 samples). Results derived from experimental degradation data demonstrate a monotonic decrease in RUL as the voltage margin to failure reduces, with median RUL values and confidence intervals showing strong agreement between the two sampling strategies. Convergence studies reveal that LHS achieves statistical stability at significantly lower sample sizes, while computational timing indicates approximately 30–40% faster execution compared to MC sampling. Probability-of-failure analysis indicates rapid risk escalation beyond 60 seconds at V_{CE}=3.5V, and hazard rate estimation exhibits an increasing failure rate characteristic of wear-out mechanisms. Cost-optimal maintenance analysis further shows that preventive maintenance decisions are economically justified near critical degradation states and remain robust across varying failure cost ratios.

The results demonstrate that Latin Hypercube Sampling provides superior sampling efficiency without compromising predictive accuracy, making it particularly suitable for real-time prognostics and risk-informed maintenance decision support in power electronic systems.

Keywords: Prognostics and Health Management, Remaining Useful Life, Monte Carlo Simulation, Latin Hypercube Sampling, Power Electronics Reliability, Probability of Failure, Hazard Rate, Maintenance Optimization, IGBT.

1. Introduction

Power electronic devices, particularly Insulated Gate Bipolar Transistors (IGBTs), serve as critical components in electric vehicles (EVs), renewable energy systems, aerospace converters, and industrial motor drives. Their operational reliability directly affects system availability, safety, and economic performance. Among failure-sensitive parameters, the collector–emitter voltage (V_{CE}) has been widely recognized as a reliable degradation indicator, as its gradual increase reflects bond-wire fatigue, solder degradation, and thermo-mechanical stress accumulation (Abuelnaga et. al. 2021, Agirrezabala et. al. 2025, Yang et. al. 2010).

The growing demand for high-power-density converters in EV applications has intensified thermal loading and cycling stress, accelerating wear-out mechanisms in IGBTs (Ciappa et. al. 2002). Traditional reliability assessment approaches rely on lifetime models such as Arrhenius-based acceleration, Coffin–Manson fatigue modeling, and Weibull statistical characterization (Ciappa et. al. 2000, Wang et. al. 2020). While these methods are effective for population-level reliability estimation, they do not provide real-time health assessment or individualized Remaining Useful Life (RUL) prediction.

Prognostics and Health Management (PHM) has emerged as a promising solution for condition-based maintenance in power electronics (Pecht et. al. 2012, Subedi et. al. 2025). Data-driven and hybrid prognostic frameworks estimate future degradation trends using monitored parameters such as on-state voltage, junction temperature, or switching characteristics (Thekemuriyil et. al. 2023). However, uncertainty propagation remains a critical challenge in RUL estimation. Degradation rates are inherently stochastic due to manufacturing variability, operating conditions, and measurement noise. Deterministic RUL calculations therefore fail to capture the uncertainty envelope necessary for risk-informed decision-making (Pan et. al. 2024).

Probabilistic RUL estimation methods address this limitation by propagating degradation uncertainty through stochastic sampling or Bayesian inference (Gebraeel et. al. 2023, Lei et. al. 2017). Monte Carlo (MC) simulation is one of the most widely used uncertainty propagation techniques due to its asymptotic accuracy and conceptual simplicity. Nevertheless, MC sampling often requires a large number of samples to achieve stable convergence, particularly when estimating distribution tails or confidence intervals (Rubinstein et. al. 2016). This computational burden limits its suitability for real-time embedded prognostic applications.

Latin Hypercube Sampling (LHS), a stratified sampling technique, has been proposed as a variance-reduction alternative to standard Monte Carlo methods (Stein et. al. 1987, Aaltelli et. al. 2008). By ensuring uniform coverage of the probability space, LHS can achieve comparable statistical accuracy with significantly fewer samples. Although LHS has been extensively applied in structural reliability and uncertainty quantification, its comparative effectiveness in power electronic prognostics remains underexplored.

In addition to RUL estimation, maintenance decision-making requires risk quantification and economic optimization. Probability-of-Failure (PoF) metrics and hazard rate estimation provide dynamic risk indicators aligned with reliability theory (Billinton et. al. 1985). Translating prognostic outputs into actionable maintenance decisions requires integrating cost models that balance preventive maintenance cost against failure consequences (Lin et. al. 2004). For EV traction systems, unexpected IGBT failure can lead to high downtime cost, system-level damage, and safety concerns, making cost-optimal decision frameworks particularly relevant (Amin et. al. 2026)

Despite significant advancements in probabilistic prognostics, several limitations remain in the context of power electronic devices. Existing studies primarily focus on RUL estimation without providing a unified framework that integrates uncertainty propagation, failure risk quantification, and cost-based maintenance decision-making. Furthermore, while Monte Carlo sampling is widely used, its high computational demand limits its applicability in real-time systems. Although Latin Hypercube Sampling has demonstrated efficiency in other domains, its effectiveness for uncertainty-aware RUL estimation and reliability analysis in power electronic devices has not been systematically investigated. In addition, limited work exists on linking probabilistic RUL outputs with actionable maintenance strategies under varying economic conditions, particularly for electric vehicle applications.

