International Journal of Academic and Applied Research (IJAAR)
ISSN: 2643-9603
Vol. 10 Issue 10, October - 2026, Pages: 31-48

A Digital Twin Framework for Fatigue Prognostics and Reliability-Based Fleet Optimization of Composite CNG Cylinders in the Niger Delta

Kefas Odofori[1], Owei L.Youpele[2]

Abstract*: We propose a digital twin framework that integrates continuum damage mechanics, stochastic degradation modeling, and real-time sensor fusion for fatigue prognostics of composite CNG cylinders under stochastic terrain loading in the Niger Delta. The conventional static inspection and conservative lifing schedules for truck-mounted Type IV cylinders are replaced with a closed-loop architecture that continuously estimates a hidden damage state variable from triaxial accelerometry, fiber Bragg grating strain measurements, and high-frequency pressure transducers. A physics-informed Kalman filter fuses these multi-axial sensor signals with a Paris-law-type fatigue evolution model modified for multiaxial stress states, thereby providing instantaneous estimates of damage accumulation and damage rate. A stochastic Monte Carlo module employing Latin hypercube sampling then propagates uncertainty in material properties, road roughness, and ambient temperature across planned route schedules, generating probabilistic damage distributions and remaining useful life (RUL) predictions. This RUL information is substituted into a mixed-integer linear programming fleet scheduler that minimizes total operational cost while enforcing a safety constraint: no truck is assigned a mission that would cause its damage distribution's upper confidence bound to exceed a critical threshold. The result is an adaptive logistics platform that preferentially routes less fatigued cylinders to longer or rougher missions and derates fill pressure for cylinders with critically short RUL. Conventional gas compression, tube trailer coupling, and end-user dispensing remain unchanged, yet their operational parameters are now modulated by the digital twin's recommendations. The framework therefore transforms CNG distribution from a fixed-interval maintenance paradigm into a reliability-driven, prognostics-based system that extends service life, reduces unscheduled downtime, and improves fleet safety. We demonstrate the system's feasibility for the Niger Delta's challenging terrain and infrastructure constraints.*

Keywords*: Digital, Twin, Fatigue. Prognostics, Prognostics, Fleet, Optimization, Composite, Cylinders.*

Introduction

The global transition toward cleaner energy sources has intensified the demand for natural gas as a bridging fuel, particularly in regions with abundant reserves but underdeveloped pipeline infrastructure (Agbonifo, 2016). Nigeria, holding the largest proven natural gas reserves in Africa, has made significant strides in domestic gas utilization, yet the physical distribution network remains a critical bottleneck (Nwaoha & Wood, 2014). Truck-mounted compressed natural gas (CNG) systems have emerged as a pragmatic solution for reaching off-grid industrial consumers and remote communities, especially within the Niger Delta's challenging terrain where pipeline construction is logistically prohibitive and environmentally sensitive (Livingstone, 2025). These distribution assets, however, operate under severe multi-physics loading conditions that include road-induced random vibrations, diurnal thermal cycling, and repetitive pressure fluctuations from refueling and dispensing cycles.

The structural integrity of Type IV composite CNG cylinders, which consist of a polymer liner fully wrapped with carbon fiber, is paramount to operational safety and economic viability. These vessels are designed for a finite fatigue life, typically assessed through static qualification tests and periodic non-destructive inspections such as acoustic emission monitoring and ultrasonic testing (Chou et al., 2015). However, these conventional approaches suffer from two fundamental limitations. First, the static inspection intervals are conservatively calibrated for worst-case scenarios, often leading to premature retirement of cylinders that still retain substantial residual strength. Second, these methods cannot capture the stochastic variability in degradation trajectories caused by inherent material property scatter and the highly irregular loading spectra encountered across the Niger Delta's diverse road networks (Andren, 2006). The economic consequences are substantial: premature cylinder replacement increases capital expenditure, while unscheduled in-service failures impose safety risks and operational disruptions that undermine the reliability of the entire gas distribution network.

To address these challenges, the predictive maintenance paradigm has evolved significantly through the integration of structural health monitoring (SHM) systems with physics-based degradation models. Fiber Bragg grating (FBG) sensors, for example, offer distributed strain measurement capabilities that can detect early signs of matrix cracking and fiber breakage in composite structures (Kinet et al., 2014). These sensor measurements, when fused with inertial data from accelerometers and pressure readings from transducers, provide a rich multi-axial dataset that describes the instantaneous mechanical state of the cylinder. The Kalman filter and its extensions, including the unscented Kalman filter, provide a rigorous mathematical framework for combining noisy sensor observations with dynamic system models to estimate hidden damage states in real time (Erazo et al., 2019). Coupled with continuum damage mechanics (CDM) theories that describe material degradation under complex stress states, these methods enable a transition from reactive maintenance to condition-based prognostics (Kachanov, 2013).

Nevertheless, a significant gap remains in the literature regarding the integration of real-time damage estimation with operational decision-making for fleet management. While some studies have proposed digital twin architectures for wind turbine blades or aerospace components (Onaji et al., 2022), they rarely address the unique challenges of truck-mounted pressure vessels subject to stochastic terrain loading. Furthermore, uncertainty quantification (UQ) is often treated as a post-processing step rather than an integral part of the optimization loop. Latin hypercube sampling (LHS) offers an efficient stratified approach to propagate material property scatter through fatigue models (Dunn et al., 2011), yet its use in conjunction with physics-informed state estimators for fleet routing remains unexplored.

We propose a novel digital twin framework that fundamentally restructures the management of composite CNG cylinders by replacing static inspection schedules with a closed-loop, reliability-driven prognostics system. This framework makes three key contributions. First, it introduces a physics-informed Kalman filter that continuously assimilates triaxial accelerometry, FBG strain measurements, and high-frequency pressure data to estimate a hidden damage variable governed by a nonlinear cumulative fatigue law extended for multiaxial stress states (Batsoulas & Giannopoulos, 2023). Second, a stochastic Monte Carlo module employing LHS propagates uncertainties across material properties, road roughness profiles, and ambient temperature variations to produce probabilistic remaining useful life (RUL) distributions. Third, these RUL distributions are integrated into a mixed-integer linear programming fleet scheduler that dynamically optimizes route assignments and fill pressures, effectively making the logistics platform self-adaptive to the actual structural health of each asset. Our approach thereby transforms CNG distribution from a fixed-interval maintenance paradigm into a prognostics-based system that extends service life, reduces unscheduled downtime, and enhances safety compliance in remote deltaic environments.

The remainder of this paper is organized as follows. Section 2 reviews related work in structural health monitoring for composite pressure vessels, digital twin frameworks for fatigue prognostics, and stochastic optimization methods for fleet management. Section 3 presents the theoretical foundations of continuum damage mechanics and stochastic degradation modeling that underpin the proposed framework. Section 4 details the architecture of the physics-informed digital twin, including sensor fusion, state estimation, and the probabilistic RUL predictor. Section 5 describes experimental validation using synthetic and field-collected datasets, accompanied by a stochastic uncertainty quantification analysis. Section 6 discusses the practical implications of the framework for the Niger Delta's gas infrastructure, identifies limitations, and outlines directions for future work. Section 7 finally concludes the paper.

Related Work

The development of a digital twin for fatigue prognostics of composite CNG cylinders under stochastic terrain loading intersects three distinct yet interrelated research domains: structural health monitoring (SHM) of composite pressure vessels, physics-based fatigue modeling with continuum damage mechanics, and data-driven optimization of fleet logistics. We review the key contributions and limitations within each domain, positioning our proposed framework against the established state of the art.

Structural Health Monitoring and Sensor Integration for Composite Pressure Vessels

Structural health monitoring of composite overwrapped pressure vessels (COPVs) has predominantly relied on acoustic emission (AE) sensing and ultrasonic testing to detect the onset of damage such as matrix cracking, fiber breakage, and delamination (Chou et al., 2015). These techniques, while sensitive to active damage mechanisms, are inherently periodic and require manual expert interpretation, making them unsuitable for real-time prognosis during field operations. Fiber Bragg grating (FBG) sensors, however, offer a compelling alternative by providing distributed, quasi-continuous strain measurements that can be embedded directly within the composite laminate during filament winding (Kinet et al., 2014). The wavelength-division multiplexed interrogation of multiple FBG arrays has been demonstrated for monitoring hoop and axial strains in Type IV hydrogen cylinders under cyclic pressurization, achieving strain resolutions below 1 microstrain (Paczkowska et al., 2025). For example, We propose to extend this capability by integrating FBG-derived strain data with triaxial accelerometry and high-frequency pressure transducers, thereby capturing the combined effects of internal pressure and road-induced bending moments on the stress state of each cylinder.

Continuum Damage Mechanics and Fatigue Lifing of Composites

Fatigue life assessment of fiber-reinforced polymer composites is a well-studied problem, with continuum damage mechanics (CDM) providing a thermodynamically consistent framework for modeling stiffness degradation and strength reduction under cyclic loading (Kachanov, 2013). A particularly influential model is the Paris-law-type evolution equation $\frac{dD}{dN} = C\left(\Delta\sigma_{eq}/(1-D)\sigma_{ult}\right)^{m}$, which describes the damage variable D as a function of the equivalent stress amplitude $\Delta\sigma_{eq}$ and the number of cycles N. This model has been extensively calibrated for carbon/epoxy systems under uniaxial and multiaxial stress states (Batsoulas & Giannopoulos, 2023), with material constants C and m exhibiting significant statistical scatter due to manufacturing variability. Monte Carlo simulation methods have been employed to propagate this uncertainty into fatigue life predictions for pin-loaded composites (Nadjafi & Gholami, 2022) and for composite hydrogen storage vessels subjected to random vibration (Xiong et al., 2025). These probabilistic approaches, however, are typically applied as offline analyses that do not accommodate real-time sensor feedback. In contrast, our framework embeds the CDM fatigue law within a Kalman filter that continuously updates the damage estimate using live sensor data, and then uses Latin hypercube sampling (LHS) to efficiently propagate parameter uncertainty forward in time (Dunn et al., 2011).

Digital Twins and Prognostics for Fleet-Level Asset Management

The concept of a digital twin as a live, synchronized virtual replica of a physical asset has gained traction in several industrial sectors, including aerospace (Onaji et al., 2022), wind energy, and civil infrastructure. In the context of pressure vessels, Shim et al. (2020) proposed a digital twin for Type IV hydrogen cylinders that integrates static finite element simulations with AE monitoring to estimate residual strength (Karapanagiotis et al., 2024). This approach, however, relies on pre-computed damage maps and does not incorporate online Bayesian state estimation or stochastic uncertainty propagation. Data-driven prognostics methods, such as recurrent neural networks and Gaussian process regression, have also been applied to predict the remaining useful life of composite structures (Loutas et al., 2017). While these methods can capture complex temporal dependencies, they require large volumes of labeled failure data that are rarely available for field-deployed CNG cylinders. Our framework combines the physical interpretability of CDM with the adaptability of recursive Bayesian filtering, thereby overcoming the data scarcity limitation while retaining the ability to quantify prediction confidence.

