

Article

Reliability-Adaptive Control of Aerospace Electromechanical Actuators with Coupled Degradation via Stochastic MPC

Le Qi 

School of Reliability and Systems Engineering, Beihang University, Beijing 100191, China; qixiaole@buaa.edu.cn

Abstract

Electromechanical Actuators (EMAs) are critical components in More-Electric Aircraft (MEA) and Reusable Launch Vehicles (RLVs), yet they remain vulnerable to jamming and fatigue failures under high-stress flight maneuvers. Existing Health-Aware Flight Control approaches often treat failure prediction and control allocation as separate processes, leading to suboptimal sortie generation rates. This paper presents a reliability-adaptive control framework that unifies trajectory tracking with online health management. Empowered by a hierarchical mission-to-control architecture, the system employs stochastic Model Predictive Control (SMPC) to actively modulate control surface deflection profiles in real time. A comparative case study on a coupled EMA drivetrain demonstrates that the proposed controller extends useful life by 65% compared to fixed-gain baselines, achieves 23% higher mission performance than reactive PID controllers, and it maintains zero constraint violations throughout the mission by optimally distributing the health budget across mission phases.

Keywords: aerospace control; electromechanical actuators; stochastic model predictive control; reliability-adaptive control; trajectory optimization; robust control

MSC: 93E20; 93B45; 90B25; 93E12

1. Introduction

The transition toward More-Electric Aircraft (MEA) and Reusable Launch Vehicles (RLVs) has driven the replacement of hydraulic systems with high-performance Electromechanical Actuators (EMAs) [1]. While EMAs offer superior weight savings and maintainability, they introduce non-deterministic failure modes—specifically jamming due to ball screw wear and gearbox fatigue—that pose severe safety risks during flight [2]. Unlike industrial systems where maintenance can be scheduled arbitrarily, aerospace actuators must sustain operability throughout rigorous flight missions. This necessitates a shift from passive prognostics to active reliability control, where the controller autonomously adapts the mission profile to preserve structural integrity [3].

Current Fault-Tolerant Control (FTC) approaches in aerospace often focus on reconfiguration after a fault occurs. However, a significant gap remains in preventive load management—using control authority to decelerate degradation before functional failure is imminent [4]. This is particularly critical in coupled EMA systems, where vibrations from a degrading gearbox can accelerate wear in the transmission interface (ball screw), creating a nonlinear feedback loop that simpler control laws fail to capture. Furthermore, the stochastic nature of gust loads and material fatigue makes deterministic life predictions unreliable for safety-critical decision-making.



Academic Editor: Duarte Valério

Received: 22 January 2026

Revised: 13 February 2026

Accepted: 20 February 2026

Published: 22 February 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

To address these challenges, this paper presents a reliability-adaptive control framework that integrates stochastic formulation with physics-based degradation modeling. The main contributions are as follows:

1. A hierarchical mission-to-control architecture that translates high-level mission requirements into dynamic health constraints.
2. A stochastic MPC formulation that explicitly handles turbulence (process noise) and fatigue uncertainty via chance constraints. Unlike standard approaches, we employ a conservative constraint-tightening mechanism motivated by the Cantelli inequality to robustly handle non-Gaussian degradation tails.
3. A multi-component case study with physics-based coupled degradation models (gearbox–ball screw interaction), demonstrating the efficacy of the proposed method in a realistic flight control scenario.

2. Related Work

2.1. From PHM to Health-Aware Flight Control

The paradigm of aerospace maintenance is shifting from scheduled time-based maintenance to condition-based strategies. However, as noted in recent aeronautics reviews [4], a critical gap remains between identifying a fault and deciding how to fly with it. While predictive algorithms can forecast Actuator Remaining Useful Life (RUL) with varying accuracy [5], they typically do not inform the guidance loop. Prescriptive maintenance (or health-aware control) aims to close this gap. Early works in aerospace FTC focused on re-configurable control after sudden failures [6]. Recent approaches using deep reinforcement learning (DRL) [7] have attempted to learn optimal policies for degrading systems. However, these data-driven methods often lack the structure to enforce explicit, verifiable safety constraints required for flight certification [8]. In contrast, our work proposes a physics-based control approach that integrates chance-constrained optimization with continuous trajectory modulation to ensure flight safety while maximizing sortie availability.

2.2. Stochastic MPC for Aerospace Reliability

Model Predictive Control (MPC) is increasingly adopted in aerospace for its ability to handle state constraints. Stochastic MPC (SMPC) extends this to handle the probabilistic nature of wind gusts and material fatigue properties [9]. The constraint tightening approach in our translation layer is conceptually related to tube MPC [10] and robust reference governors [11], which ensure robustness by confining trajectories to safe sets. In the context of RLVs, SMPC has been explored for trajectory optimization under uncertainty [12]. However, most existing works treat actuator health as a static constraint or assume independent failure modes [13]. This paper leverages SMPC to actively manage the “Reliability Budget” of the aircraft actuators during the mission [14].

2.3. Coupled Degradation in Electromechanical Actuators (EMAs)

EMAs in RLVs and MEAs are complex assemblies where component interactions drive failure. Liu et al. [15] emphasized that degradation is rarely isolated; for instance, gearbox backlash significantly accelerates ball screw wear due to impact loading. Standard reliability models that treat components as independent [16] often underestimate the system-level risk. Shi et al. [17] have shown that neglecting these coupling effects leads to non-conservative life estimates—a critical oversight for flight certification. Our work explicitly models this vibration-wear coupling (Paris–Archard interaction) within the control horizon, enabling the autopilot to mitigate cascading failures before they jeopardize the vehicle. Unlike existing health-aware SMPC approaches [18,19] that treat degradation as a “passive” state constraint (i.e., enforcing $h(t) \geq h_{min}$ without modeling the inter-component coupling

C), our framework actively exploits the coupling dynamics. By embedding the vibration–wear interaction matrix directly into the prediction horizon, the controller can trade off “upstream” health (gearbox) to preserve “downstream” life (ball screw)—a capability absent in decoupled RUL-constrained formulations.

3. Problem Formulation

3.1. Coupled EMA Dynamics

Consider an EMA system composed of N critical components (motor, gearbox, and ball screw). The health state $h_i(t) \in [0, 1]$ of the i -th component evolves according to an Itô stochastic differential equation:

$$dh_i(t) = -f_i(\mathbf{u}(t), \mathbf{h}(t), \theta_{env}) dt + \sigma_i(\mathbf{h}(t)) dW_i(t) \quad (1)$$

where $\mathbf{u}(t)$ is the control input vector (current/torque), $\mathbf{h}(t) \in \mathbb{R}^N$ represents the system health state vector, θ_{env} denotes environmental parameters (gust load intensity and ambient temperature), $W_i(t)$ is a standard Wiener process on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$, and $\sigma_i(\mathbf{h}) \geq 0$ is the state-dependent diffusion coefficient. The drift term $f_i \geq 0$ ensures $\mathbb{E}[dh_i] \leq 0$, preserving the physical constraint that expected material damage is irreversible. The diffusion coefficient is modeled as $\sigma_i(\mathbf{h}) = \sigma_0 \sqrt{f_i(\cdot)}$, ensuring that variance scales with degradation intensity—a property consistent with empirical fatigue scatter observations [20].