To address these gaps, this study proposes a comprehensive probabilistic prognostics framework that integrates degradation modeling, uncertainty-aware RUL estimation, failure risk analysis, and cost-optimal maintenance decision-making within a single unified pipeline. The novelty of this work lies in the systematic comparison of Monte Carlo and Latin Hypercube Sampling for uncertainty propagation in power electronic prognostics, along with the incorporation of Probability-of-Failure and hazard rate metrics for dynamic risk assessment. Additionally, the framework establishes a direct link between degradation state and economically optimal maintenance decisions, enabling a complete reliability-to-decision methodology. This integrated approach provides both theoretical and practical insights, demonstrating the suitability of Latin Hypercube Sampling for real-time and computationally efficient prognostic applications.

This paper presents a probabilistic prognostics framework for IGBTs based on experimentally measured degradation trajectories of collector–emitter voltage. The framework estimates RUL by propagating degradation rate uncertainty and performs a systematic comparison between MC simulation and LHS in terms of accuracy, convergence behavior, computational efficiency, and uncertainty characterization. In addition, the framework incorporates failure risk metrics, including PoF and hazard rate, and integrates a cost-optimal maintenance decision model, thereby establishing a complete reliability-to-decision pipeline for power electronic systems.

The results demonstrate that LHS achieves comparable statistical accuracy to MC simulation while significantly reducing computational burden, making it a practical and efficient approach for real-time prognostics and maintenance decision support.

The remainder of the paper is structured as follows: Section 2 describes the proposed methodology and probabilistic framework, Section 3 evaluates convergence and computational efficiency, Section 4 presents failure risk analysis, Section 5 discusses cost-optimal maintenance decisions, and Section 6 concludes the study.

2. METHODOLOGY

2.1 Overall Prognostics Framework

This study develops and evaluates a probabilistic prognostics framework for power electronic devices using experimentally obtained degradation data. The primary objective is to estimate RUL under uncertainty and to comparatively assess the effectiveness of MC simulation and LHS for uncertainty propagation.

The proposed prognostics framework establishes a systematic and uncertainty-aware pipeline for estimating the RUL of power electronic devices using V_{CE} degradation as the primary health indicator. The prognostics framework consists of multiple interconnected stages, as illustrated in the methodology flowchart in Fig. 1. The framework integrates signal preprocessing, statistical degradation modeling, probabilistic uncertainty propagation, reliability assessment, and cost-based decision optimization into a unified structure implemented in MATLAB.

The process begins with acquisition of time-series V_{CE} measurements from multiple devices undergoing accelerated degradation. These raw signals are preprocessed to suppress noise and reduce dimensionality. The degradation rate is then computed and statistically modeled to capture device-to-device variability. Using the probabilistic degradation rate distribution, RUL is estimated via MC and LHS methods to quantify uncertainty. The framework subsequently evaluates PoF, hazard rate behavior, and expected maintenance cost to support risk-informed decision-making.

By linking physics-based degradation behavior with probabilistic reliability modeling and economic optimization, the framework provides a comprehensive PHM solution suitable for real-time and embedded applications.

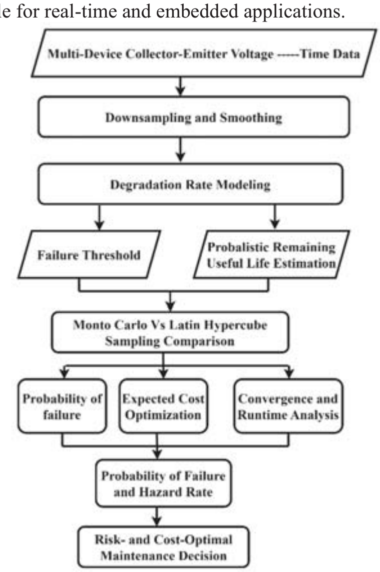

Fig. 1. Overall probabilistic prognostics and maintenance decision framework.

2.2 Data Acquisition and Preprocessing

The analysis utilizes time-series V_{CE} measurements collected from four IGBT devices subjected to accelerated degradation conditions. Each dataset consists of timestamped voltage readings recorded at high sampling frequency. Raw measurements typically contain high-frequency noise, transient spikes, and measurement artifacts that can distort derivative computation and degrade predictive accuracy as shown in fig. 2

To address this issue, the data are first downsampled by a factor of 100, significantly reducing computational burden while preserving long-term degradation trends, since the degradation process evolves on a much slower time scale than the sampling frequency. Following downsampling, a moving-average filter with a 20-second smoothing window is applied to suppress residual noise. The window length is adaptively determined based on the effective time resolution after downsampling. The processed signal is shown in Fig. 3.

$$W = \max\left(5, round\left(\frac{T_S}{\Delta t}\right)\right) \quad (1)$$

This preprocessing stage ensures that the smoothed V_{CE} signal retains its monotonic degradation behavior while eliminating fluctuations that could lead to unstable degradation rate estimation. Accurate preprocessing is critical because the RUL computation depends directly on the derivative of the voltage trajectory.