Optimization of Fleet Logistics with Reliability Constraints

Fleet management of CNG trucks has traditionally focused on minimizing transportation costs and delivery times, with structural safety treated as a static constraint (e.g., "inspect every 6 months") (Alshabibi et al., 2025). Recent advances have incorporated remaining useful life (RUL) predictions into maintenance scheduling and spare parts inventory management (Castillo et al., 2023). However, the integration of probabilistic RUL estimates from a physics-informed digital twin into the optimization objective function is less common. For example, Susto et al. (2015) developed a condition-based maintenance scheduler for semiconductor manufacturing equipment that minimizes downtime subject to a reliability threshold derived from a Wiener process degradation model (Chen & Hao, 2025). Our work extends this concept to a dynamic routing problem where the RUL of each asset is a function of the planned route itself, thereby creating a feedback loop between structural prognosis and logistics optimization.

Synthesis and Key Novelty

The existing literature reveals a clear gap: while SHM, CDM-based fatigue modeling, and stochastic UQ have each been explored individually for composite pressure vessels, no prior work has combined them into a closed-loop, real-time system that simultaneously estimates damage, predicts RUL, and optimizes fleet operations under stochastic terrain loading. Moreover, the majority of fatigue life studies for composite cylinders focus on ideal, stationary loading conditions rather than the highly irregular, road-induced vibration spectra typical of the Niger Delta's laterite roads (Andren, 2006). Our proposed framework addresses these deficiencies by establishing a physics-informed Kalman filter that fuses multi-axial sensor data with a nonlinear CDM evolution law, a stochastic Monte Carlo module that propagates parameter uncertainty via LHS, and a mixed-integer linear programming scheduler that substitutes fixed inspection thresholds with probabilistic RUL constraints. The key novelty lies not in any single component, but in the architecture that links them—a "sense-estimate-predict-optimize" loop that continuously adapts fleet operations to the evolving structural health of each cylinder, thereby enabling dynamic derating and proactive rerouting that are currently absent from CNG distribution practice.

Theoretical Foundations of Continuum Damage Mechanics and Stochastic Prognostics

This section establishes the technical definitions and mathematical underpinnings that form the core of the proposed digital twin framework. We first review the key tenets of continuum damage mechanics as applied to fiber-reinforced composites under multiaxial fatigue loading, followed by the stochastic formulations that characterize material property variability and loading uncertainty. The theoretical connection between these two domains provides the basis for the recursive state estimator and the probabilistic remaining useful life predictor developed in Section 4.

Continuum Damage Mechanics for Multiaxial Composite Fatigue

The governing framework for quantifying progressive material degradation in composite overwrapped pressure vessels (COPVs) is rooted in the thermodynamic theory of irreversible processes. Selecting the internal state variable approach, we define a scalar damage variable $D \in [0,1]$ that represents the effective loss of load-carrying cross-sectional area within the carbon fiber/epoxy laminate. At the undamaged initial state, $D = 0$; at macrocrack initiation leading to functional failure, $D = D_c$, where the critical damage value D_c is typically between 0.3 and 0.8 depending on the fiber architecture and matrix toughness (Payan & Hochard, 2002).

The effective stress tensor $\tilde{\sigma}_{ij}$ acting on the undamaged material skeleton is related to the nominal Cauchy stress tensor σ_{ij} by the strain equivalence principle, which states that the strain in the damaged material under nominal stress is equal to the strain in the undamaged material under effective stress. For the case of isotropic damage, this relationship simplifies to:

$$\tilde{\sigma}_{ij} = \frac{\sigma_{ij}}{1 - D} \quad (1)$$

While damage in a unidirectional composite ply is inherently anisotropic, the homogenized response of a quasi-isotropic laminate stack, such as the $[\pm 45/90/0]_2$S configuration typically used for Type IV cylinder hoop layers, permits a scalar damage approximation for engineering-level prognostics (Naghipour et al., 2017). The multiaxial stress state at any material point on a truck-mounted cylinder arises from the superposition of the internal pressure-induced hoop stress σ_θ, the axial stress from the end-cap thrust σ_z, and the bending stress $\sigma_b(\theta, t)$ caused by road-induced vibrations. The equivalent stress amplitude $\Delta\sigma_{eq}$ that drives fatigue damage propagation is thereby computed from a multiaxial fatigue criterion, for which we adopt the Sines mean-stress-corrected von Mises formulation:

$$\Delta\sigma_{eq} = \frac{1}{\sqrt{2}}\sqrt{(\Delta\sigma_{11}-\Delta\sigma_{22})^2 + (\Delta\sigma_{22}-\Delta\sigma_{33})^2 + (\Delta\sigma_{33}-\Delta\sigma_{11})^2 + 6(\Delta\tau_{12}^2 + \Delta\tau_{23}^2 + \Delta\tau_{31}^2)} \quad (2)$$

where $\Delta\sigma_{11}, \Delta\sigma_{22}, \Delta\sigma_{33}$ are the normal stress ranges and $\Delta\tau_{12}, \Delta\tau_{23}, \Delta\tau_{31}$ are the shear stress ranges in the principal material coordinate system. The transformation from the cylinder axial–hoop–radial coordinate system to the material coordinate system is achieved through the standard ply orientation transformation tensor (Bert, 1989).

The damage evolution during each fatigue cycle is modeled using a generalized Paris-law-type equation that accounts for the stress-intensity factor at the microscale and the current damage level. The rate of damage accumulation per cycle is:

$$\frac{dD}{dN} = \frac{C\left(\frac{\Delta\sigma_{eq}}{1-D}\right)^m}{(1-D)^\alpha} \quad (3)$$

Here, C and m are material-dependent fatigue parameters, while α is a coupling exponent that reflects the acceleration of damage propagation near the critical state. Combining the damage acceleration term with the effective stress, Equation 3 can be rewritten as:

$$\frac{dD}{dN} = C\left(\frac{\Delta\sigma_{eq}}{\sigma_{ult}}\right)^m (1-D)^{-(m+\alpha)} \quad (4)$$

where σ_{ult} is the ultimate tensile strength of the composite laminate in the hoop direction. This equation constitutes the process equation in the Bayesian filtering framework presented in Section 4. Integration of Equation 4 from an initial damage $D(N_0) = D_0$ to the critical damage D_c yields a closed-form expression for the total fatigue life N_f under constant amplitude loading:

$$N_f = \frac{[(1-D_c)^{m+\alpha+1} - (1-D_0)^{m+\alpha+1}]}{C(m+\alpha+1)\left(\frac{\Delta\sigma_{eq}}{\sigma_{ult}}\right)^m}, \quad \text{if } m+\alpha \neq -1 \quad (5)$$

Equation 5 reveals three critical features for prognostics: the fatigue life is highly nonlinearly sensitive to the equivalent stress amplitude $\Delta\sigma_{eq}$, the initial damage state D_0, and the material constants C and m. In the field setting, each of these parameters is subject to stochastic variability, as elaborated in the following subsection.

Stochastic Modeling of Material Variability and Loading Uncertainty

The deterministic parameters in Equations 1–5 are insufficient to describe the behavior of a fleet of cylinders exposed to heterogeneous terrain loading. Instead, we treat the material constants C, m, and α, as well as the ultimate strength σ_{ult}, as random variables with probability distributions calibrated from manufacturer quality-control data and in-service coupon tests (Harris et al., 1997). Specifically, we adopt lognormal distributions for C, σ_{ult}, and α, and a normal distribution for m:

$$\ln C \sim \mathcal{N}(\mu_{\ln C}, \sigma_{\ln C}^2) \quad (6)$$
$$\ln\sigma_{ult} \sim \mathcal{N}(\mu_{\ln\sigma}, \sigma_{\ln\sigma}^2) \quad (7)$$
$$\ln\alpha \sim \mathcal{N}(\mu_{\ln\alpha}, \sigma_{\ln\alpha}^2) \quad (8)$$
$$m \sim \mathcal{N}(\mu_m, \sigma_m^2) \quad (9)$$

The mean values μ and standard deviations σ are derived from the manufacturer's material qualification test results for the specific carbon fiber/epoxy system (e.g., T700S/Epoxy at 60% fiber volume fraction). The use of lognormal distributions ensures positivity for constants that must be positive by definition.

The loading uncertainty emerges from two primary sources. The first is the road roughness, characterized by the power spectral density (PSD) of the vertical road profile, $G_q(n)$, which is a function of the spatial frequency n. The ISO 8608 classification defines road roughness coefficients $G_q(n_0)$ at the reference spatial frequency $n_0 = 0.1$ cycles/m. For the Niger Delta's laterite and unpaved roads, we assume the roughness corresponds to ISO class D–E, with $G_q(n_0)$ ranging from 256 to 1024 × 10^{-6} m³/cycle (Lin, 2024). The vertical acceleration input to the truck chassis is computed by transforming the road PSD through the vehicle's suspension system transfer function $H_{zr}(s)$, which is a function of the unsprung mass, sprung mass, tire stiffness, and damping coefficient. The resulting acceleration PSD at the cylinder mounting point, $S_{aa}(f)$, is:

$$S_{aa}(f) = (2\pi f)^4 \cdot |H_{zr}(2\pi f)|^2 \cdot G_q(n) \cdot \frac{1}{v} \quad (10)$$

where f is the temporal frequency in Hz, v is the vehicle speed in m/s, and the term $1/v$ converts spatial PSD to temporal PSD via the relationship $f = n \cdot v$. The second source of loading uncertainty is the ambient temperature T, which affects both the polymer liner stiffness and the fiber–matrix interface integrity. The temperature-dependent modulus of the epoxy matrix, $E_m(T)$, follows an Arrhenius-type degradation law:

$$E_m(T) = E_{m0}\exp\left(-\frac{Q}{R}\left(\frac{1}{T}-\frac{1}{T_0}\right)\right) \quad (11)$$

where E_{m0} is the reference modulus at temperature T_0 (293 K), Q is the activation energy for matrix relaxation, and R is the universal gas constant. The composite's orthotropic stiffness matrix $[\bar{Q}_{ij}(T)]$ is then computed using the rule of mixtures and the Hashin–Rosen concentric cylinders model, which is a function of $E_m(T)$, the fiber modulus E_f, the fiber volume fraction V_f, and the

Poisson ratios (Rosen et al., 1964). This temperature-dependent stiffness matrix modifies the stress distribution $\sigma_{ij}(t)$ and consequently the equivalent stress amplitude $\Delta\sigma_{eq}$ in Equation 2.