Crucially, the degradation rate f_i depends not only on the component’s own state but also on the health of coupled components (e.g., gearbox vibration affecting the ball screw), modeled by a coupling matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$:

$$f_i(\cdot) = f_{i,base}(\mathbf{u}) \cdot \left(1 + \sum_{j \neq i} C_{ij}(1 - h_j) \right) \quad (2)$$

where $C_{ij} \geq 0$ quantifies the influence of component j ’s degradation on component i ’s wear rate. Note that since $h_j \in [0, 1]$ and $C_{ij} \geq 0$, each term $(1 - h_j) \in [0, 1]$ is non-negative, guaranteeing that the coupling factor satisfies $1 + \sum_{j \neq i} C_{ij}(1 - h_j) \geq 1 > 0$ for all feasible health states. The linear superposition in Equation (2) represents a first-order Taylor expansion of the general coupling function $g_i(1 - h_1, \dots, 1 - h_N)$ around the healthy state ($h_j = 1, \forall j$), which is valid when degradation levels remain moderate ($1 - h_j < 0.5$). This approximation is justified by experimental tribology studies [21] showing that vibration-induced wear amplification is approximately linear in the source component’s damage level for early-to-mid-life operation. The coupling matrix is asymmetric ($C_{ij} \neq C_{ji}$) to capture unidirectional effects (e.g., gearbox vibration affects ball screw wear, but not vice versa). For the EMA system studied, we set $C_{12} = 0.8$ (gearbox-to-screw coupling) and $C_{21} = 0.1$ (weak reverse coupling) to model the asymmetric vibration transfer observed in coupled actuators. These parameters can be identified online via extended Kalman filtering for specific hardware configurations.

Remark (Linear Approximation Error Bound): The truncation error of the Taylor expansion in Equation (2) is $\mathcal{O}(\|(1 - \mathbf{h})\|^2)$. Denoting the true coupling function as $g_i(\mathbf{d})$, where $d_j = 1 - h_j$ represents damage, the approximation error satisfies the following:

$$|g_i(\mathbf{d}) - (1 + \sum_j C_{ij}d_j)| \leq \frac{1}{2} \|\nabla^2 g_i\|_\infty \cdot \|\mathbf{d}\|^2 \quad (3)$$

Assuming typical early-to-mid-life degradation profiles where damage accumulation curvature remains moderate, the nonlinearity of the coupling function can be bounded

by $\|\nabla^2 g_i\|_\infty \approx 0.3$ for $\|\mathbf{d}\| < 0.7$. Thus, with our terminal health $h \approx 0.5$ (i.e., $d \approx 0.5$), the approximation error is bounded by $\frac{1}{2}(0.3)(0.5)^2 \approx 3.75\%$ —acceptable for control purposes. Near failure ($h \rightarrow 0$), crack propagation enters the unstable growth regime where fracture mechanics dominates and linear coupling breaks down; however, the controller is designed to avoid this regime by maintaining $h_i \geq h_{req} > 0$. We note that $h \approx 0.53$ is reached only in the final phase of the mission; for approximately 90% of the operational duration, $h > 0.6$ ($d < 0.4$), which is well within the validity range. Even at the terminal state, the linearization error remains bounded: As shown in Remark 1 (Section 4.2), the cumulative error over the prediction horizon $H = 10$ stays below 2% of the predicted variance for all $h_i \in [0.2, 1.0]$, covering the terminal condition with significant margin. Furthermore, the closed-loop correction of the SMPC compensates for residual model mismatch at every time step via state feedback, preventing the accumulation of approximation errors. This robustness to model errors is empirically confirmed by the sensitivity analysis in Section 6, where the controller successfully completed the mission even under a quadratic coupling model ($C_{ij} \propto (1 - h)^2$), a much larger structural mismatch than the linearization error at $h = 0.53$.

3.2. Physics-Based Degradation Models

To ensure certification-grade rigor, we employ established physics-of-failure models:

- Gearbox (Fatigue): This is modeled using Paris's law [22] for crack propagation in gear teeth:

$$\frac{da}{dt} = C_{paris}(\Delta K)^m \quad (4)$$

where a is the crack length, C_{paris} is Paris's law coefficient, and ΔK is the stress intensity factor, driven by maneuver loads and vibration.

- Ball Screw (Wear): This is modeled using Archard's Law [23] for sliding/rolling wear:

$$\frac{dV}{dt} = k_{archard} \frac{F \cdot v}{H_d} \quad (5)$$

where V is the wear volume, $k_{archard}$ is Archard's wear coefficient, F is the axial load (aerodynamic force), v is the sliding velocity, and H_d is hardness.

- Thermal–mechanical Coupling: The temperature rise due to ball screw friction reduces lubricant viscosity inside the gearbox, increasing contact stress:

$$\mu(T) = \mu_0 e^{-\beta(T-T_0)} \quad (6)$$

These models represent a *grey-box* approach—unifying fundamental aerospace material physics with data-driven coupling parameters. The dimensionless health state $h_i(t) \in [0, 1]$ represents the remaining structural margin (1 = healthy; 0 = functional failure).

Remark (Thermal-State Treatment): The actuator temperature $T \in \theta_{env}$ is modeled as a quasi-static parameter rather than an explicit state variable. Under steady-state flight conditions, actuator temperature equilibrates within seconds (thermal time constant $\tau_{th} \approx 5$ s), while degradation evolves over thousands of cycles (degradation time constant $\tau_{deg} \approx 10^3$ s). This three-order-of-magnitude time-scale separation ($\tau_{th}/\tau_{deg} \approx 10^{-3}$) justifies treating T as an algebraic function of control input: $T = T_0 + \kappa u^2$, where κ captures resistive heating. The viscosity effect $\mu(T)$ thus enters the degradation rate f_i as a control-dependent multiplicative factor, preserving the form of Equation (1) without expanding the state dimension.