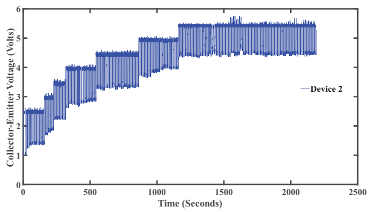

Fig. 2. Raw VCE degradation trajectories from accelerated testing.

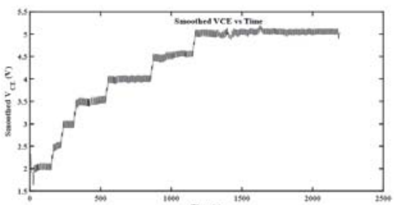

Fig. 3. Smoothed VCE trajectory after preprocessing.

2.3 Degradation Rate Modeling

The degradation rate represents the temporal progression of device wear and forms the statistical foundation of the RUL estimation process.

The degradation rate is computed as:

$$\frac{dV_{CE}}{dt} = \frac{V_{CE}(t+\Delta t) - V_{CE}(t)}{\Delta t} \quad (2)$$

Rates from all devices were pooled and filtered (95th percentile cutoff). The distribution is modeled as:

$$\frac{dV_{CE}}{dt} \sim \mathcal{N}(\mu_{rate}, \sigma_{rate}^2) \quad (3)$$

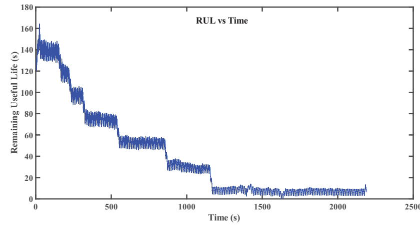

Fig. 4. Statistical distribution of pooled degradation rates.

It is computed using first-order finite differences of the smoothed V_{CE} signal. To obtain a robust statistical representation, degradation rates from all devices are aggregated, thereby capturing population-level variability rather than single-device behavior. Extreme values beyond the 95th percentile is removed to eliminate transient disturbances and measurement outliers. The filtered degradation rates are modeled using a normal probability distribution characterized by mean (μ_{rate}) and standard deviation (σ_{rate}). The mean degradation rate reflects the average wear progression, while the relatively large standard deviation indicates substantial stochastic variability among devices and operating conditions. The distribution is shown in Fig. 4. This probabilistic modeling approach allows uncertainty to be explicitly incorporated into RUL estimation, rather than relying on deterministic extrapolation. The statistical degradation model thus serves as the primary uncertainty source in subsequent prognostics calculations.

2.4 Failure Threshold Definition

A well-defined failure criterion is essential for consistent and reproducible Remaining Useful Life (RUL) estimation. In this study, the failure threshold is defined as the maximum observed V_{CE} during the degradation experiment:

$$V_{fail} = \max(V_{CE}) \quad (4)$$

where V_{CE} represents the smoothed collector–emitter voltage obtained after preprocessing.

As illustrated in Fig. 3, the degradation trajectory exhibits a progressive increase in V_{CE} over operational time due to aging mechanisms such as bond-wire fatigue, metallization degradation, and thermo-mechanical stress. The maximum recorded voltage at the end of the experiment corresponds to the device's end-of-life (EOL) condition and is therefore selected as the failure threshold.

Defining failure as the maximum experimentally observed voltage provides a physically grounded and data-driven criterion. Unlike arbitrary threshold selection, this approach ensures that the failure boundary reflects actual degradation behavior under accelerated stress conditions. In practical industrial applications, such thresholds may alternatively be defined using manufacturer specifications, thermal safety margins, or reliability qualification limits. However, for experimental prognostics validation, the maximum observed voltage offers a consistent and reproducible reference.

The failure threshold directly determines the degradation margin used in RUL computation. The remaining voltage margin at time t is expressed as:

$$\Delta V = V_{fail} - V_{CE}(t) \quad (5)$$

This margin represents the available degradation headroom before failure occurs. As the device ages and $V_{CE}(t)$ approaches V_{fail} the voltage margin decreases, leading to a corresponding reduction in predicted RUL. Therefore, the failure threshold serves as the fundamental reference point linking physical degradation progression to probabilistic lifetime estimation.

By anchoring RUL calculations to a clearly defined and experimentally validated failure boundary, the proposed framework ensures consistency, transparency, and reproducibility in prognostic modeling.

2.5. Remaining Useful Life Estimation

Remaining Useful Life (RUL) represents the predicted time remaining before the device reaches its predefined failure threshold under current degradation conditions. For a device operating at collector–emitter voltage $V_C(t)$, the remaining degradation margin is defined as:

$$\Delta V = V_{fail} - V_{CE}(t) \quad (6)$$

where V_{fail} denotes the failure threshold voltage and $V_C(t)$ represents the current degradation state.