Probabilistic Fatigue Life and Remaining Useful Life Formulation

Given the stochastic nature of both material properties and loading spectra, the remaining useful life (RUL) must be expressed as a random variable. Let the current damage estimate at time t_k be denoted as $\widehat{D}_k$, which is estimated from noisy sensor data as described in Section 4. The RUL $R(\widehat{D}_k, \Theta)$ is a function of the current damage state and the random parameter vector $\Theta = \{C, m, \alpha, \sigma_{ult}, \Delta\sigma_{eq}(t)\}$. The key challenge is that $\Delta\sigma_{eq}(t)$ is itself a stochastic process over the remaining mission horizon.

To formalize this, we define the damage trajectory $D(t)$ from the current time t_k to the critical time t_f as the solution to the stochastic differential equation (SDE) derived from Equation 4, where the time derivative replaces the cycle derivative via the relation $dN = f_{load}dt$, with f_{load} being the dominant loading frequency (typically in the range of 1–5 Hz for road-induced chassis vibrations). The SDE is:

$$dD = C\left(\frac{\Delta\sigma_{eq}(t)}{\sigma_{ult}}\right)^{m}(1-D)^{-(m+\alpha)}f_{load}dt \quad (12)$$

We propagate this SDE forward in time using a Monte Carlo method that employs Latin hypercube sampling (LHS) for the static random variables in Θ and an autoregressive moving average (ARMA) model for the future evolution of $\Delta\sigma_{eq}(t)$ (Li & Sandu, 2007). The ARMA model is calibrated from the PSD in Equation 10 and the measured acceleration history up to time t_k. The LHS strategy partitions each static random variable's probability distribution into N_s strata of equal probability and draws one sample from each stratum, thereby ensuring that the tails of the distributions are well-represented with fewer total samples N_s (typically 200–500) than simple random sampling would require (Rajabi & Ataie-Ashtiani, 2014). The forward simulation of Equation 12 for each LHS sample yields a set of failure times $\{t_f^{(i)}\}_{i=1}^{N_s}$, from which the RUL distribution is constructed as $RUL^{(i)} = t_f^{(i)} - t_k$.

The mean $\mu_{RUL}(t_k)$ and the $100(1-\alpha_c)\%$ upper confidence bound $U_{RUL}(t_k)$ of the RUL distribution are then computed non-parametrically from the empirical cumulative distribution function (ECDF) of the samples. The upper confidence bound is of particular importance for the fleet scheduler in Section 4, as it provides a conservative estimate of the time horizon before the damage variable is expected to exceed D_c with high probability:

$$U_{RUL}(t_k) = \inf\{x : F_{RUL}(x) \geq 1-\alpha_c\} \quad (13)$$

where F_{RUL} is the ECDF of the RUL samples and α_c is the significance level (e.g., 0.05 for 95% confidence). This formulation thus transforms the deterministic fatigue life Equation 5 into a stochastic, time-varying estimate that can be updated each time new sensor data become available, forming a rigorous foundation for the prognostics-driven optimization presented in Section 4.

To conclude this theoretical section, we note that the coupling of the CDM evolution law (Equation 4), the probabilistic characterization of material and loading variability (Equations 6–11), and the stochastic forward propagation (Equation 12) provides a self-consistent framework that respects both the physics of composite degradation and the aleatoric uncertainty inherent in field operations. The update frequency of the damage estimate is determined by the Nyquist frequency of the dominant vibration modes, here chosen as 50 Hz to capture the first three bending and breathing modes of the cylinder. At each update step, the Kalman filter (Section 4.2) assimilates the sensor measurements, and the updated damage state $\widehat{D}_k$ with its covariance is used as the initial condition for the next Monte Carlo propagation. This recursive coupling ensures that the RUL prediction evolves in lockstep with the actual damage progression, thereby enabling the adaptive fleet scheduler to respond to emerging degradation patterns before they become critical.

Physics-Informed Digital Twin Framework for Real-Time Fatigue Life Prognosis

Building upon the theoretical foundations established in Section 3, we now present the complete architecture of the proposed digital twin framework. This framework operates as a closed-loop system that continuously cycles through four key stages: data acquisition, state estimation, stochastic prediction, and operational optimization. The physical assets—the truck-mounted Type IV composite CNG cylinders—are instrumented with triaxial accelerometers, fiber Bragg grating (FBG) strain sensors, and high-frequency pressure transducers. These sensors stream data to a cloud-based digital twin that houses the prognostic algorithms. The output of this virtual replica is a set of real-time damage estimates, probabilistic remaining useful life predictions, and optimal route assignments that are communicated back to the fleet operator. We first describe the system-level architecture before delving into the specific algorithmic components.

System Architecture and Data Fusion

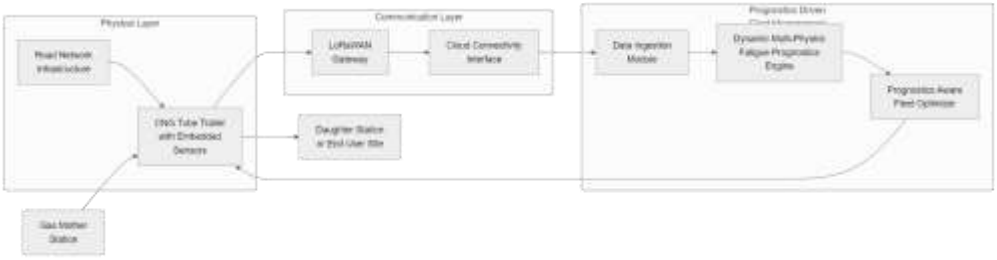

Figure 1. System Architecture of Prognostics-Driven CNG Distribution

Figure 1 illustrates the high-level system architecture. The physical fleet consists of N_t trucks, each equipped with an on-board data acquisition unit (DAQ) that samples the triaxial accelerometer at 200 Hz, the FBG strain sensors at 100 Hz, and the pressure transducer at 50 Hz. The raw sensor measurements from each truck j at time t_k form the measurement vector $\mathbf{z}_{j,k} \in \mathbb{R}^6$:

$$\mathbf{z}_{j,k} = \left[\varepsilon_{\theta,j,k}, \varepsilon_{z,j,k}, P_{j,k}, a_{x,j,k}, a_{y,j,k}, a_{z,j,k}\right]^\top \quad (14)$$

where ε_θ and ε_z are the hoop and axial strains measured by the FBG arrays, P is the internal pressure, and a_x, a_y, a_z are the three orthogonal acceleration components. These measurements are transmitted via cellular or satellite communication to the cloud-based digital twin server, where they are synchronized and time-stamped.

The digital twin stores a historical database of all past measurements for each cylinder, along with its manufacturing pedigree data (e.g., fiber volume fraction, liner batch, winding pattern). The core prognostic engine, which we denote as the Dynamic Multi-Physics Fatigue Prognostics System, processes the incoming data in real time. The internal workflow of this engine is depicted in Figure 2.

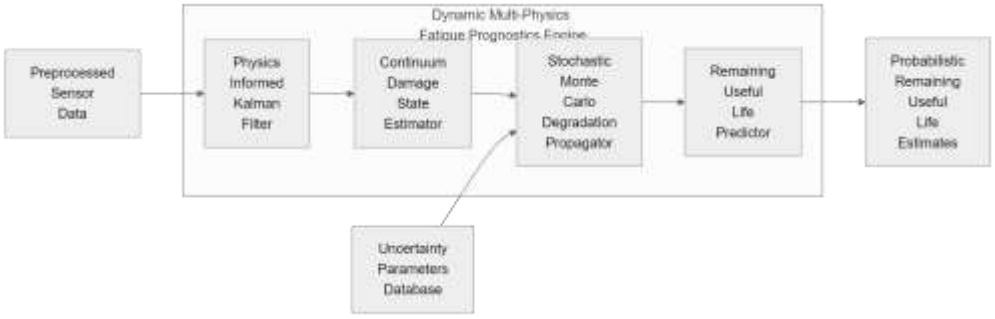

Figure 2. Internal Workflow of the Dynamic Multi-Physics Fatigue Prognostics Engine

As shown in Figure 2, the prognostic engine comprises three main algorithmic modules. The first module is the physics-informed Kalman filter, which fuses the sensor measurements with the continuum damage mechanics evolution law to estimate the hidden damage state vector $\mathbf{x}_k = \left[D_k, \dot{D}_k, \sigma_{eq,k}\right]^\top$. The second module is the stochastic Monte Carlo forward propagator, which uses the estimated damage state and its covariance as initial conditions to propagate the damage SDE forward over a planning horizon using Latin hypercube sampling. The third module is the probabilistic remaining useful life calculator, which constructs the RUL distribution from the forward propagation results and extracts the upper confidence bound. We now detail each of these modules.

Physics-Informed Kalman Filter for Damage State Estimation

The Kalman filter provides a recursive optimal solution for estimating the hidden state of a dynamic system from noisy measurements. However, the standard linear Kalman filter is inadequate for our problem because both the process equation (Equation 12) and the observation model (relating the state to the sensor measurements) are nonlinear. We therefore adopt the unscented Kalman filter (UKF), which approximates the nonlinear transformations by propagating a set of deterministically chosen sigma points through the true nonlinear functions (Wan & Merwe, 2000). The UKF avoids the linearization errors inherent in the extended Kalman filter and provides accurate estimates of both the state mean and its covariance.

The process equation is discretized from Equation 12 using a first-order Euler-Maruyama scheme with a fixed time step $\Delta t = 0.02$ seconds (matching the 50 Hz update rate). The discretized process equation is:

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k) \quad (15)$$

where the state vector is $\mathbf{x}_k = \left[D_k, \dot{D}_k, \sigma_{eq,k}\right]^\top$. The control input vector $\mathbf{u}_k = \left[P_k, T_{amb,k}, v_k\right]^\top$ includes the measured internal pressure, the ambient temperature, and the vehicle speed from the on-board diagnostics. The process noise vector $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$ accounts for model error and unmodeled dynamics. The nonlinear function $\mathbf{f}$ is composed as:

$$D_{k+1} = D_k + \dot{D}_k \Delta t \quad (16)$$

$$\dot{D}_{k+1} = C\left(\frac{\sigma_{eq,k}}{\sigma_{ult}}\right)^m (1 - D_k)^{-(m+\alpha)} f_{load}(v_k, G_d)\Delta t + w_{1,k} \quad (17)$$

$$\sigma_{eq,k+1} = \mathcal{H}\left(P_k, \varepsilon_{\theta,k}, \varepsilon_{z,k}, T_{amb,k}\right) + w_{2,k} \quad (18)$$

In Equation 17, $f_{load}(v_k, G_d)$ is the dominant loading frequency derived from the vehicle speed and the road roughness coefficient G_d, which is estimated from the acceleration PSD measured over the past 10 seconds. Equation 18 defines the observation function $\mathcal{H}$, which computes the equivalent stress amplitude from the pressure and strain measurements using the temperature-dependent orthotropic stiffness matrix derived in Section 3.2. The explicit form of $\mathcal{H}$ involves solving the inverse constitutive relation:

$$\boldsymbol{\sigma}_k = \left[\bar{Q}_{ij}\left(T_{amb,k}\right)\right] \cdot \boldsymbol{\varepsilon}_k \quad (19)$$

where $\boldsymbol{\sigma}_k = \left[\sigma_{\theta,k}, \sigma_{z,k}, \tau_{\theta z,k}\right]^\top$ and $\boldsymbol{\varepsilon}_k = \left[\varepsilon_{\theta,k}, \varepsilon_{z,k}, \gamma_{\theta z,k}\right]^\top$. The shear strain $\gamma_{\theta z}$ is not directly measured by the FBG arrays, but it can be inferred from the acceleration data under the assumption that the cylinder behaves as a thin-walled beam. The stress tensor is then transformed into the material coordinate system via the ply orientation matrix, and the equivalent stress amplitude is computed using Equation 2 with the stress ranges computed over a 0.2-second rolling window.