3.3. Optimization Objective

The goal is to find a control policy $\pi : \mathbf{h}(t) \rightarrow \mathbf{u}(t)$ that maximizes cumulative performance J_{perf} while ensuring that the system survives until a target time T_{target} with probability α :

$$\max_{\mathbf{u}(t)} \int_0^{T_{target}} P(\mathbf{u}(t), \mathbf{h}(t)) dt \quad (7)$$

$$\text{s.t. } \mathbb{P}(h_i(t) \geq h_{req}(t)) \geq \alpha, \quad \forall i, t \quad (8)$$

To address the timescale separation between slow degradation dynamics and fast control dynamics, we apply hierarchical decomposition to the global objective (7). This splits the problem into two timescale-separated layers:

1. Slow-Time Scale (Reliability Planning): Determine a feasible health reference trajectory $h_{req}(t)$ that satisfies the chance constraint $\mathbb{P}(h_i(t) \geq h_{req}(t)) \geq \alpha$.
2. Fast-Time Scale (Performance Maximization): Maximize instantaneous performance subject to tracking this reference:

$$\max_{\mathbf{u}(t)} J_{perf}(\mathbf{u}(t)) \quad (9)$$

$$\text{s.t. } h_i(t) \geq h_{req}(t) - \epsilon(t) \quad (10)$$

where $\epsilon(t)$ is a slack variable allowing for short-term optimal deviations.

This decomposition naturally maps to the hierarchical architecture proposed in Section 4.

3.4. The Reliability Paradox: Why Derating Increases Sortie Generation

A naive autopilot might attempt to maximize instantaneous control authority at each timestep. However, this locally optimal strategy is globally suboptimal due to the nonlinear coupling between maneuver intensity and failure risk. Consider Paris's law: A 10% increase in stress intensity ΔK causes $1.1^m \approx 1.31 \times$ faster crack growth (for $m = 2.8$). The cumulative damage from aggressive high-G maneuvers early in life forfeits future operational availability—a form of reliability debt.

The key insight is that maximizing mission effectiveness (area under the performance curve) requires limiting peak loads during high-health periods to preserve actuation authority for the degraded-state regime. This inverts the traditional control objective: Instead of tracking a fixed gain, we track a health consumption rate, treating structural integrity as a depletable resource to be invested optimally across the service life.

4. Methodology

4.1. Hierarchical Mission-to-Control Architecture

The proposed framework (Figure 1) aligns with standard cyber-physical system (CPS) hierarchies (e.g., ISA-95 [24]) but introduces a critical innovation: the explicit probabilistic feedback of reliability constraints into the control loop. Rather than treating maintenance as an external schedule, the architecture embeds degradation physics directly into the decision process via a hierarchical translation mechanism:




Figure 1. Hierarchical mission-to-control architecture. The colored dashed boxes indicate hierarchical levels: blue = Decision & Planning Level, green = Control & Estimation Level, red = Physical Level. The key innovation is the translation layer, which actively converts stochastic reliability demands into deterministic control references.

4.1.1. Decision and Planning Level

This top level handles the stochastic nature of mission requirements using a dedicated translation layer:

- **Demand Layer (Goal Setting):** Defines the high-level mission profile (e.g., “Complete 1000 cycles”) and reliability constraints. It serves as the reference input.
- **Translation Layer (Prescription):** This layer serves the role of a Reference Governor [11] by actively converting high-level mission reliability checks into feasible state trajectories. It computes a Target Health Trajectory $h_{req}(t)$ compatible with the system’s physical degradation limits, ensuring that the control targets are not only probabilistically safe but also physically reachable.

4.1.2. Control and Estimation Level

This level ensures that the system tracks the translated prescription despite physical uncertainties:

- **Algorithm Layer (Execution):** Executes the stochastic MPC to compute optimal control actions $\mathbf{u}(t)$ that minimize the deviation between predicted health and the translated reference trajectory.
- **Model Layer (Prediction):** Contains the physics-based coupled degradation models. It predicts future state distributions based on current estimates.
- **Data Layer (Perception):** Fuses raw sensor data to estimate unobservable health states via state estimation (e.g., extended Kalman filter). This layer bridges the physical measurements to the model-based prediction.

4.1.3. Physical Level

This level interfaces with real-world hardware:

- Physical Layer (Plant): The multi-component EMA (motor–gearbox–ball screw) executing deflection commands and generating sensor measurements.

This hierarchical separation allows for plug-and-play flexibility: The *translation layer* isolates the complex probabilistic mission logic from the deterministic tracking controller, creating a scalable architecture for reliability-adaptive systems.

4.2. Stochastic Model Predictive Control

We employ a rolling-horizon approach. The continuous-time SDE (1) is discretized using the Euler–Maruyama scheme with sampling time Δt :

$$h_i(k+1) = h_i(k) - f_i(\mathbf{u}(k), \mathbf{h}(k))\Delta t + \sigma_i(\mathbf{h}(k))\sqrt{\Delta t}\xi_i(k) \quad (11)$$

where $\xi_i(k) \sim \mathcal{N}(0, 1)$ denotes i.i.d. standard normal random variables. At each discrete time step k , we solve a finite-horizon optimal control problem over the prediction horizon H .

To handle the chance constraints' tractably, we propagate the first two moments of the health state's distribution. Let $\bar{\mathbf{h}}(k) = \mathbb{E}[\mathbf{h}(k)]$ denote the mean and $\Sigma(k) = \text{Cov}[\mathbf{h}(k)]$ denote the covariance. Under the linearization of f_i around the current mean trajectory,

$$\Sigma(k+1) = \mathbf{A}(k)\Sigma(k)\mathbf{A}(k)^T + \mathbf{Q}(k) \quad (12)$$

where $\mathbf{A}(k) = \mathbf{I} - \nabla_{\mathbf{h}}\mathbf{f}|_{\bar{\mathbf{h}}(k)}\Delta t$ is the Jacobian, and $\mathbf{Q}(k) = \text{diag}(\sigma_i^2(\bar{\mathbf{h}}(k)))\Delta t$ is the process noise covariance.

Uncertainty Characterization: The stochastic diffusion term $\sigma_i(\cdot)$ serves as a lumped parameter capturing the aggregate effect of the following:

1. Process Noise (Turbulence): High-frequency aerodynamic load variations ($\sigma_{env} \approx 15\%$ of nominal load).
2. Model Mismatch: Aleatoric uncertainty in material properties and epistemic uncertainty from unmodeled dynamics (e.g., actuator backlash).

By treating these combined uncertainties as a state-dependent diffusion process, the controller's safety margin α dynamically adapts to the growing variance over the prediction horizon.

Remark 1. (Linearization Error Bound): The linearization error $\epsilon_k = h_i(k) - \bar{h}_i(k) - \nabla_{\mathbf{h}}f_i \cdot (\mathbf{h}(k-1) - \bar{\mathbf{h}}(k-1))$ is bounded by $|\epsilon_k| \leq \frac{1}{2}L\|\mathbf{h}(k-1) - \bar{\mathbf{h}}(k-1)\|^2\Delta t$, where L is the Lipschitz constant of the gradient $\nabla_{\mathbf{h}}f_i$ (i.e., $\|\nabla f_i(\mathbf{h}) - \nabla f_i(\mathbf{h}')\| \leq L\|\mathbf{h} - \mathbf{h}'\|$). For our parameterization with $m = 2.8$, we have $L < 0.5$ over the operating range $h_i \in [0.2, 1.0]$. Combined with the slow degradation dynamics ($\Delta t \cdot f_i < 10^{-3}$), the cumulative linearization error over horizon $H = 10$ remains below 2% of the predicted variance.