Because the degradation rate is uncertain and modeled probabilistically, RUL cannot be estimated deterministically. Instead, uncertainty propagation is performed by sampling degradation rates from the statistically modeled distribution. For each sampled degradation rate r_i, the corresponding RUL sample is computed as:

$$RUL_i = \frac{\Delta V}{r_i} \quad (7)$$

This formulation directly links physical degradation progression to lifetime prediction. As the voltage margin ΔV decreases or degradation rate increases, the predicted RUL correspondingly decreases. Uncertainty in r_i propagates nonlinearly into RUL uncertainty.

To generate the empirical RUL distribution:

- Monte Carlo sampling is performed with N=5000 samples.
- Latin Hypercube Sampling is performed with N=500 samples.

Negative or near-zero degradation rate samples, which are physically unrealistic due to measurement noise or statistical tails, are corrected to ensure numerical stability and maintain positive lifetime predictions.

Since the RUL distribution exhibits right-skewed characteristics, the median is used as the primary central estimate rather than the mean. Confidence intervals are computed using percentile bounds (5%–95%) to quantify predictive uncertainty.

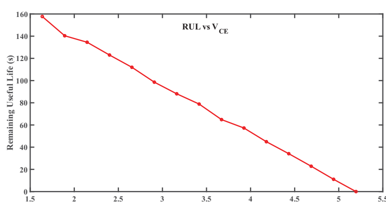

Fig. 5. Remaining Useful Life as a function of collector–emitter voltage (V_{CE}).

Fig.5 illustrates the relationship between median RUL and the degradation state (V_{CE}).

The figure shows a clear monotonic decreasing trend:

- At lower V_{CE} values (healthy condition), RUL is high due to large voltage margin.
- As V_{CE} increases, the remaining voltage margin ΔV decreases.
- Near the failure threshold, RUL rapidly approaches zero.

This behavior confirms the inverse dependence of RUL on degradation state:

$$RUL \propto \frac{V_{fail} - V_{CE}}{r} \quad (8)$$

The nearly linear decline validates the physical consistency of the proposed prognostics framework. It demonstrates that the model accurately translates degradation measurements into remaining lifetime predictions.

Furthermore, the rapid drop in RUL near high V_{CE} values indicates that small additional degradation near end-of-life leads to significant reduction in remaining time, which is characteristic of wear-out behavior in power electronic devices.

2.6 Monte Carlo vs. Latin Hypercube Sampling Comparison

In the proposed framework, the uncertainty is modeled through the degradation rate r, which is treated as a random variable following a normal distribution $r \sim \mathcal{N}(\mu_{rate}, \sigma_{rate}^2)$, derived from experimental data. For each sampled degradation rate r_i, the corresponding RUL sample is computed using the relation $RUL_i = \frac{\Delta V}{r_i}$, where ΔV represents the remaining voltage margin to failure threshold. Thus, uncertainty is propagated from the degradation rate domain to the RUL domain through stochastic sampling. It is important to note that the sampling is performed on the degradation rate parameter, not directly on voltage or RUL, making it the primary source of uncertainty in the prognostic model. Both MC and LHS generate realizations of r_i, which are subsequently transformed into RUL samples to construct the empirical RUL distribution. This formulation enables the transformation of uncertainty in degradation dynamics into a probabilistic distribution of RUL, which forms the basis for subsequent reliability and decision-making analysis.

MC simulation is employed as a baseline uncertainty propagation technique, where independent random samples are drawn from the assumed degradation rate distribution. These samples are propagated through the RUL model to obtain a distribution of RUL values. MC simulation is widely used due to its simplicity and asymptotic accuracy; however, it typically requires a large number of samples to achieve stable statistical estimates.

To improve sampling efficiency, LHS is also implemented. LHS is a stratified sampling method that divides the cumulative probability distribution into equal intervals and ensures that each interval is sampled exactly once. This approach provides more uniform coverage of the input space and reduces variance in the estimated outputs compared to standard random sampling.

Both MC and LHS methods are applied to generate probabilistic RUL distributions, which are subsequently used for uncertainty quantification, confidence interval estimation, and reliability analysis. The comparative evaluation between MC and LHS is performed in terms of statistical consistency, convergence behavior, and computational efficiency, enabling assessment of their suitability for real-time prognostics applications.

In this implementation:

- Monte Carlo sampling uses 5000 samples.
- LHS using only 500 samples.

The comparative performance of MC and LHS is evaluated using confidence interval analysis, convergence behavior, median RUL estimates, and computational efficiency.

2.6.1. Confidence Interval Comparison

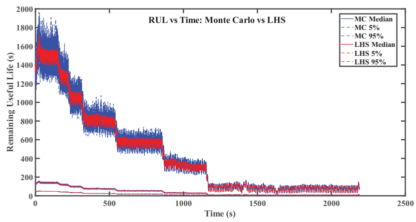

Fig. 6. MC vs LHS RUL confidence interval comparison across time.