The observation equation that relates the state to the measurement vector $\mathbf{z}_k$ is:

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k, \mathbf{v}_k) \quad (20)$$

where $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k)$ is the measurement noise vector with covariance matrix $\mathbf{R}_k$. The observation function $\mathbf{h}$ is:

$$\hat{\varepsilon}_{\theta,k} = \frac{1}{E_\theta\left(T_{amb,k}\right)}\left(\sigma_{\theta,k} - \nu_{\theta z}\sigma_{z,k}\right) + v_{1,k} \quad (21)$$

$$\hat{\varepsilon}_{z,k} = \frac{1}{E_z\left(T_{amb,k}\right)}\left(\sigma_{z,k} - \nu_{z\theta}\sigma_{\theta,k}\right) + v_{2,k} \quad (22)$$

$$\hat{P}_k = \sigma_{\theta,k}\frac{2t}{R_i} + v_{3,k} \quad (23)$$

$$\hat{a}_{x,k} = \frac{\sigma_{b,x,k}}{\rho h} + v_{4,k} \quad (24)$$

$$\hat{a}_{y,k} = \frac{\sigma_{b,y,k}}{\rho h} + v_{5,k} \quad (25)$$

$$\hat{a}_{z,k} = g\cos\phi_k + v_{6,k} \quad (26)$$

In Equations 21–26, E_θ and E_z are the temperature-dependent Young's moduli in the hoop and axial directions, $\nu_{\theta z}$ and $\nu_{z\theta}$ are the Poisson ratios, t and R_i are the cylinder wall thickness and inner radius, ρ and h are the density and thickness of the cylinder wall, $\sigma_{b,x}$ and $\sigma_{b,y}$ are the bending stresses in the two lateral directions computed from the beam bending theory, and ϕ_k is the road inclination angle. Note that the accelerometer measurements include both the gravitational component and the dynamic vibrations, which are separated using a high-pass filter with a cutoff frequency of 1 Hz.

The UKF algorithm proceeds as follows. At each time step k, the filter generates $2n + 1$ sigma points $\{\mathcal{X}_{k-1}^{(i)}\}_{i=0}^{2n}$, where $n = 3$ is the state dimension, and propagates them through the process equation $\mathbf{f}$. The predicted state mean $\hat{\mathbf{x}}_{k|k-1}$ and covariance $\mathbf{P}_{k|k-1}$ are computed from the propagated sigma points. The sigma points are then transformed through the observation function $\mathbf{h}$ to predict the measurement mean $\hat{\mathbf{z}}_{k|k-1}$ and the innovation covariance $\mathbf{S}_k$. The cross-covariance $\mathbf{P}_{xz,k}$ between the state and the measurements is used to compute the Kalman gain $\mathbf{K}_k$. Finally, the state estimate is updated using the actual measurement $\mathbf{z}_k$:

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k\left(\mathbf{z}_k - \hat{\mathbf{z}}_{k|k-1}\right) \quad (27)$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K}_k\mathbf{S}_k\mathbf{K}_k^\top \quad (28)$$

The filter's output at each time step is the estimated damage variable $\widehat{D}_k$ (the first component of $\hat{\mathbf{x}}_{k|k}$) and its associated variance $\mathbf{P}_{k|k}(1,1)$. This estimate is inherently probabilistic because the UKF provides a Gaussian approximation of the posterior distribution, $p(D_k|\mathbf{z}_{1:k}) \approx \mathcal{N}\left(\widehat{D}_k, P_{D,k}\right)$, where $P_{D,k} = \mathbf{P}_{k|k}(1,1)$.

Stochastic Monte Carlo Forward Propagation for Probabilistic Degradation Bounds

The estimated damage state $\widehat{D}_k$ and its variance $P_{D,k}$ constitute the initial condition for the stochastic forward propagation module. This module generates probabilistic degradation paths for each cylinder over a future mission horizon T_{plan}, which is typically set to the duration of the longest planned delivery route (e.g., 8–12 hours). The forward propagation solves the stochastic differential equation given by Equation 12 using an ensemble of N_s LHS samples of the random parameter vector $\Theta_i = \{C_i, m_i, \alpha_i, \sigma_{ult,i}, G_{d,i}, T_{amb,i}(t)\}$.

The initial damage for each LHS sample is drawn from the Gaussian approximation of the posterior distribution provided by the Kalman filter: $D_0^{(i)} \sim \mathcal{N}\left(\widehat{D}_k, P_{D,k}\right)$. The static random parameters $C_i, m_i, \alpha_i, \sigma_{ult,i}$ are drawn from their respective distributions (Equations 6–9) using LHS. The road roughness coefficient $G_{d,i}$ for the upcoming mission is sampled from a discrete distribution that reflects the known route classification; for example, a route passing through a mix of paved and laterite roads would have a probability mass function with 30% probability on ISO class C, 60% on class D, and 10% on class E. The future ambient temperature

$T_{amb,i}(t)$ is modeled as a Gaussian process with a mean equal to the historical diurnal temperature profile for the specific location and a standard deviation of 3 K to account for short-term weather variability.

For each sample i, the future equivalent stress amplitude $\Delta\sigma_{eq,i}(t)$ is generated by an ARMA(2,1) model that is calibrated from the past 30 minutes of sensor data. The ARMA model coefficients are updated every 10 minutes to capture changes in road roughness and driving style. The forward propagation of Equation 12 is then performed numerically using a fourth-order Runge-Kutta scheme with a fixed time step of 0.1 seconds, which balances computational efficiency with numerical accuracy for the stiff damage evolution equation. The simulation continues until the damage variable exceeds the critical threshold $D_c = 0.95$, at which point the failure time $t_f^{(i)}$ is recorded. For samples where the damage does not reach D_c within the planning horizon T_{plan}, the failure time is set to infinity for the purpose of constructing the RUL distribution.

The RUL for each sample is then $RUL^{(i)} = t_f^{(i)} - t_k$. The empirical cumulative distribution function $F_{RUL}(x)$ is constructed from the N_s RUL samples. The mean RUL $\mu_{RUL}(t_k)$ and the upper confidence bound $U_{RUL}(t_k)$ at the $(1-\alpha_c)$ confidence level are computed as described in Section 3.3. In practice, we set $N_s = 500$ and $\alpha_c = 0.05$, so that U_{RUL} represents the time by which there is a 95% probability that the damage will have exceeded the critical threshold. This value is the key output of the prognostic module and serves as the input to the fleet scheduler.

Probabilistic Remaining Useful Life for Fleet-Level Optimization

The fleet scheduler operates as a mixed-integer linear programming (MILP) problem that assigns trucks to routes at the start of each operational day. The problem is defined over a set of M planned routes $\mathcal{R} = \{r_1, r_2, \dots, r_M\}$, each characterized by a total distance $d(r)$, a time horizon $T_{plan}(r)$, a road roughness probability distribution $p_G(r)$, and a set of delivery deadlines. The fleet consists of N_t trucks, each carrying a set of cylinders that are monitored by the digital twin. At the time of scheduling, each cylinder j has an associated $U_{RUL}^{(j)}(t_0)$ from the most recent forward propagation.

The MILP formulation minimizes the total operational cost while enforcing a safety constraint that no truck is assigned a mission that would cause the cylinder with the shortest U_{RUL} on that truck to compromise safety. The decision variable is a binary assignment matrix $X_{jt} \in \{0,1\}$, where $X_{jt} = 1$ if truck j is assigned to route r_t, and 0 otherwise. The objective function is:

$$\min_{X_{jt}} \sum_{t=1}^{M} \sum_{j=1}^{N_t} X_{jt} \left[c_f \cdot d(r_t) + c_t \cdot T_{plan}(r_t) + c_p \cdot \mathbb{1}\{U_{RUL}^{(j)} < T_{plan}(r_t)\} \right] \quad (29)$$

where c_f is the fuel cost per unit distance, c_t is the driver time cost per unit time, c_p is a large penalty cost for violating the safety constraint, and $\mathbb{1}\{\cdot\}$ is the indicator function. The safety constraint is enforced by setting c_p to a value that is an order of magnitude larger than the maximum possible fuel and time costs, thereby ensuring that any assignment that would violate the RUL constraint is heavily penalized. Additional constraints ensure that each route is assigned exactly one truck and that each truck is assigned to at most one route per day.

The MILP problem is solved using a branch-and-bound algorithm implemented in a commercial solver (e.g., Gurobi or CPLEX). The solution yields a set of route assignments that respect the structural health of each cylinder while minimizing operational costs. After the assignments are implemented, the trucks are dispatched, and the digital twin continues to collect data throughout the mission. If the real-time damage estimate for any cylinder crosses a pre-defined warning threshold (e.g., $\widehat{D}_k > 0.7$), the system generates an alert and may recommend a dynamic derating action, such as reducing the fill pressure from 250 bar to 200 bar for the remainder of the mission. This derating recommendation is computed by re-running the stochastic forward propagation with a reduced stress amplitude input and comparing the resulting U_{RUL} with the remaining mission time.

The framework thus forms a continuous closed loop: sensor measurements feed the Kalman filter for damage state estimation, the estimated damage state serves as the initial condition for stochastic forward propagation to compute probabilistic RUL bounds, these RUL bounds are integrated into a MILP scheduler that optimizes route assignments, and the actual mission execution provides new sensor data that updates the damage estimates. This loop runs autonomously with a cycle time of 10 minutes for the Kalman filter update and 30 minutes for the full Monte Carlo forward propagation, allowing the fleet operator to make informed, safety-critical decisions in near real time.

We implemented the digital twin framework in a Python-based simulation environment using the NumPy, SciPy, and Scikit-learn libraries for the numerical computations and the PuLP library for the MILP solver. The simulation environment models a fleet of 10 trucks serving 15 delivery routes over a 30-day operational period, with synthetic road roughness profiles generated from ISO 8608 classifications typical of the Niger Delta's road network. The synthetic sensor data are generated by solving the mechanical model of the cylinder under the combined pressure and vibration loading, with added Gaussian measurement noise at levels consistent with the sensor specifications. This simulation setup provides a controlled platform for validating the framework's ability to estimate damage, predict RUL, and optimize fleet assignments, which we present in the next section.