Remark 2. (Gaussian Assumption Validity): The Euler–Maruyama discretization with additive Gaussian increments $\xi_i(k)$ yields conditionally Gaussian state updates. While fatigue life often follows a Weibull distribution, the Gaussian approximation is valid here because we update the conditional probability density function at every step via filtering, rather than predicting the entire lifetime distribution from $t = 0$. Empirical validation via Monte Carlo simulation ($N = 10^4$ trajectories) confirms that the Kolmogorov–Smirnov statistic between the one-step-ahead conditional health distribution and the fitted Gaussian remains below $D_{KS} < 0.03$ for $h_i > 0.3$. The threshold $h_i = 0.3$ was selected as the point at which the KS statistic first exceeds 0.03, indicating a statistically significant departure from Gaussianity due to boundary effects at $h = 0$.

For degraded states ($h_i \leq 0.3$), the distribution exhibits positive skewness due to the reflecting boundary at $h = 0$. To maintain the prescribed safety level despite this non-Gaussianity, we apply a conservative confidence adjustment:

$$\alpha'(h_i) = \begin{cases} \alpha & \text{if } h_i > 0.3 \\ 1 - \frac{1-\alpha}{2} & \text{if } h_i \leq 0.3 \end{cases} \quad (13)$$

For $\alpha = 0.95$, this yields $\alpha' = 0.975$ in the degraded regime. This adjustment halves the admissible risk, $\alpha' = 1 - (1 - \alpha)/2$, i.e., the tolerable failure probability is reduced from $(1 - \alpha)$ to $(1 - \alpha)/2$. This conservative tightening is motivated by the Cantelli inequality, which provides distribution-free one-sided tail bounds. While a full Cantelli-based formulation would require variance-dependent quantile adjustment, we adopt this simpler fixed tightening as a practical conservative approximation. The factor of 2 represents the minimal conservative correction, analogous to a Bonferroni-type adjustment for a single additional source of uncertainty (non-Gaussianity). By increasing the safety buffer ($z_{\alpha'}$ vs. z_{α}), this covers the additional uncertainty introduced by the Weibull-like fat tails and linearization errors bounded in Remark 1.

The chance constraint $\mathbb{P}(h_i(k) \geq h_{req}(k)) \geq \alpha$ is converted to a deterministic constraint using the Gaussian quantile:

$$\bar{h}_i(k) - \Phi^{-1}(\alpha) \sqrt{\Sigma_{ii}(k)} \geq h_{req}(k) \quad (14)$$

This formulation ensures that the controller maintains a safety buffer proportional to the uncertainty, preventing premature failure even under stochastic disturbances. The complete algorithm is summarized in Algorithm 1.

Algorithm 1 Reliability-Adaptive SMPC

Require: Current health $\mathbf{h}(k)$, target trajectory $h_{req}(\cdot)$, horizon H , confidence α

Ensure: Optimal control $\mathbf{u}^*(k)$

```

1: Initialize:  $\bar{\mathbf{h}}_0 \leftarrow \mathbf{h}(k), \Sigma_0 \leftarrow \mathbf{0}$ 
2: for each candidate action  $\mathbf{u} \in \mathcal{U}$  do
3:    $J(\mathbf{u}) \leftarrow 0$ 
4:   for  $j = 0$  to  $H - 1$  do
5:     // Propagate mean dynamics
6:      $\bar{\mathbf{h}}_{j+1} \leftarrow \bar{\mathbf{h}}_j + \mathbf{f}(\mathbf{u}, \bar{\mathbf{h}}_j)\Delta t$ 
7:     // Propagate covariance (Eq. 12)
8:      $\Sigma_{j+1} \leftarrow \mathbf{A}_j \Sigma_j \mathbf{A}_j^T + \mathbf{Q}_j$ 
9:     // Compute safety margin
10:     $h_{safe} \leftarrow \bar{h}_i - \Phi^{-1}(\alpha) \sqrt{\Sigma_{ii}}$ 
11:    // Accumulate cost
12:    if  $h_{safe} < h_{req}(k + j)$  then
13:       $J(\mathbf{u}) \leftarrow J(\mathbf{u}) + \lambda(h_{req} - h_{safe})^2$ 
14:    end if
15:     $J(\mathbf{u}) \leftarrow J(\mathbf{u}) - w_{perf} \cdot P(\mathbf{u}, \bar{\mathbf{h}}_j)$ 
16:  end for
17: end for
18:  $\mathbf{u}^*(k) \leftarrow \arg \min_{\mathbf{u} \in \mathcal{U}} J(\mathbf{u})$ 
19: return  $\mathbf{u}^*(k)$ 

```

4.3. Parametric Multi-Objective Optimization

While the translation layer provides a baseline reference $h_{req}(t)$, the lower-level controller retains the freedom to exploit the health slack for performance gain. We formulate

this as a slack exploitation problem using parametric multi-objective optimization (PMOO). The optimization minimizes deviation from the reference while maximizing performance:

$$J(\mathbf{u}) = w_{\text{track}} \|\bar{\mathbf{h}}_{k+1} - h_{\text{req}}(k+1)\|^2 - w_{\text{perf}} J_{\text{perf}}(\mathbf{u}) \quad (15)$$

where $\mathcal{U} = \{\mathbf{u} : u_{\min} \leq u_i \leq u_{\max}\}$ is the compact control constraint set. The performance metric is defined as $J_{\text{perf}}(\mathbf{u}, \mathbf{x}) = u_{\text{load}} \cdot u_{\text{velocity}} \cdot h_{\text{ball screw}}$, representing the instantaneous effective power delivered to the load. This health-weighted formulation implies that operation with degraded components yields lower utility due to stiffness loss and backlash. The term “Mission Performance” reported in Section 5 refers to the cumulative integral of this effective work capacity $\int_0^T J_{\text{perf}} dt$ over the proficient life of the actuator.

Remark 3. (Non-Convexity and Solution Strategy): The feasible set defined by chance constraints (7) is generally non-convex due to the nonlinear dependence of $\Sigma_{ii}(k)$ on the control sequence. However, for a fixed control horizon $N_u = 1$, the optimization reduces to evaluating a finite set of candidate actions $|A| = 9$ (discretized control inputs). We employ exhaustive enumeration over this discrete set, guaranteeing global optimality within the discretization resolution. For finer discretization or larger N_u , sequential quadratic programming (SQP) with multiple random restarts (10 initializations) is used, with convergence verified by comparing objective values across restarts. In our experiments, all restarts converged to solutions within 0.1% of each other, suggesting a benign optimization landscape for this problem class.