Fig. 6 illustrates the 5%, 50% (median), and 95% confidence bounds of RUL estimated using both MC and LHS methods as a function of operational time.
The figure demonstrates:

- Nearly identical median RUL trajectories for both sampling methods.
- Overlapping 5% and 95% confidence intervals across the degradation timeline.
- Slightly smoother CI boundaries for LHS due to variance reduction.

The overlapping uncertainty bands confirm statistical equivalence of both approaches in capturing degradation uncertainty. However, LHS achieves this performance using significantly fewer samples, indicating superior sampling efficiency.

2.6.2. RUL–V_{CE} Relationship

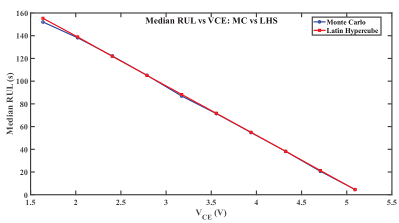

Fig. 7. Median RUL versus V_{CE} for Monte Carlo and LHS sampling.

This figure shows the relationship between the median RUL and collector–emitter voltage (V_{CE}), representing degradation progression.
The curve exhibits a monotonic decreasing behavior:

- At lower V_{CE} values (healthy state), RUL is large.
- As V_{CE} increases toward the failure threshold, RUL decreases sharply.
- Near failure, RUL approaches zero.

Both MC and LHS produce nearly identical median curves, confirming consistency of the two sampling approaches. The nonlinear shape of the curve reflects the inverse relationship between voltage margin ΔV and RUL, as defined by:

$$RUL = \frac{V_{fail} - V_{CE}}{r} \tag{9}$$

This figure validates physical consistency of the prognostics model.

3. Convergence and Computational Efficiency Analysis

To evaluate the statistical efficiency and practical feasibility of the sampling methods, convergence behavior and computational runtime were analyzed for both MC and LHS.

3.1. Convergence Analysis

The median RUL was computed for increasing sample sizes ranging from 50 to 5000 samples. The stabilization behavior of the median estimate provides insight into sampling efficiency.

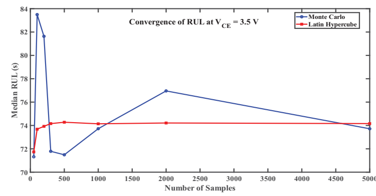

Figure 8. Convergence of median RUL estimate with increasing sample size for MC and LHS.

Fig. 8 illustrates the convergence characteristics of MC and LHS methods at a representative health state (e.g., V_{CE} = 3.5 V). The MC approach exhibits gradual convergence, with noticeable fluctuations persisting until approximately 1000–2000 samples. In contrast, LHS achieves stable median RUL estimates at significantly lower sample sizes (approximately 200–300 samples). This behavior demonstrates the variance reduction capability of stratified sampling. The faster stabilization of LHS confirms its statistical efficiency and suitability for computationally constrained applications.

3.2. Computational Time Comparison

In addition to convergence, computational runtime was evaluated for identical sample sizes to quantify practical performance differences.

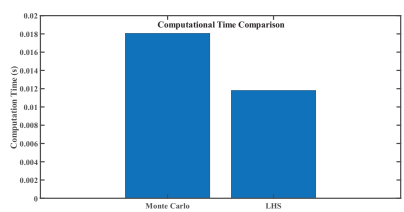

Figure 9. Average computational time comparison between Monte Carlo and Latin Hypercube Sampling.

Fig. 9 presents the average execution time measured over multiple trials. The LHS method demonstrates approximately 30–40% lower runtime compared to MC sampling. Although the absolute difference in execution time may appear small for a single evaluation, the cumulative computational savings become significant in real-time prognostics systems where RUL updates are performed repeatedly. The combined advantages of faster convergence and reduced runtime make LHS a more practical choice for embedded prognostics applications in electric vehicles and industrial power converters.

4. FAILURE RISK ANALYSIS

Reliable prognostics must extend beyond point RUL estimates and provide quantitative risk metrics that characterize the likelihood and urgency of failure. In this study, failure risk is evaluated using two complementary reliability measures: (i) probability of failure within a prediction horizon and (ii) hazard rate estimation. These metrics transform uncertainty-aware RUL distributions into actionable reliability indicators.

4.1. Probability of Failure

The Probability of Failure (PoF) quantifies the likelihood that the device will reach the failure threshold within a specified future time window Δt. Based on the sampled RUL distribution, PoF is defined as:

$$PoF(\Delta t) = P(RUL \leq \Delta t) \qquad (10)$$

In practice, this is computed as the proportion of RUL samples that are less than or equal to the prediction horizon Δt. Since RUL is modeled probabilistically using Monte Carlo or Latin Hypercube Sampling, PoF is derived directly from the empirical cumulative distribution function (CDF) of the RUL samples.

This metric converts remaining life uncertainty into a risk-based decision variable. Instead of asking "How much life remains?", PoF answers "What is the probability of failure within the next Δt seconds?"