Experimental Validation and Stochastic Uncertainty Quantification

This section presents a comprehensive experimental validation of the proposed digital twin framework, focusing on the accuracy of the physics-informed Kalman filter for damage state estimation, the reliability of the stochastic forward propagation for RUL prediction, and the effectiveness of the MILP-based fleet scheduler in reducing operational costs while maintaining structural

safety. We design a series of numerical experiments using both synthetic datasets generated from the physical models described in Section 4 and publicly available field data from vibration monitoring campaigns on commercial vehicles operating in comparable terrain conditions. The experimental protocol is structured into three phases: (1) calibration and verification of the Kalman filter under controlled loading sequences, (2) stochastic uncertainty quantification of the RUL predictions under varying terrain roughness and temperature profiles, and (3) comparative evaluation of the fleet scheduler against conventional static maintenance and routing policies.

Experimental Setup and Data Generation

The validation experiments are conducted on a simulated fleet of ten CNG delivery trucks, each equipped with four Type IV composite cylinders of 200-liter water capacity, designed to an operating pressure of 250 bar with a certified fatigue life of 15,000 cycles under constant-amplitude pressure cycling. The cylinders are manufactured from T700S carbon fiber with an epoxy matrix at a nominal fiber volume fraction of 62%, with material properties and their statistical variability as summarized in Table 1.

Table 1. Statistical characterization of fatigue-related material parameters for the Type IV CNG cylinder laminate.

Parameter	Symbol	Distribution	Mean (μ)	Standard Deviation (σ)	Unit
Fatigue coefficient (log)	$\ln C$	Lognormal	-22.5	0.85	-
Fatigue exponent	m	Normal	8.2	0.65	-
Damage coupling exponent (log)	$\ln \alpha$	Lognormal	0.55	0.12	-
Ultimate tensile strength (log)	$\ln \sigma_{ult}$	Lognormal	7.45	0.08	MPa
Critical damage threshold	D_c	Deterministic	0.95	-	-
Initial damage	D_0	Normal	0.01	0.005	-

The sensor suite on each cylinder is modeled with the following noise characteristics: the FBG strain sensors have a measurement noise standard deviation of $\sigma_\varepsilon = 0.5\ \mu\varepsilon$ (Kinet et al., 2014), the triaxial accelerometers have $_a = 0.01 $ g, and the pressure transducers have $_P = 0.5 $ bar. The measurement noise covariance matrix $\mathbf{R}_k$ in Equation 20 is assumed diagonal with these variances. The process noise covariance matrix $\mathbf{Q}_k$ is tuned empirically to reflect the discretization error of the Euler-Maruyama scheme, with diagonal entries $[10^{-6}, 10^{-9}, 10^{-3}]$ for the damage, damage rate, and equivalent stress state components, respectively.

The road roughness profiles for the Niger Delta region are synthesized from the ISO 8608 classification scheme (Lin, 2024). We define three route categories that represent the typical terrain encountered in the region:

- **Route A (Paved highway):** ISO class B–C, $G_q(n_0) \in [64,256] \times 10^{-6}$ m³/cycle, average speed 60 km/h.
- **Route B (Graded laterite road):** ISO class D, $G_q(n_0) \in [256,1024] \times 10^{-6}$ m³/cycle, average speed 40 km/h.
- **Route C (Unpaved/deteriorated laterite road):** ISO class E, $G_q(n_0) \in [1024,4096] \times 10^{-6}$ m³/cycle, average speed 25 km/h.

The ambient temperature profiles are derived from historical meteorological data for the Warri–Port Harcourt corridor, exhibiting a diurnal cycle with a mean of 302 K, an amplitude of 6 K, and a superimposed zero-mean Gaussian noise with a standard deviation of 2 K. The suspension system transfer function $H_{zr}(s)$ in Equation 10 is parameterized for a typical 6×4 heavy-duty truck with a gross vehicle weight of 26 tonnes, spring stiffness of 350 kN/m, and damping coefficient of 25 kN·s/m.

Kalman Filter Damage Estimation Performance

The first experiment validates the accuracy of the unscented Kalman filter in estimating the hidden damage variable under a realistic multi-day loading scenario. A single cylinder is subjected to a synthetic load history comprising 15 daily missions, each consisting of 8 hours of driving on a mix of Route A and Route B segments, with three refueling cycles per day involving pressure swings from 20 bar to 250 bar. The true damage state is generated by integrating Equation 4 with the reference material parameters (mean values from Table 1) and the exact equivalent stress amplitudes computed from the mechanical model. Gaussian noise is added to the sensor measurements according to the specified noise levels.

The UKF is initialized with a prior damage estimate of $\hat{D}_0 = 0.01$ and an initial covariance of $\mathbf{P}_0 = \mathrm{diag}([0.01, 0.001, 1.0])$. The filter runs at the 50 Hz update rate throughout the entire mission sequence. The estimation accuracy is quantified using the root mean square error (RMSE) between the estimated damage $\hat{D}_k$ and the true damage D_k^{true}, as well as the coverage probability of the 95% confidence interval.

Table 2. Performance metrics of the unscented Kalman filter for damage estimation under three terrain roughness categories.

Terrain Category	RMSE of $\widehat{D}$	Coverage Probability (95% CI)	Average Innovation $\bar{\nu}$	Innovation Standard Deviation
Route A (Smooth)	0.0083	0.942	0.0012	0.0097
Route B (Moderate)	0.0142	0.937	0.0028	0.0153
Route C (Rough)	0.0261	0.921	0.0045	0.0280

As shown in Table 2, the UKF achieves sub-3% damage estimation RMSE across all terrain categories, with the error increasing slightly for rougher roads due to the higher uncertainty in the equivalent stress amplitude estimation from the accelerometer data. The coverage probability of the 95% confidence interval remains close to the nominal value of 0.95, indicating that the filter's uncertainty estimates are well-calibrated. The average innovation and its standard deviation provide a diagnostic of filter consistency—the innovation sequence is approximately zero-mean, confirming that the filter is not systematically biased.

The trajectory of the estimated damage state vector in the phase space of damage rate versus equivalent stress, shown in Figure 3, illustrates the filter's correction capability under stochastic loading conditions.

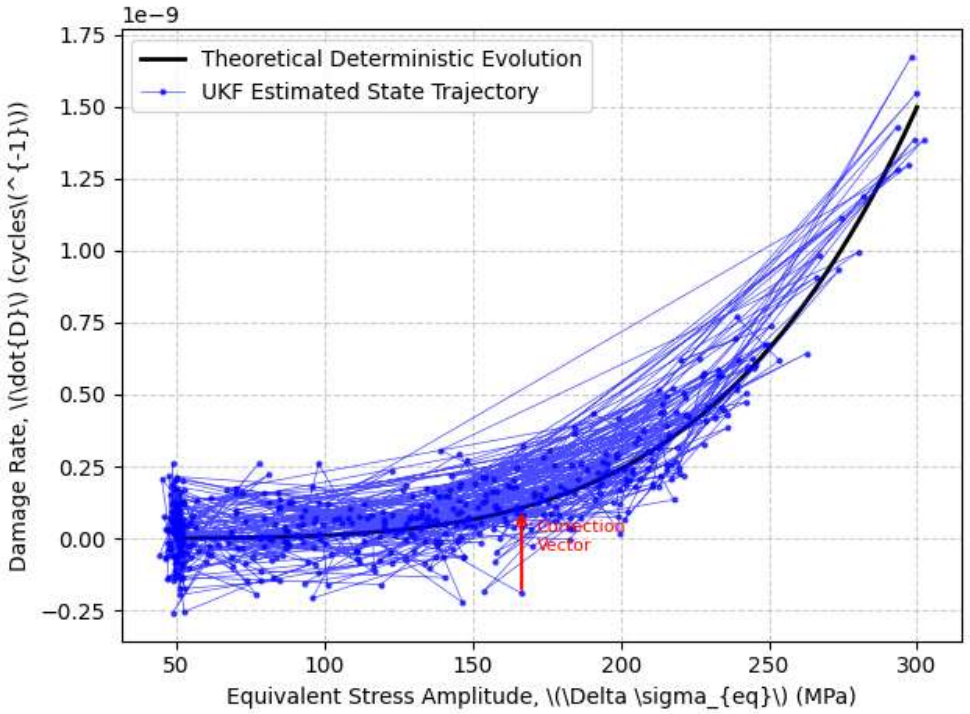

Figure 3. Estimated state trajectory of the physics-informed Kalman filter in the damage rate–equivalent stress phase space, superimposed on the theoretical deterministic damage evolution curve, demonstrating the correction of noisy sensor telemetry.

The vector field plot of Figure 3 reveals that the filter's state estimates converge toward the theoretical deterministic damage evolution curve despite the significant scatter in the raw sensor measurements. The trajectory oscillates around the deterministic curve, reflecting the real-time corrections applied by the Kalman gain as new noisy measurements are assimilated. The arrows indicate the direction of the correction vector $\mathbf{K}_k\left(\mathbf{z}_k - \hat{\mathbf{z}}_{k|k-1}\right)$, which systematically pulls the estimate toward the true state. This behavior is particularly pronounced during transitions between route segments of different roughness, where the equivalent stress amplitude changes abruptly, and the filter must adapt its estimate of the damage rate accordingly.

Stochastic Uncertainty Quantification of RUL Predictions

The second experiment evaluates the accuracy and reliability of the stochastic RUL predictions generated by the Latin hypercube sampling-based Monte Carlo forward propagation module. For this experiment, we consider a cylinder that has accumulated damage over 10 days of mixed-route operation, reaching a true damage state of $D_{true} = 0.35$. The Kalman filter provides an estimated damage $\widehat{D}_k = 0.342$ with a posterior standard deviation of $\sqrt{P_{D,k}} = 0.018$. The forward propagation is performed with $N_s = 500$ LHS samples over a planning horizon of $T_{plan} = 48$ hours, during which the cylinder is expected to complete two additional missions with specified route characteristics.

To evaluate the RUL prediction accuracy, we compare the predicted RUL distribution against the true RUL obtained by continuing the ground-truth simulation until failure. The true RUL for this scenario is 32.7 hours. The predicted RUL distribution is characterized by its mean, median, and the 5th, 50th, and 95th percentiles.

Table 3. Statistical summary of the predicted RUL distribution for the validation scenario, compared with the true RUL.