The trade-off is resolved via weighted scalarization where the tracking penalty dominates: $w_{\text{track}}/w_{\text{perf}} = 100$. This effectively enforces the reliability constraint as a soft barrier that stiffens as the system approaches the reference limit, allowing the exploitation of health margins only when safety is not compromised. This constraint acts as a dynamic “health budget”, enforced via the chance constraint $\mathbb{P}(h_i(t) \geq h_{\text{req}}(t)) \geq \alpha$, which allows the optimizer to aggressively consume margin early in the mission while guaranteeing terminal safety.

5. Case Study

5.1. Setup

We validate the framework on a canonical primary flight control actuator (PFCA) architecture (Figure 2). This configuration—comprising a high-speed AC motor, a planetary reduction gearbox, and a ball screw transmission—is representative of the jam-critical stabilization systems found in modern MEA platforms. The degradation parameters are calibrated to reflect the fatigue characteristics of aerospace-grade hardened steel couples (e.g., AISI 9310 gearing driving a 52100 steel ball screw). Crucially, the model captures cross-domain coupling: pitting fatigue in the gearbox induces high-frequency vibration, which accelerates fretting wear in the ball screw track.

5.2. Implementation Details

The simulation parameters are calibrated to reflect realistic aerospace operating conditions. The mission profile consists of $T = 8000$ flight cycles, where each cycle represents a complete control surface deflection sequence (extension–retraction) under stochastic gust loads modeled as Gaussian white noise with a standard deviation of $\sigma_{\text{env}} = 0.15$. System failure is defined as any component’s health reaching the critical threshold ($h_i < 0$), corresponding to gearbox tooth fracture (crack length > 40 mm) or ball screw seizure (backlash > 0.3 mm).



Figure 2. Canonical primary flight control actuator (PFCA) architecture: The coupled drivetrain comprises a high-speed AC motor, planetary gearbox (AISI 9310 steel), and ball screw transmission (52100 steel). Gearbox fatigue-induced vibration propagates to the ball screw, accelerating fretting wear—the coupling mechanism central to this study.

The degradation parameters are set as follows: Paris’s law coefficient $C_{\text{paris}} = 5 \times 10^{-10}$ with exponent $m = 2.8$ for the gearbox; Archard’s wear coefficient $k_{\text{archard}} = 5 \times 10^{-8}$ for the ball screw. The vibration coupling coefficient is $C_{12} = 0.8$, representing the amplification factor by which gearbox-induced vibrations accelerate ball screw fretting wear. The SMPC controller operates with a prediction horizon of $H = 10$ steps and confidence level of $\alpha = 0.95$ for the chance constraints, ensuring a 95% probability of meeting health targets under process uncertainty. For reproducibility, all Monte Carlo simulations use sequential random seeds starting from seed $s_0 = 42$ (i.e., run i uses seed $42 + i$). Table 1 summarizes all key simulation and control parameters.

Table 1. Summary of key simulation and control parameters

Parameter	Symbol	Value
Mission Duration	T_{target}	8000 cycles
Env. Noise Std. Dev.	σ_{env}	0.15
Paris’s Law Coeff.	C_{paris}	5×10^{-10}
Paris’s Exponent	m	2.8
Archard Coeff.	k_{archard}	5×10^{-8}
Vibration Coupling	C_{12}	0.8
MPC Horizon	H	10 steps
Chance Constraint	α	0.95

5.3. Results

Figure 3 visualizes the system’s evolution in a three-dimensional state space where the axes represent the gearbox’s health (x), ball screw’s health (y), and instantaneous sortie effectiveness (z). The color gradient encodes time progression (purple/dark = early mission; yellow/bright = late mission). This visualization reveals two key observations: (1) The gearbox’s and ball screw’s health declines in a coordinated manner, confirming the vibration–wear coupling mechanism; (2) as both health states degrade, the system’s

instantaneous performance (z-axis) decreases correspondingly, reflecting the physical reality that worn components deliver reduced output.




Figure 3. 3D state-space trajectory showing gearbox health (x-axis), ball screw health (y-axis), and instantaneous sortie effectiveness (z-axis, performance metric). Color gradient indicates time progression (purple/dark = $t \approx 0$; yellow/bright = $t \approx 8000$). The coordinated decline of both health axes confirms the coupling mechanism, while the decreasing z-values reflect performance degradation as components wear.

The trajectory's smooth, continuous nature indicates stable controller behavior throughout the mission. The slight nonlinearity in the path reflects the exponential nature of Paris's Law ($m = 2.8$) and the controller's adaptive response to changing health margins.

Figure 4 details the component health evolution. Unlike the baseline, the stochastic MPC exhibits three distinct behavioral phases:

- Phase I (Conservative, $t < 3000$): Early in life, when components are healthy and degradation rates are inherently low, the controller operates conservatively (low control inputs). This preserves health margins for later phases when coupling effects will accelerate wear, effectively “banking” reliability budgets for future use.
- Phase II (Ramp-Up, $3000 < t < 6000$): As the mission progresses, the controller gradually increases operating intensity. It balances the need to generate performance against accumulating degradation, carefully monitoring the coupling feedback between gearbox vibrations and ball screw wear.
- Phase III (Aggressive, $t > 6000$): In the final phase, the controller operates at near-maximum capacity, spending the remaining health budget before the target time. This “spend-down” strategy ensures that the health investment made in early phases is fully utilized, rather than wasted as unused margin at mission end.

This emergent “bank-then-spend” behavior is the key advantage over reactive controllers like PID, which maintain constant operating intensity throughout.




Figure 4. Coupled degradation profiles under stochastic MPC. The three components exhibit distinct degradation rates: motor (slowest, thermal-limited), gearbox (Paris’s law fatigue), and ball screw (fastest, Archard wear amplified by gearbox vibration). Three behavioral phases emerge: Phase I (conservative, $t < 3000$) preserves the health margin; Phase II (ramp-up, $3000 < t < 6000$) gradually increases intensity; Phase III (aggressive, $t > 6000$) spends the remaining budget. Monte Carlo validation ($N = 10$ runs, $\sigma_{env} = 0.15$) showed $< 2\%$ variance in trajectories, indicating robust controller performance.

5.4. Baseline Comparison

Table 2 compares the proposed stochastic MPC against three baselines representing progressively more capable control strategies:

- **Fixed-Gain Control:** Open-loop strategy operating at 80% rated capacity—a common industrial practice that prioritizes throughput over longevity. This represents the “run-to-failure” mentality where operators lack real-time health feedback.
- **PID Health Tracking:** Closed-loop reactive controller using only instantaneous health observations. Critically, the PID has no knowledge of the mission duration T_{target} and must estimate degradation rates online, representing practical limitations of feedback-only control.
- **Deterministic MPC:** Predictive controller using the same horizon ($H = 10$) and physics-based degradation model as the SMPC but without stochastic uncertainty handling. This isolates the contribution of predictions from stochastic health awareness.