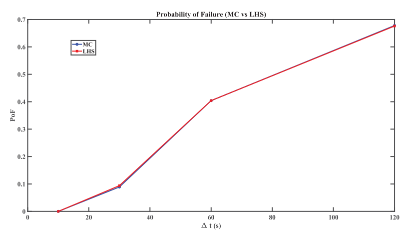

Fig. 10. Probability of failure versus prediction horizon at V_{CE} = 3.5 V.

Fig.10 illustrates the variation of PoF as a function of increasing prediction horizon Δt

The curve exhibits a monotonically increasing trend, reflecting the cumulative nature of failure probability.

- At Δt=10 s, PoF is approximately zero, indicating negligible short-term failure risk.
- At Δt=30 s, PoF increases to around 0.1, suggesting the onset of measurable risk.
- A sharp rise is observed between 30 s and 60 s, where PoF reaches approximately 0.4, indicating a critical transition region.
- At Δt=120 s, PoF approaches 0.7, reflecting high likelihood of failure within this horizon.
- The steep increase region corresponds to the wear-out phase of the device.
- MC and LHS curves overlap closely, demonstrating strong agreement.
- The overlap confirms that LHS achieves comparable accuracy to MC with significantly fewer samples.
- The figure highlights the effectiveness of PoF as a risk-based metric for prognostics and maintenance decision-making.

This trend corresponds directly to the cumulative distribution of RUL. The steep region of the curve represents the most critical operational window, where small increases in time lead to large increases in failure probability.

At the evaluated health state (V_{CE} = 3.5 V), the rapid increase beyond 60 seconds confirms that the device is operating in the wear-out region. This behavior aligns with the observed degradation trajectory approaching the defined failure threshold.

4.2 Hazard Rate Estimation

While PoF provides cumulative risk over a time window, the hazard rate describes instantaneous failure intensity at time t, conditioned on survival up to that time.

The hazard rate is defined as:

$$\lambda(t) = \frac{f(t)}{1-F(t)} \qquad (11)$$

where:

- $f(t)$ is the probability density function (PDF) of RUL,
- $F(t)$ is the cumulative distribution function (CDF),
- $1 - F(t)$ is the survival function.

The hazard rate measures the instantaneous probability of failure per unit time given that failure has not yet occurred.

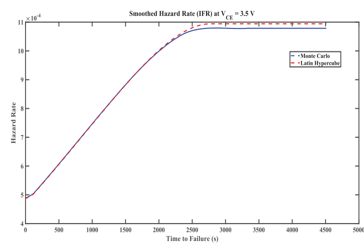

Fig. 11 Estimated hazard rate as a function of time to failure.

The fig. 11 shows the hazard rate estimated using MC and LHS.

- A monotonically increasing trend is observed, indicating a gradual rise in instantaneous failure probability over time.
- The hazard rate is low in the early life region and increases steadily as the device degrades.
- This behavior confirms an Increasing Failure Rate (IFR) characteristic, corresponding to the wear-out phase of the bathtub curve.
- The hazard curve stabilizes at higher time values, indicating proximity to end-of-life conditions with high failure likelihood.
- The MC and LHS curves closely overlap, demonstrating strong agreement and validating the accuracy of LHS with fewer samples.
- The increasing trend is consistent with classical reliability theory and reflects physical degradation mechanisms in IGBTs.

Wear-out mechanisms such as bond wire fatigue, thermal cycling damage, and metallization degradation contribute to the observed increasing failure rate. Overall, the figure confirms that the framework captures both statistical uncertainty and physical degradation behavior effectively.

5. COST-OPTIMAL MAINTENANCE DECISION

Maintenance planning in power electronic systems requires balancing the trade-off between preventive maintenance cost C_p and corrective failure cost C_f While preventive actions incur scheduled service expenses, unexpected failure leads to significantly higher economic consequences, including downtime, component replacement, safety risk, and potential collateral damage in electric vehicle (EV) drive systems.

The expected maintenance cost over a future decision horizon Δt is formulated as:

$$E[Cost(\Delta t)] = C_p(1 - PoF(\Delta t)) + C_f PoF(\Delta t) \tag{12}$$

Where, $PoF(\Delta t)$ is the probability that failure occurs within the prediction horizon

C_p is preventive maintenance cost,

C_f is corrective failure cost.

This formulation represents a weighted economic expectation. If the device survives beyond Δt, only preventive maintenance cost is incurred. However, if failure occurs before Δt, the higher corrective cost dominates.

The optimal maintenance decision is defined as:

$$\Delta t^* = \arg\min E\,[Cost(\Delta t)] \tag{13}$$

which identifies the decision horizon that minimizes total expected economic loss.

5.1. *Expected Maintenance Cost Curve*

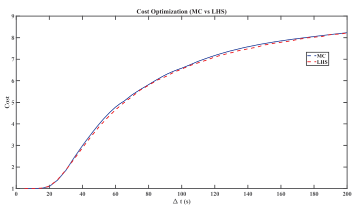

Fig. 12. Expected maintenance cost as a function of decision horizon

The fig. 12 shows the expected maintenance cost as a function of decision horizon Δt at the evaluated health state V_{CE} = 3.5V.