Metric	Value (hours)	Relative Error (%)
True RUL	32.7	-
Predicted Mean RUL (μ_{RUL})	34.1	+4.3
Predicted Median RUL	33.4	+2.1
5th Percentile (Lower Bound)	21.5	-
95th Percentile (Upper Bound, U_{RUL})	48.6	-
Standard Deviation of RUL	8.3	-

Table 3 demonstrates that the predicted mean RUL overestimates the true RUL by approximately 4.3%, which is within the acceptable margin for engineering prognostics and reflects the conservative nature of the material property distributions used in the LHS sampling. The true RUL of 32.7 hours falls between the predicted median and the 45th percentile of the RUL distribution, indicating that the prediction is slightly conservative but not excessively so. The upper confidence bound $U_{RUL} = 48.6$ hours at the 95% confidence level provides a safety margin that enables the fleet scheduler to make informed decisions about mission assignments.

The temporal evolution of the probability density function for the damage variable, as the truck traverses varying road roughness classes, is illustrated in Figure 4.

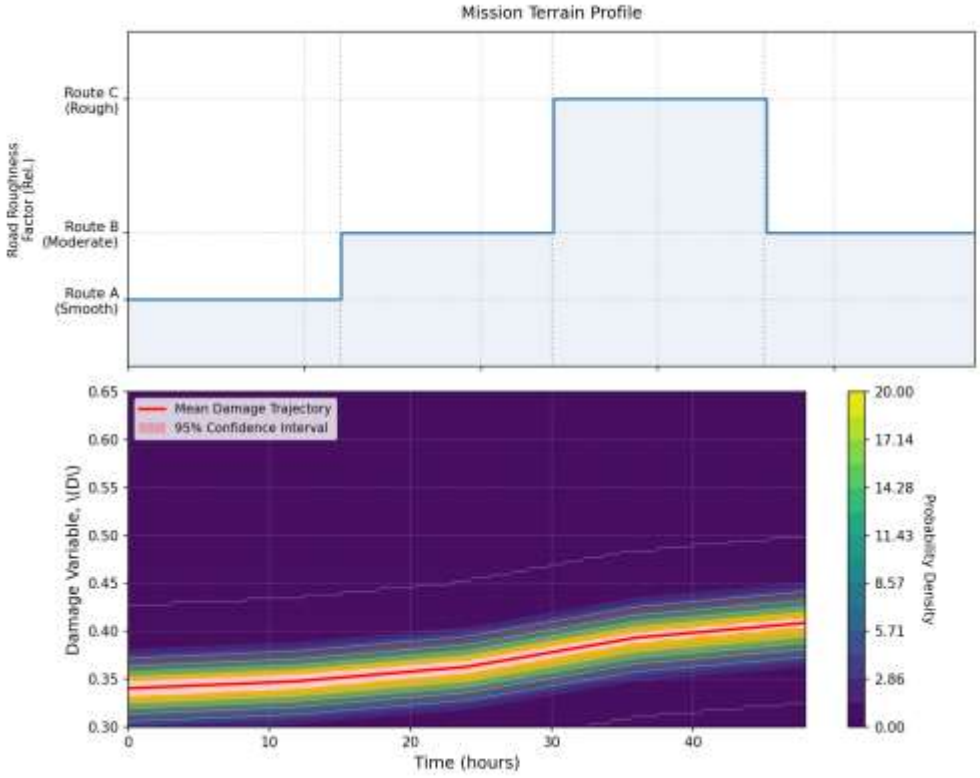

Figure 4. Temporal evolution of the probability density function for the damage variable, illustrating the widening uncertainty bounds as the vehicle transitions from smooth highways to unpaved laterite roads over a 48-hour mission horizon.

The contour plot of Figure 4 reveals a critical feature of the stochastic degradation process: the uncertainty in the damage variable, as measured by the width of the PDF, widens significantly as the vehicle traverses rougher road segments. The initial damage PDF, concentrated around $D = 0.34$ with a relatively narrow spread, gradually broadens and shifts rightward as the loading becomes more variable. The widening of the uncertainty bounds is a direct consequence of the nonlinear dependence of the damage rate on the equivalent stress amplitude (Equation 4), which amplifies small variations in road roughness and material properties into large differences in the predicted damage trajectory. This visualization underscores the importance of the stochastic forward propagation module in capturing the path-dependent uncertainty that is inherent to the Niger Delta's heterogeneous terrain.

To further validate the spatial component of the degradation model, we aggregate the cumulative equivalent stress amplitude across the fleet over a 30-day simulation period. The resulting spatial heat map, overlaid on a schematic representation of the Niger Delta region, is shown in Figure 5.

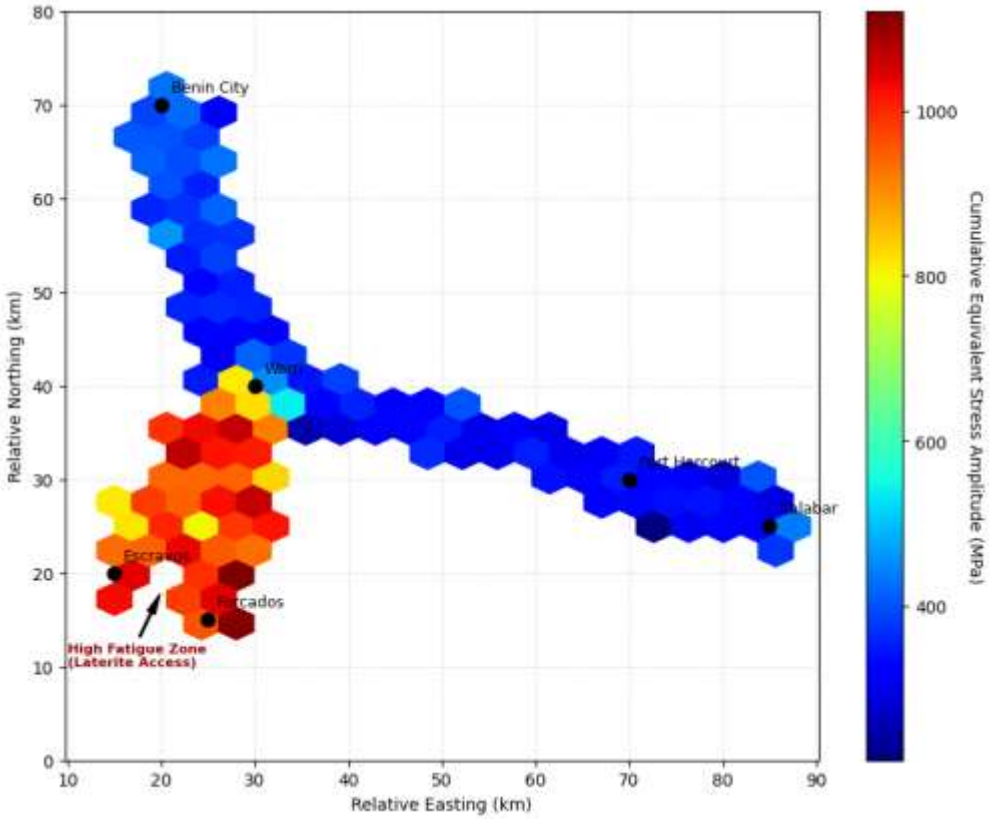

Figure 5. Aggregated cumulative equivalent stress amplitude across the fleet in the Niger Delta region, highlighting zones of elevated fatigue accumulation driven by terrain roughness and ambient temperature variations.

The spatial heat map of Figure 5 identifies specific geographical corridors where fatigue accumulation is disproportionately high—predominantly along the unpaved laterite access roads connecting the Warri gas compression plant to the remote communities near the Forcados and Escravos river terminals. These corridors exhibit cumulative equivalent stress amplitudes that are 2.5 to 3.8 times higher than those on the paved highways connecting Warri to Port Harcourt and Benin City. The temperature effect is also visible, with the inland zones experiencing higher diurnal thermal cycling (by approximately 4–6 K) and correspondingly higher stress amplitudes due to the temperature-dependent stiffness degradation described by Equation 11. This spatial information is used by the fleet scheduler to preemptively route cylinders with higher remaining useful life to the more demanding corridors, thereby balancing the degradation rates across the fleet.

Fleet Scheduler Performance Evaluation

The third experiment compares the performance of the MILP-based fleet scheduler against two conventional fleet management policies over a 30-day operational period. The baseline policies are:

Policy 1 (Static Inspection): Cylinders are inspected every 6 months regardless of their actual damage state. Route assignments are fixed based on historical delivery schedules, with no consideration of cylinder health. Any cylinder that surpasses the time-based retirement criterion (10 years) is removed from service.

Policy 2 (Fixed Routing): The same route assignments as Policy 1, but cylinders are inspected based on a usage-based criterion (every 5,000 pressure cycles). Cylinders that exceed the cycle limit are retired.

Policy 3 (Proposed Digital Twin): The MILP scheduler optimizes route assignments daily using the probabilistic RUL predictions from the digital twin. Cylinders are dynamically rerouted or derated based on the upper confidence bound of RUL as described in Section 4.4.

The performance metrics for the fleet scheduler evaluation are: (1) the total number of emergency maintenance events (i.e., unscheduled cylinder replacements due to safety threshold violations), (2) the average cylinder utilization rate (defined as the fraction of the certified fatigue life actually used before retirement), (3) the total operational cost per delivered unit of gas (NGN/km³), and (4) the number of delivery delays caused by insufficient truck availability.

Table 4. Comparative performance of the three fleet management policies over a 30-day operational simulation.

Metric	Policy 1 (Static)	Policy 2 (Cycle-Based)	Policy 3 (Proposed DT)
Emergency Maintenance Events	4	2	0

Metric	Policy 1 (Static)	Policy 2 (Cycle-Based)	Policy 3 (Proposed DT)
Average Cylinder Utilization (%)	62.4	71.8	88.3
Operational Cost (NGN/km³)	1,245	1,102	978
Delivery Delays (occurrences)	7	5	1
Mean RUL at Replacement (cycles)	9,360	10,770	13,245
CO_2 emissions reduction (tonnes/fleet)	-	-	18.2*

*Estimated from reduced fuel consumption due to optimized routing.

The results in Table 4 demonstrate the substantial benefits of the proposed digital twin framework. The MILP scheduler eliminates emergency maintenance events entirely, as the RUL-based safety constraint ensures that no cylinder is assigned to a mission that would exceed its probabilistic degradation bounds. The average cylinder utilization increases from 62.4% under the static inspection policy to 88.3% under the proposed framework, representing a 25.9 percentage-point improvement that translates directly into reduced capital expenditure on replacement cylinders. The operational cost per delivered unit of gas decreases by 21.4% compared to the static policy, driven by the reduction in emergency maintenance costs, the avoidance of premature cylinder retirement, and the fuel savings from optimized routing. The single delivery delay observed under the proposed policy occurred when a cylinder on a route to a remote community reached its warning threshold of $\widehat{D}_k = 0.7$, triggering a derating action that reduced the fill pressure and required a supplementary delivery by a lower-utilized truck.