Table 2. Performance comparison ($T_{target} = 8000$, $N = 20$ runs, $\sigma_{env} = 0.15$). Values shown as mean $\pm$ std where variation exists.

Strategy	Lifetime (Cycles)	Performance (Units)	Terminal Health	Target Error
Fixed-Gain Control	4841 ± 7	2149 ± 3	0.00 (Dead)	+39.5% (Fail)
PID Health Tracking	8000	1960 ± 0	0.73	$< 0.1\%$
Deterministic MPC	4839 ± 0	2147 ± 0	0.00 (Dead)	+39.5% (Fail)
Reliability-Adaptive MPC	8000	2412 ± 1	0.53	$< 0.1\%$

The results reveal a key finding: The deterministic MPC fails at a statistically comparable point to the fixed-gain controller ($t = 4839$), despite possessing identical predictive capability to the SMPC. Without chance constraints to enforce safety margins under uncertainty, the deterministic MPC operates at maximally aggressive settings and drives the ball screw to failure. This demonstrates that predictive capability alone is insufficient—it is the stochastic health-awareness (chance constraints + uncertainty propagation) that en-

ables both mission survival and the 23% performance gain over PID (2412 vs. 1960 units). The SMPC optimally distributes the health budget across mission phases, operating aggressively in early phases and conservatively as coupling effects amplify wear rates.

6. Discussion

The results demonstrate that the reliability-adaptive MPC outperforms all baselines in both survival and mission effectiveness. The fixed-gain strategy, prioritizing high throughput, depletes the ball screw backlash margin prematurely at $t = 4841$. The proposed SMPC extends useful life by 65% ($8000/4841 \approx 1.65\times$) compared to this baseline. Notably, the deterministic MPC—which shares the same prediction horizon and physics model—fails identically to the fixed-gain controller. This confirms that the performance gain does not arise from predictive capability alone but specifically from the stochastic health-awareness mechanism (chance constraints with uncertainty propagation). The SMPC achieves **23%** higher performance than the PID controller (2412 vs. 1960 units) while maintaining full mission survival. This advantage arises from the SMPC’s ability to optimally distribute the health budget across the mission: It operates more aggressively in early phases (when components are healthy and degradation is slow) and conservatively in later phases (when coupling effects amplify wear rates). The PID controller, lacking predictive capability, maintains a suboptimal constant operating point throughout.

The near-linear health trajectories in Figure 4 initially appeared to be counter-intuitive given the exponential nature of Paris’s law ($m > 2$). Extensive verification confirmed that this is not a modeling error but an emergent property of the prescriptive control: The controller actively “fights” the exponential growth by progressively derating the load to maintain a constant damage rate. This emergent “active linearization” proved to be a key feature, as it effectively transforms a complex nonlinear degradation process into a deterministic linear trajectory, simplifying long-term RUL prediction.

6.1. Computational Complexity

A key concern for real-time control is computational feasibility. We measured the execution time of the SMPC algorithm on a standard desktop workstation (Intel Core i7-10700K CPU @ 3.80 GHz, 32 GB RAM). Figure 5 shows that the average computation time grows linearly with H , remaining below 2.5 ms even for $H = 25$. This is well within the typical control cycle time (10–100 ms) for mechatronic systems. While validated on desktop hardware, the $\mathcal{O}(|A| \cdot H)$ complexity and absence of iterative solvers imply sub-10ms execution on modern flight-certified multicore processors (e.g., NXP QorIQ P4080, NXP Semiconductors, Eindhoven, The Netherlands; Xilinx Zynq UltraScale+, Xilinx Inc., San Jose, CA, USA), confirming real-time applicability for embedded flight control systems.

6.2. Computational Trade-Off: Gaussian vs. Sampling Methods

While fatigue life distributions often exhibit non-Gaussian tails (e.g., Weibull), we employ a Gaussian moment-propagation approach to prioritize real-time feasibility. Full non-Gaussian methods such as particle MPC or polynomial chaos expansion (PCE) offer higher fidelity but incur prohibitive computational costs: Our moment-propagation controller executes in approximately 1 ms per control cycle (90 model evaluations over the prediction horizon), whereas a particle MPC implementation with $N_p = 1000$ particles would require approximately 960 ms per cycle—nearly $1000\times$ slower and far exceeding the <10 ms cycle time requirement for high-rate flight control. Moreover, the Gaussian approximation is inherently conservative for bounded health states ($h \in [0, 1]$): It assigns non-zero probability to infeasible regions ($h < 0$), resulting in tighter-than-necessary chance constraints. The additional constraint tightening via α' (Section 4.2) further in-

creases this conservatism, ensuring that the actual constraint violation rate is strictly below the prescribed α . This conservative property is empirically supported by the robustness studies above: Across all Monte Carlo trials with $\pm 20\%$ parameter mismatch and under quadratic coupling mismatch, zero health constraint violations were observed, confirming that the Gaussian + α' formulation provides adequate safety margins even under significant model uncertainty.




Figure 5. Computational complexity: average execution time vs. prediction horizon. The SMPC remains computationally efficient for real-time control.

The linear scaling observed in Figure 5 results from our implementation of a receding horizon control strategy with a control Horizon $N_u = 1$. Instead of a computationally prohibitive full tree search $O(|A|^H)$, we efficiently evaluate the discrete action set $|A| = 9$ by propagating the dynamics forward over the prediction horizon H . This reduces the complexity to $O(|A| \cdot H)$, verifying the observed sub-millisecond execution times and ensuring real-time feasibility.

Remark 4. (Greedy Strategy Viability): The choice of a single-step control horizon $N_u = 1$ requires careful justification given the coupled degradation dynamics. We establish sufficient conditions for greedy optimality.

Theoretical Motivation 1 (Greedy Strategy Viability): Let the degradation dynamics be given by Equation (2) with coupling matrix $\mathbf{C}$. If the coupling satisfies the unidirectional dominance condition

$$C_{ij} > 0 \implies C_{ji} < C_{ij} \quad \forall i \neq j \quad (16)$$

then the value function $V(\mathbf{h}, t) = \max_{\pi} \mathbb{E}[\int_t^T P(\mathbf{u}, \mathbf{h}) d\tau \mid \mathbf{h}(t) = \mathbf{h}]$ is component-wise monotonically increasing in $\mathbf{h}$, i.e., $\partial V / \partial h_i > 0$ for all i .

Proof Sketch: Under unidirectional dominance, degradation in the “upstream” component (e.g., gearbox) accelerates wear in the “downstream” component (ball screw), but not vice versa. This creates a hierarchical damage propagation structure where preserving upstream health always reduces the total system’s degradation. Consequently, for any control action u , higher current health h_i leads to the following: (1) lower immediate degradation rate f_i and (2) lower coupling-induced degradation in downstream components. Both effects increase future cumulative performance, establishing the monotonicity of V . This argument provides the physical intuition for the greedy strategy’s success: In this

specific architecture, preserving “upstream” health yields monotonic benefits. While a rigorous viscosity solution proof is beyond this scope, the empirical verification in Figure 6 confirms the validity of this heuristic.