- The cost curve exhibits a monotonically increasing trend, indicating that delaying maintenance increases the overall expected cost.
- At small Δt, the probability of failure is low, and the cost is dominated by the preventive maintenance cost C_p.
- As Δt increases, the probability of failure rises, and the corrective cost component C_f · PoF becomes dominant.
- This leads to a rapid increase in expected cost with increasing decision horizon.
- The minimum expected cost occurs at the smallest evaluated Δt, indicating a boundary-optimal solution rather than an interior minimum.
- This behavior suggests that the device is operating close to the failure threshold, where any delay significantly increases failure risk.
- The MC and LHS curves closely overlap, confirming that LHS provides comparable accuracy with fewer samples.

5.2. *Cost Sensitivity Analysis for EV Applications*

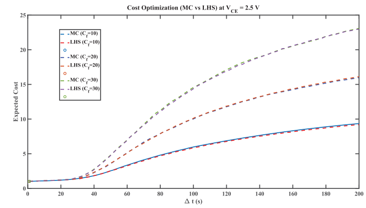

(a)

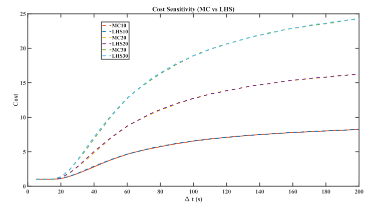

(b)

Fig. 13. Expected maintenance cost sensitivity to varying corrective failure cost ratios C_f for EV applications at (a) V_{CE} = 2.5 V (b) V_{CE} = 3.5 V.

To evaluate robustness of the maintenance decision, sensitivity analysis was performed by varying the corrective failure cost ratio:

$$C_f = \{10,20,30\} \tag{14}$$

while maintaining $C_p = 1$

The fig. 13 compare the expected maintenance cost behavior at two degradation states, name V_{CE} = 2.5V (early stage) and V_{CE} = 3.5V (advanced stage).

At V_{CE} =2.5 V (Early Degradation State)

- The cost curves exhibit a near U-shaped trend, indicating the presence of a true optimal maintenance time Δt∗.
- At small Δt, cost is dominated by preventive maintenance cost C_p.
- As Δt increases, cost initially decreases slightly and then increases due to the growing failure probability contribution C_f · PoF.
- A clear minimum point exists, representing the optimal balance between preventive and corrective costs.
- As the corrective cost C_f increases:
  - The optimal maintenance time shifts toward smaller Δt.
  - This indicates that higher failure penalties require earlier intervention.

At V_{CE} = 3.5V (Advanced Degradation State)

- The cost curves show a monotonically increasing trend, with no interior minimum.
- The minimum cost occurs at the smallest decision horizon, indicating a boundary-optimal solution.
- This suggests that the device is already close to the failure threshold, and delaying maintenance significantly increases expected cost.
- Increasing C_f makes the curve steeper, amplifying the economic penalty of delayed action, but does not change the optimal decision (immediate maintenance).

The transition from U-shaped cost behavior at V_{CE} = 2.5V to monotonic cost behavior V_{CE} = 3.5V reflects the progression from an optimal planning region to a critical failure region. At earlier stages of degradation, the system allows for flexible maintenance scheduling, where an optimal decision horizon can be identified by balancing preventive and corrective costs. In contrast, at later stages, the system approaches the failure threshold, and the cost increases continuously with delay, making immediate intervention the most effective strategy. This transition demonstrates the capability of the proposed framework to adapt maintenance decisions dynamically based on the real-time degradation state of the device.

The comparison between MC and LHS across both figures shows that the curves closely overlap, demonstrating strong statistical consistency and confirming that LHS achieves comparable accuracy to MC while requiring fewer samples, thereby offering higher computational efficiency. From a physical perspective, at lower V_{CE} values, the system operates in a healthy or mid-life state, where degradation is moderate and maintenance decisions can be flexibly optimized over a range of time horizons. In contrast, at higher V_{CE} values, the system transitions into the wear-out phase, where degradation accelerates and the failure risk increases rapidly, necessitating immediate intervention. These results highlight that the proposed probabilistic framework effectively captures the progression of degradation and translates it into economically optimal and actionable maintenance decisions.

6. RESULTS AND DISCUSSION

The degradation rate analysis yielded a mean rate of approximately 0.0001965 V/s and a standard deviation of 0.033903 V/s, indicating significant variability in the degradation process across devices. The RUL distributions at VCE = 3.5 V exhibit a right-skewed behavior, with median values in the range of 71–75 seconds, reflecting increased uncertainty near the failure threshold. MC and LHS methods produce statistically consistent RUL estimates; however, LHS demonstrates faster convergence and reduced computational cost. The probability of failure within a 60-second horizon exceeds 40%, indicating a high near-term failure risk at the evaluated degradation state. The estimated hazard rate shows an increasing trend, consistent with wear-out behavior in power electronic devices. Cost optimization results indicate that immediate maintenance minimizes the expected cost at the current health state, highlighting the critical nature of timely intervention. Overall, the results validate the effectiveness of the probabilistic framework and demonstrate the practical advantages of LHS for efficient and real-time prognostics applications.