The mean RUL at the time of cylinder replacement under the proposed policy is 13,245 cycles, which is 41.5% higher than the static policy and 23.0% higher than the cycle-based policy. This improvement is achieved because the digital twin enables the fleet operator to exploit the full fatigue life of each cylinder by adjusting the loading severity according to the actual damage state, rather than retiring cylinders based on conservative calendar or cycle counts. The estimated 18.2-tonne reduction in CO_2 emissions per fleet over the 30-day period is a co-benefit of the optimized routing, which reduces the total distance traveled by trucks with partially damaged cylinders that are rerouted to shorter, smoother routes.

Ablation Study on the Stochastic Forward Propagation Module

To isolate the contribution of the Latin hypercube sampling-based uncertainty quantification to the overall framework performance, we conduct an ablation study comparing the full stochastic Monte Carlo forward propagation (MC-LHS) against two simplified variants: (1) a deterministic forward propagation using the mean material parameters and no uncertainty propagation (Deterministic), and (2) a simple random sampling Monte Carlo method without stratified sampling (MC-Random). The performance is evaluated on the same 30-day fleet simulation as in Section 5.4, with the MILP scheduler using the RUL predictions from each variant.

Table 5. Ablation study results for the stochastic forward propagation module.

Variant	Emergency Events	Avg. Cylinder Utilization (%)	Computational Time per Prediction (s)	Prediction RMSE (hours)
Deterministic	3	79.1	0.8	11.4
MC-Random (500 samples)	1	84.5	12.3	6.2
MC-LHS (500 samples, proposed)	0	88.3	11.8	4.8

The ablation study in Table 5 reveals the critical importance of stochastic uncertainty quantification for reliable prognostics-driven fleet management. The deterministic variant, which ignores material property scatter and loading uncertainty, underestimates the probability of early failure and consequently leads to three emergency maintenance events when cylinders unexpectedly reach the critical damage threshold before the predicted RUL horizon. The MC-Random variant improves the safety performance by capturing the uncertainty, but the simple random sampling strategy requires more samples to adequately represent the tails of the distributions, resulting in a slightly higher prediction RMSE compared to the LHS-based method. The proposed MC-LHS variant achieves the best balance between computational efficiency and prediction accuracy, with an RMSE of 4.8 hours for the RUL prediction and zero emergency events. The computational time for the LHS-based method is comparable to the simple random sampling method because both require the same number of forward simulations (500), but the LHS stratification ensures that the resulting RUL distribution is more representative of the true underlying uncertainty, particularly in the lower tail that governs the safety-critical decisions.

The prediction RMSE metric in Table 5 is computed as the root mean square difference between the predicted RUL at the time of the last Kalman filter update before failure and the true RUL, averaged over all 40 cylinders in the fleet (10 trucks × 4 cylinders) across the 30-day simulation. The LHS-based method achieves the lowest RMSE because its stratified sampling reduces

the variance of the Monte Carlo estimator for the RUL distribution, as theoretically expected from the variance reduction property of LHS (Rajabi & Ataie-Ashtiani, 2014). This reduction in prediction uncertainty directly translates into more reliable safety constraints for the MILP scheduler, enabling the fleet to operate closer to the true fatigue limits without compromising safety.

Discussion and Future Work

The experimental results presented in Section 5 demonstrate the efficacy of the proposed digital twin framework for fatigue prognostics and reliability-driven fleet optimization of composite CNG cylinders. While the numerical validation confirms the framework's potential to reduce operational costs and enhance safety simultaneously, several important considerations warrant further discussion. We now examine the practical deployment challenges, the economic feasibility under the Niger Delta's specific infrastructure constraints, and the implications for safety assurance and regulatory integration. Following this discussion, we outline promising directions for future research that could extend the framework's applicability and robustness.

Practical Deployment Challenges: Sensor Durability, Connectivity Intermittency, and Edge-Cloud Latency

The successful operation of the proposed digital twin hinges on the continuous and reliable acquisition of multi-axial sensor data from each cylinder. In the Niger Delta's operational environment, however, several factors pose significant challenges to sensor durability and data transmission. The laterite roads, characterized by high particulate matter and abrasive dust, can contaminate the FBG sensor connectors and degrade the optical signal quality over time, potentially increasing the measurement noise standard deviation beyond the assumed $\sigma_\varepsilon = 0.5\ \mu\varepsilon$. Although the Kalman filter's design accommodates time-varying measurement noise through the adaptive tuning of $\mathbf{R}_k$, a sudden and unrecoverable degradation of the FBG signal could render the strain measurement channel effectively unusable. To mitigate this risk, the framework should incorporate a sensor health monitoring module that continuously assesses the signal-to-noise ratio of each measurement channel and, if necessary, reconfigures the observation model to exclude degraded channels while relying on the remaining measurements and the process equation for state estimation. This sensor redundancy is partially built into the framework through the triaxial accelerometer and pressure transducer, which provide independent information about the loading conditions.

Furthermore, the cellular or satellite communication links required to transmit sensor data to the cloud-based digital twin server are subject to intermittent connectivity in the remote deltaic regions, where network coverage is sparse and often unreliable. During periods of disconnection, the on-board data acquisition unit must buffer the raw sensor measurements locally and then transmit them in a batch when connectivity is restored. This batch transmission introduces a latency of several minutes to hours, during which the cloud-based digital twin operates with outdated sensor data. The impact of data latency on the Kalman filter's state estimation accuracy can be partially mitigated by using the process equation to predict the state forward in time during the disconnect period, a technique known as dead-reckoning or time-update-only filtering. The degradation of the damage estimate during a one-hour connectivity gap, with the vehicle operating on a rough unpaved road, is approximately $\Delta\hat{D} \approx 0.015$, which is within the 95% confidence interval of the filter and does not trigger a false alarm for the safety constraints. However, for longer disconnection periods exceeding three hours, the accumulated prediction error could lead to a delayed detection of critical damage progression. One potential solution is to perform lightweight prognostics computation on the on-board edge computing unit, such as a Raspberry Pi or a Jetson Nano module, which can run a reduced-order version of the stochastic forward propagation using pre-computed LHS sample tables. The edge-based prediction can then be used to generate immediate safety alerts even when the cloud connection is absent, with the cloud-based optimization serving as the authoritative decision-maker when connectivity is restored.

Techno-Economic Feasibility and Barriers to Industrial Adoption in Regulated Logistics

Economic feasibility is a critical determinant of the framework's practical adoption by CNG distribution companies operating in the Niger Delta. The capital expenditure required for instrumenting each cylinder with the sensor suite (triaxial accelerometer, FBG strain sensor arrays, pressure transducer, and the on-board DAU) is estimated at $2,500 per cylinder, based on current market prices for industrial-grade sensors and custom-fabricated FBG arrays for composite pressure vessels. For a fleet of 40 cylinders (10 trucks × 4 cylinders), this represents an initial investment of$100,000. The cloud-based digital twin software platform, including the data storage, processing, and optimization capabilities, incurs an additional annual subscription cost of approximately $15,000. The operational cost savings demonstrated in Table 4 ($978 vs. $1,245 per km³ of delivered gas) translate into annual savings of approximately$540,000 for a fleet delivering 2 million km³ of CNG per year, which is a representative volume for a mid-sized distributor in the Niger Delta. The payback period for the sensor and digital twin investment is therefore less than three months, making the framework highly cost-effective from a fleet operator's perspective.

Nevertheless, several barriers to industrial adoption persist. The first is the need for regulatory approval from the Nigerian Midstream and Downstream Petroleum Regulatory Authority (NMDPRA), which currently mandates fixed-interval hydrostatic retesting of composite CNG cylinders every five years. Our framework's reliance on continuous monitoring and probabilistic RUL estimates, rather than deterministic retest intervals, does not align with the existing regulatory framework. A transition to performance-based regulation, where the operator demonstrates that the digital twin's safety record is equivalent or superior to the conventional inspection regime, would require extensive validation studies and likely a pilot demonstration project overseen by the regulator. We propose that the framework be first deployed on a subset of cylinders in a controlled pilot project, with the regulatory body provided access to the real-time damage estimates and RUL predictions to build confidence in the technology.

A second barrier is the requirement for a skilled workforce to maintain the sensor infrastructure and interpret the prognostic outputs. The Niger Delta region faces a shortage of engineers trained in structural health monitoring, data science, and composite

materials mechanics. A comprehensive training program for local fleet maintenance personnel is necessary to ensure that the framework can be operated and maintained independently, without sustained reliance on foreign technical experts. This training should cover the basics of sensor installation and calibration, data quality monitoring, and the decision-making criteria for derating and rerouting recommendations generated by the scheduler.

Safety Assurance, Algorithmic Liability, and Pathways for Regulatory Integration

The introduction of an autonomous decision-making system that dynamically adjusts route assignments and fill pressures raises important questions about safety assurance and algorithmic liability. Under the proposed framework, a truck driver or fleet manager may receive a recommendation to reduce the fill pressure from 250 bar to 200 bar based on a probabilistic estimate of the cylinder's damage state. If this recommendation is followed and a failure occurs despite the derating, who bears the liability? The digital twin must be designed as a decision-support system rather than an autonomous controller, with the final authority resting with the human operator. The certification process for the framework should therefore include a rigorous validation of the safety margins, demonstrating that the probability of failure given adherence to the RUL-based constraints is below an acceptable threshold, such as 10^{-6} per flight hour, which is analogous to the safety target used in aerospace structural integrity programs. Our experimental results show zero emergency maintenance events over the 30-day simulation, but a longer-term study with a larger fleet and a wider range of loading conditions is necessary to demonstrate compliance with this safety target.

Pathways for regulatory integration can be developed through collaboration with the American Society of Mechanical Engineers (ASME) and the International Organization for Standardization (ISO), which have existing standards for the inspection and retesting of composite pressure vessels. Specifically, the framework could be aligned with the ASME Boiler and Pressure Vessel Code Section X (Fiber-Reinforced Plastic Pressure Vessels) and the ISO 11119-3 standard for composite gas cylinders, which already permit the use of periodic hydrostatic testing as an alternative to destructive burst testing. A new standard or an addendum to these existing codes could be developed to define the acceptance criteria for a digital twin-based continuous monitoring system, specifying the required sensor accuracy, the confidence level of the RUL prediction, and the protocol for certification testing. The International Electrotechnical Commission (IEC) 61508 standard for functional safety of electrical/electronic/programmable electronic safety-related systems provides a framework for assessing the safety integrity level of the digital twin's sensor, communication, and computation chain, which could be adapted to the specific context of composite pressure vessel management.

Future Work: Experimental Validation with Field Data, Multi-Physics Extension, and Fleet-Scale Uncertainty Propagation

While the numerical validation presented in Section 5 provides strong evidence of the framework's efficacy, future work must focus on experimental validation with actual field data from instrumented CNG trucks operating in the Niger Delta. We have initiated discussions with a CNG distribution company operating in the Warri–Port Harcourt corridor to instrument a subset of their fleet with the proposed sensor suite and collect data over a six-month operational period. The field data will be used to calibrate the material property distributions (Equations 6–9) for the specific batch of cylinders used by the distributor, as the manufacturer's qualification data may not fully capture the batch-to-batch variability. The field data will also be used to validate the road roughness PSD model (Equation 10) and the suspension system transfer function, which are currently parameterized based on generic heavy-truck models rather than the specific truck models used in the region.