For our EMA system, the coupling matrix satisfies $C_{12} = 0.8 > C_{21} = 0.1$, confirming unidirectional dominance. Combined with the monotonic decrease in health ($dh_i \leq 0$), this establishes that any action preserving more health today cannot be suboptimal—validating $N_u = 1$. Figure 6 provides empirical verification: experiments with $N_u \in \{1, 2, 3\}$ yielded identical terminal health states, while computation time increased by approximately $80\times$ for $N_u = 3$.




Figure 6. Control horizon (N_u) verification ($T_{target} = 8000$ cycles): **(left)** lifetime distribution (blue bars)—all N_u values achieve target lifetime; the red dashed line indicates the target lifetime of 8000 cycles; **(center)** terminal ball screw health (green bars, $h \approx 0.53$)—identical across all N_u , validating greedy optimality of $N_u = 1$; **(right)** total simulation time (orange bars) showing $\sim 80\times$ increase for $N_u = 3$ due to exponential growth in action sequences (9^{N_u}).

6.3. Recursive Feasibility and Stability

For safety-critical aerospace applications, formal guarantees on controller behavior are essential. We establish recursive feasibility and stability properties of the proposed SMPC. Detailed proofs are provided in Appendix A.

Proposition 1. (Recursive Feasibility). *If the optimization problem (7) and (8) is feasible at time k with state $\mathbf{h}(k)$ and the realized disturbance satisfies $\|\xi(k)\| \leq \bar{\xi}$ (bounded noise), then the problem remains feasible at time $k + 1$.*

Proof. Let $\mathbf{u}^*(k)$ be the optimal control at time k , yielding predicted state $\bar{\mathbf{h}}(k + 1)$. The realized state is $\mathbf{h}(k + 1) = \bar{\mathbf{h}}(k + 1) + \delta(k)$, where $|\delta_i(k)| \leq \sigma_i \sqrt{\Delta t} \bar{\xi}$. Chance constraint (8) ensures that $\bar{h}_i(k + 1) \geq h_{req}(k + 1) + \Phi^{-1}(\alpha) \sqrt{\Sigma_{ii}(k + 1)}$. For $\alpha = 0.95$, the buffer of $\Phi^{-1}(0.95) \approx 1.65$ standard deviations exceeds the bounded perturbation with probability $\geq \alpha$, maintaining constraint satisfaction. Since the control set $\mathcal{U}$ is compact and degradation is monotonic ($h_i(k + 1) \leq h_i(k)$), a feasible action always exists: the minimum-degradation action u_{min} . $\square$

Proposition 2. (Practical Stability). *Define the Lyapunov candidate $V(\mathbf{h}, k) = \sum_{i=1}^N (h_i(k) - h_{req}(k))^2$. Under the SMPC policy, $\mathbb{E}[V(\mathbf{h}, k + 1) | \mathbf{h}(k)] \leq V(\mathbf{h}, k) + \epsilon_k$, where $\epsilon_k \rightarrow 0$ as $h_i(k) \rightarrow h_{req}(k)$.*

Proof Sketch: The controller minimizes tracking error $\|\bar{\mathbf{h}}_{k+1} - h_{req}(k + 1)\|^2$ (Equation (15)). Near the reference trajectory, the control actively reduces deviation. Far from the reference, performance objectives may temporarily increase deviation, but the high weighting ratio $w_{track}/w_{perf} = 100$ ensures that tracking dominates. The bounded process noise contributes $\mathbb{E}[\delta^T \delta] = \text{tr}(\mathbf{Q})$, which decreases as $f_i \rightarrow 0$ near the healthy state. Thus, the system exhibits *practical stability*: trajectories remain within a neighborhood of the reference, with the neighborhood size proportional to the noise intensity.

These properties ensure that the SMPC provides formal safety guarantees suitable for flight-critical applications, complementing the empirical validation in Section 7.

7. Extended Validation

To address the need for robust validation, we conducted three additional studies: a comparative analysis against a robust baseline, a Monte Carlo robustness study, and a parameter mismatch analysis.

Sensor and Communication Robustness: To address operational uncertainties beyond parameter mismatch, we validated the controller under (a) Gaussian sensor noise ($\sigma = 0.1$) injected into the health feedback loop and (b) feedback delays ($\tau = 50$ steps). The system achieved a 100% survival rate in both scenarios, confirming that the conservative safety margins ($\alpha = 0.95$) effectively absorb moderate sensing and latency errors.

7.1. Benchmarking Control Strategies

We compared the proposed reliability-adaptive MPC against a fixed-gain baseline:

- **Fixed-Gain:** Operates at constant 80% rated capacity, representing common industrial practice without health feedback.
- **Reliability-Adaptive MPC:** The proposed stochastic controller with chance constraints and physics-based degradation awareness.

Figure 7 shows that the fixed-gain controller fails at $t = 4841$ due to accelerated ball screw wear, while the adaptive controller survives the full mission by dynamically modulating control inputs. The three-phase behavior (conservative early operation, gradual ramp-up, and aggressive end-of-life exploitation) is clearly visible in the control action plots.




Figure 7. Comparative analysis: fixed-gain vs. reliability-adaptive MPC. The fixed-gain controller (top) operates at a constant 80% capacity and fails at $t = 4841$ due to ball screw wear. The adaptive controller (bottom) modulates control inputs based on the health state, surviving the full 8000-cycle mission while achieving 12% higher cumulative performance.

Remark 5. (Robustness and Sensitivity Analysis):

Robustness to Identification Uncertainty: In practical deployment, degradation parameters are subject to identification errors. To assess robustness against these uncertainties, we introduced a

systematic parameter mismatch where the plant's degradation rate parameters (Paris's law coefficient C_{paris} and Archard wear coefficient k_{archard}) differed from the controller's nominal values by $\pm 20\%$. Monte Carlo validation ($N = 50$) confirmed 100% mission success across the full mismatch range. However, the terminal health margin decreased with increasing positive mismatch (where the plant degrades faster than the controller expects): from $h_{\text{final}} \approx 0.61$ at nominal to $h_{\text{final}} \approx 0.44$ at $+20\%$ mismatch. This demonstrates the graceful degradation of safety margins rather than abrupt failure. The receding-horizon feedback mechanism compensates for the model drift by replanning at each time step using the measured plant state, correcting the trajectory even when the internal model underestimates wear rates.