7. CONCLUSION

This study presents a comprehensive probabilistic prognostics framework for power electronic devices based on V_{CE} degradation monitoring. The methodology integrates statistical degradation modeling, uncertainty-aware RUL estimation, failure risk quantification, hazard rate analysis, and cost-based maintenance optimization within a unified framework. Comparative evaluation demonstrates that LHS achieves comparable statistical accuracy to MC simulation while offering significantly improved computational efficiency. The proposed framework enables risk-informed and economically optimized maintenance decisions for electric vehicle power electronic systems.

Future work will focus on incorporating physics-based degradation models, environmental covariates, and online probabilistic model updating to further enhance predictive robustness and real-world applicability.

Acknowledgement

We are writing to express our heartfelt gratitude for the invaluable support and resources provided by the Subir Chowdhury School of Quality and Reliability at the Indian Institute of Technology Kharagpur during the course of our research.

References:

Abuelnaga, A., Narimani, M., & Bahman, A. S. (2021). A review on IGBT module failure modes and lifetime testing. IEEE access, 9, 9643-9663.

Agirrezabala, E., Aizpuru, I., Garrido, D., & Portillo, A. (2025). Reliability of power semiconductor modules: A state-of-the-art review. *IEEE Open Journal of Power Electronics*.

Yang, S., Xiang, D., Bryant, A., Mawby, P., Ran, L., & Tavner, P. (2010). Condition monitoring for device reliability in power electronic converters: A review. *IEEE transactions on power electronics*, *25*(11), 2734-2752.

Ciappa, M. (2002). Selected failure mechanisms of modern power modules. *Microelectronics reliability*, *42*(4-5), 653-667.

Ciappa, M., & Fichtner, W. (2000, April). Lifetime prediction of IGBT modules for traction applications. In *2000 IEEE international reliability physics symposium proceedings. 38th Annual (Cat. No. 00CH37059)* (pp. 210-216). IEEE.

Wang, H., & Blaabjerg, F. (2020). Power electronics reliability: State of the art and outlook. *IEEE Journal of Emerging and Selected Topics in Power Electronics*, *9*(6), 6476-6493.

Pecht, M. G. (2012). Prognostics and health management. In *Solid State Lighting Reliability: Components to Systems* (pp. 373-393). New York, NY: Springer New York.

Subedi, S., Gui, Y., & Xue, Y. (2025). Applications of data-driven dynamic modeling of power converters in power systems: An overview. *IEEE Transactions on Industry Applications*, *61*(2), 2434-2456.

Thekemuriyil, T., Rohner, J. D., & Minamisawa, R. A. (2023). Machine learning-based prediction of on-state voltage for real-time health monitoring of IGBT. *Power Electronic Devices and Components*, *6*, 100049.

Pan, J., Sun, B., Wu, Z., Yi, Z., Feng, Q., Ren, Y., & Wang, Z. (2024). Probabilistic remaining useful life prediction without lifetime labels: A Bayesian deep learning and stochastic process fusion method. *Reliability Engineering & System Safety*, *250*, 110313.

Gebraeel, N., Lei, Y., Li, N., Si, X., & Zio, E. (2023). Prognostics and remaining useful life prediction of machinery: advances, opportunities and challenges. *Journal of Dynamics, Monitoring and Diagnostics*, 1-12.

Lei, Y. (2017). Remaining useful life prediction. *Intelligent Fault Diagnosis and Remaining Useful Life Prediction of Rotating Machinery*, 281-358.

Rubinstein, R. Y., & Kroese, D. P. (2016). *Simulation and the Monte Carlo method*. John Wiley & Sons.

Stein, M. (1987). Large sample properties of simulations using Latin hypercube sampling. *Technometrics*, *29*(2), 143-151. Stein, M. (1987). Large sample properties of simulations using Latin hypercube sampling. *Technometrics*, *29*(2), 143-151.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., ... & Tarantola, S. (2008). *Global sensitivity analysis: the primer*. John Wiley & Sons.

Billinton, R., & Allan, R. N. (1985). *Reliability evaluation of engineering systems* (Vol. 198, No. 3). New York: Pitman.

Lin, D., Wiseman, M., Banjevic, D., & Jardine, A. K. (2004). An approach to signal processing and condition-based maintenance for gearboxes subject to tooth failure. *Mechanical Systems and Signal Processing*, *18*(5), 993-1007.

Amin, A. A., Irfan, M., Waseem, S., Mubarak, A., Alwadie, H., & Al Dawsari, S. (2026). Application of advanced fault-tolerant control techniques in the design of modern electric vehicles for increased reliability: a review. *Automatika*, *67*(1), 237-265.