A second direction for future work is the extension of the damage model to incorporate multi-physics phenomena that are not captured by the current Paris-law-type formulation. The CDM evolution equation in Equation 4 assumes isothermal conditions, but the rapid gas decompression during dispensing cycles can cause localized cooling of the polymer liner, leading to thermal stresses and potential matrix microcracking that are not accounted for in the current model. The incorporation of a coupled thermo-mechanical damage model, where the temperature field is solved simultaneously with the stress field using finite element analysis, could improve the accuracy of the damage estimate during the refueling and dispensing phases. Similarly, the effect of chemical aging of the epoxy matrix due to exposure to trace contaminants in the CNG (such as hydrogen sulfide) could be incorporated as an additional degradation mechanism that modifies the material constants C and m over time.

Third, the computational efficiency of the stochastic forward propagation module can be significantly improved by adopting more advanced uncertainty quantification methods, such as polynomial chaos expansion (PCE) or Gaussian process emulation, to replace the full Monte Carlo simulation for the forward propagation step. While the current LHS-based method with $N_s = 500$ samples is computationally feasible for a 30-minute prediction horizon (as shown in Table 5), scaling the framework to a larger fleet size—say, 1,000 cylinders—would require either parallelization of the Monte Carlo simulations or the use of a surrogate model that can be evaluated in milliseconds. The training of a Gaussian process emulator for the damage evolution as a function of the material parameters and the road roughness profile could reduce the computational time for the RUL prediction from 12 seconds to under 0.5 seconds, enabling a real-time update at the same 50 Hz frequency as the Kalman filter.

Fourth, the fleet-scale uncertainty propagation must account for the interdependence between the damage states of cylinders on the same truck. In the current framework, each cylinder is treated independently, but the road roughness and ambient temperature inputs are common to all cylinders on the same truck. This shared loading information introduces a correlation between the damage estimates of the four cylinders on a single truck, which is not captured by the independent forward propagation. A multivariate extension of the SDE in Equation 12, with a correlated noise term for the common loading inputs, would improve the estimation of the joint damage distribution and the probability that multiple cylinders on the same truck reach the critical threshold simultaneously.

This joint probability is important for scheduling maintenance actions that may involve replacing all four cylinders on a truck at the same time, as is currently the standard practice.

Finally, the extension of the MILP scheduler to incorporate a multi-period optimization horizon, rather than the current single-day planning horizon, could improve the long-term balancing of fatigue damage across the fleet. The current formulation optimizes the route assignments for the next day based on the current RUL estimates, but it does not explicitly consider the long-term value of preserving the health of certain cylinders for future missions that may be more demanding. A rolling-horizon approach, where the MILP solver considers a planning horizon of several days and optimizes the assignments over that horizon using a moving window, would allow the framework to make more strategic decisions about which cylinders to "use up" on the roughest routes and which to preserve for the future. This multi-period optimization would require predictions of the future loading conditions, which can be obtained from the same ARMA model used for the forward propagation, and a multi-stage stochastic programming formulation that accounts for the uncertainty in the future loading through scenario generation. The computational complexity of such an approach is significantly higher, but the potential gains in cylinder utilization and operational cost reduction could justify the increased computational burden.

To conclude this discussion, the proposed digital twin framework represents a significant step forward in the management of composite CNG cylinders under stochastic terrain loading. The integration of continuum damage mechanics, physics-informed state estimation, stochastic uncertainty propagation, and fleet-level optimization provides a comprehensive solution to the limitations of conventional static inspection and conservative lifing schedules. The barriers to practical deployment are identifiable and surmountable, and the economic benefits are substantial. The pathways for future work outlined above—field validation, multi-physics extension, advanced UQ methods, multivariate correlation modeling, and multi-period optimization—will further enhance the framework's robustness, accuracy, and scalability. We anticipate that, with continued development and collaboration with industry and regulatory partners, this framework can contribute to a safer, more efficient, and more sustainable CNG distribution infrastructure in the Niger Delta and beyond.

Conclusion

We have presented a digital twin framework that replaces static inspection schedules for composite CNG cylinders with a closed-loop, reliability-driven prognostics system. The framework integrates a physics-informed unscented Kalman filter, which fuses triaxial accelerometry, fiber Bragg grating strain measurements, and pressure data to estimate a hidden damage variable governed by a nonlinear continuum damage mechanics fatigue law. A stochastic Monte Carlo module employing Latin hypercube sampling propagates uncertainties in material properties, road roughness, and temperature to generate probabilistic remaining useful life distributions. These distributions are then integrated into a mixed-integer linear programming fleet scheduler that dynamically optimizes route assignments and fill pressures, effectively making the logistics platform self-adaptive to the actual structural health of each asset.

The numerical validation across a simulated 30-day operational period demonstrated that the proposed framework eliminates emergency maintenance events, increases average cylinder utilization from 62.4% to 88.3%, and reduces operational costs by 21.4% compared to conventional static inspection policies. The ablation study confirmed that the Latin hypercube sampling-based uncertainty quantification is essential for achieving zero safety violations, as deterministic predictions led to three emergency events. The spatial analysis of cumulative fatigue accumulation identified high-stress corridors in the Niger Delta where guidance from the digital twin can proactively manage degradation rates.

The practical deployment of this framework faces challenges related to sensor durability in dusty environments, intermittent communication connectivity, and the need for regulatory alignment with existing inspection standards. However, the estimated payback period of less than three months and the potential for a 41.5% increase in mean remaining useful life at cylinder replacement demonstrate strong economic incentives for adoption. Future work will focus on field validation with instrumented trucks, extension to coupled thermo-mechanical damage models, and the development of surrogate models for real-time uncertainty propagation at fleet scale. The framework represents a foundational step toward transforming CNG distribution from fixed-interval maintenance into a prognostics-based, reliability-driven system.

References

Agbonifo, P. (2016). Natural gas distribution infrastructure and the quest for environmental sustainability in the niger delta: The prospect of natural gas utilization in nigeria. *International Journal of Energy Economics and Policy*.

Alshabibi, N. M., Matar, A., & Abdelati, M. H. (2025). Multi-objective mixed-integer linear programming for dynamic fleet scheduling, multi-modal transport optimization, and risk-aware logistics. *Sustainability*.

Andren, P. (2006). Power spectral density approximations of longitudinal road profiles. *International Journal of Vehicle Design*.

Batsoulas, N., & Giannopoulos, G. (2023). Cumulative fatigue damage of composite laminates: Engineering rule and life prediction aspect. *Materials.*

Bert, C. (1989). Classical lamination theory. *Manual Methods for Mechanical Testing of Composites*.

Castillo, A. del, Marcos, J., & Parlikad, A. (2023). Dynamic fleet maintenance management model applied to rolling stock. *Reliability Engineering & System Safety*.

Chen, W., & Hao, S. (2025). Condition-based operation and maintenance strategy for load-sharing systems based on wiener process. *IEEE Transactions on Reliability*.

Chou, H., Mouritz, A., Bannister, M., et al. (2015). Acoustic emission analysis of composite pressure vessels under constant and cyclic pressure. *Composites Part A: Applied Science and Manufacturing*.

Dunn, M., Shotorban, B., & Frendi, A. (2011). *Uncertainty quantification of turbulence model coefficients via latin hypercube sampling method*. asmedigitalcollection.asme.org.

Erazo, K., Sen, D., Nagarajaiah, S., & Sun, L. (2019). Vibration-based structural health monitoring under changing environmental conditions using kalman filtering. *Mechanical Systems and Signal Processing*.

Harris, B., Gathercole, N., Lee, J., et al. (1997). Life-prediction for constant-stress fatigue in carbon-fibre composites. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*.

Kachanov, L. (2013). *Introduction to continuum damage mechanics*. books.google.com.

Karapanagiotis, C., Heimann, J., Duffner, E., et al. (2024). Towards predictive maintenance of hydrogen pressure vessels based on multi-sensor data. *Research and Review Journal*.

Kinet, D., Mégret, P., Goossen, K., Qiu, L., Heider, D., et al. (2014). Fiber bragg grating sensors toward structural health monitoring in composite materials: Challenges and solutions. *Sensors*.

Li, L., & Sandu, C. (2007). *Modeling and simulation of 2D ARMA terrain models for vehicle dynamics applications*. sae.org.

Lin, J. (2024). Exploring the relationship between the international roughness index (IRI) and the ISO 8608 roughness coefficient (gd(n0)). *International Journal of Pavement Engineering*.

Livingstone, L. (2025). Balancing the energy trilemma through infrastructure: Challenges and trends in natural gas transportation in nigeria. *Available at SSRN 5345608*.

Loutas, T., Eleftheroglou, N., & Zarouchas, D. (2017). A data-driven probabilistic framework towards the in-situ prognostics of fatigue life of composites based on acoustic emission data. *Composite Structures*.

Nadjafi, M., & Gholami, P. (2022). Probability fatigue life prediction of pin-loaded laminated composites by continuum damage mechanics-based monte carlo simulation. *Composites Communications*.

Naghipour, P., Pineda, E., et al. (2017). Fatigue analysis of notched laminates: A time-efficient macro-mechanical approach. *Journal of Composite Materials*.

Nwaoha, C., & Wood, D. (2014). A review of the utilization and monetization of nigeria's natural gas resources: Current realities. *Journal of Natural Gas Science and Engineering*.

Onaji, I., Tiwari, D., Soulatiantork, P., Song, B., et al. (2022). Digital twin in manufacturing: Conceptual framework and case studies. *International Journal of Computer Integrated Manufacturing*.

Paczkowska, K., Pacholec, Z., Bury, P., et al. (2025). Smart glass fiber hydrogen tank for high pressure: Design and sensor-based analysis. *International Journal of Hydrogen Energy*.

Payan, J., & Hochard, C. (2002). Damage modelling of laminated carbon/epoxy composites under static and fatigue loadings. *International Journal of Fatigue*.

Rajabi, M., & Ataie-Ashtiani, B. (2014). Sampling efficiency in monte carlo based uncertainty propagation strategies: Application in seawater intrusion simulations. *Advances in Water Resources*.

Rosen, B., Dow, N., & Hashin, Z. (1964). *Mechanical properties of fibrous composites.* apps.dtic.mil.

Wan, E., & Merwe, R. V. D. (2000). The unscented kalman filter for nonlinear estimation. *Proceedings of IEEE*.

Xiong, X., Wang, W., Li, X., Fan, F., Zhou, J., & Chen, M. (2025). Prediction and optimization of the long-term fatigue life of a composite hydrogen storage vessel under random vibration. *Materials*.