Sensitivity to Coupling Non-Linearity: To address concerns regarding the linear coupling approximation (Equation (2)), we conducted a sensitivity analysis where the simulation plant utilized a quadratic coupling model ($C_{ij} \propto (1 - h)^2$). Unlike the linear model, the quadratic form is convex in damage: Coupling-induced load amplification grows superlinearly as health decreases, causing the controller—which internally assumes linear coupling—to progressively underestimate the true coupling strength in the later mission phase. As a result, the controller completed the full mission (Life > 8000 cycles) but with significantly reduced terminal health ($h_{\text{final}} \approx 0.14$ vs. 0.70 under nominal linear coupling). This 80% reduction in safety margin reflects the controller consuming its “reliability budget” to compensate for the unmodeled nonlinearity—a form of graceful degradation rather than abrupt failure. Importantly, $h_{\text{final}} = 0.14$ remains above the structural failure threshold ($h = 0$), confirming that the closed-loop feedback mechanism prevents catastrophic failure even under severe model mismatch. This result demonstrates that while the linear assumption is adequate for the nominal operating regime (where $h > 0.6$ for 90% of the mission), robustness margins should be considered when coupling nonlinearity is suspected in the terminal phase.

7.2. Value of Coupling Knowledge

To isolate the contribution of physics-based coupling awareness, we conducted an ablation study under amplified coupling conditions (vibration transfer coefficient increased to $3 \times$ nominal, simulating misalignment or lubrication failure). This stress test deliberately creates a more challenging scenario than the main case study.

The results, shown in Figure 8, demonstrate that under nominal conditions, both controllers survive—the system has a sufficient design margin that even a coupling-unaware controller can complete the mission. However, under amplified coupling with accelerated wear, the decoupled controller—which ignores gearbox-to-screw vibration transfer—fails to anticipate cascading damage. It maintains aggressive actuation, leading to premature failure. The coupled controller, anticipating the interaction, proactively modulates the load to mitigate vibration-induced wear, extending useful life by approximately 77.6%. This confirms that the gain is not merely from the MPC algorithm but specifically from the physics-based awareness of component interactions.

Remark 6. (Practical Applicability): This finding has important implications for deployment. The coupling-aware model provides graceful degradation under off-nominal conditions—precisely when robustness matters most. In nominal operations, the additional modeling complexity incurs no penalty (both approaches succeed). However, during abnormal conditions such as misalignment, lubrication degradation, or thermal excursions, the coupling-aware controller provides a critical safety margin. This “insurance” property justifies the modeling investment for safety-critical aerospace applications where off-design conditions are inevitable over a system's lifecycle.




Figure 8. Ablation study (amplified coupling, $3\times$ nominal): Under strong coupling conditions simulating off-nominal operations, the coupled controller (green) extends mission life by 77.6% compared to the decoupled baseline (red), which ignores interaction physics. **(Left):** Ball screw health trajectories showing the earlier failure of the decoupled approach. **(Right):** Lifetime comparison (decoupled: 4175 cycles; coupled: 7414 cycles).

8. Conclusions

This paper presented a reliability-adaptive control framework for coupled aerospace EMAs. By integrating stochastic MPC with a rigorous six-layer architecture, we demonstrated its ability to actively manage reliability under uncertainty. The results confirm that health-aware control can significantly extend a system's life and enhance sortie generation, offering a viable path towards self-optimizing flight control systems.

The key contributions are: (1) a physics-based coupled degradation model incorporating Paris's and Archard's laws with thermal–vibration interactions; (2) a chance-constrained SMPC formulation that maintains prescribed reliability targets under stochastic disturbances; and (3) comprehensive validation demonstrating 65% life extension compared to fixed-operation baselines, with real-time computational feasibility.

Future Work

Future research will focus on scaling this framework to fleet-level implementation, where multiple distributed aircraft coordinate to optimize global mission objectives. Integrating distributed SMPC with consensus protocols could enable coordinated maintenance scheduling across the fleet, maximizing overall readiness. Additionally, exploring online parameter estimation for the coupling matrix $\mathbf{C}$ via extended Kalman filtering would further enhance robustness against unmodeled degradation dynamics.

These extensions would advance the state of the art in prescriptive maintenance, moving closer to fully autonomous, self-optimizing aerospace ecosystems.

Scalability to Multi-Actuator Systems: The proposed hierarchical architecture naturally extends to full aircraft-level control allocation. In a multi-actuator setting, each EMA runs a local SMPC instance, estimating its own health and computing dynamic capability bounds $[\underline{u}(t), \bar{u}(t)]$ that satisfy its local health budget. The Mission Level (Level 1) then acts as a health-aware control allocator, distributing the total force demand $\mathbf{F}_{cmd}$ among N actuators by solving $\min \|\mathbf{B}\mathbf{u} - \mathbf{F}_{cmd}\|^2$ subject to $\underline{u}_i \leq u_i \leq \bar{u}_i$. Since the computationally intensive SMPC runs in parallel on distributed controllers, the system latency scales as $O(1)$, making the approach viable for large-scale fleet deployment.

Regarding observability, the framework assumes that health states are estimable via EKF; in practice, this requires assessing sensor reliability and calculating observability Gramians for the chosen sensor suite, which is a subject of our ongoing flight-test validation. Finally, regarding certification of adaptive controllers, the recursive feasibility guarantee

(Proposition 2) provides the deterministic safety bounds required to satisfy airworthiness standards, offering a predictable fallback behavior even under stochastic uncertainty.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/math14040737/s1>.

Funding: This research received no external funding.

Data Availability Statement: The simulation code (Python) files used to generate all figures and tables in this manuscript are provided in the Supplementary Materials. The implementation includes degradation models, controller algorithms, and figure generation scripts to enable the full reproducibility of the reported results.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Formal Proofs

Appendix A.1. Proof Sketch for Theoretical Motivation 1 (Value Function Monotonicity)

Statement: Under unidirectional dominance ($C_{ij} > 0 \Rightarrow C_{ji} < C_{ij}$), the value function $V(\mathbf{h}, t) = \max_{\pi} \mathbb{E}[\int_t^T P(\mathbf{u}, \mathbf{h}) d\tau \mid \mathbf{h}(t) = \mathbf{h}]$ satisfies $\frac{\partial V}{\partial h_i} > 0$ for all i .

Proof Sketch: We establish monotonicity through structural analyses of the HJB equation, supported by numerical verification.

Step 1 (Regularity Assumptions): We assume that the value function $V(\mathbf{h}, t)$ is continuously differentiable in the interior of the domain $\mathcal{H} = (0, 1]^N \times [0, T)$. This holds when (i) the drift f_i and diffusion σ_i are Lipschitz continuous in $\mathbf{h}$, (ii) the control set $\mathcal{U}$ is compact, and (iii) the performance function P is continuous. Our parameterization satisfies these conditions for $h_i > 0$.

Step 2 (HJB Structure): Under the regularity assumption, V satisfies the HJB equation:

$$-\frac{\partial V}{\partial t} = \max_{\mathbf{u} \in \mathcal{U}} \left\{ P(\mathbf{u}, \mathbf{h}) - \sum_i f_i(\mathbf{u}, \mathbf{h}) \frac{\partial V}{\partial h_i} + \frac{1}{2} \sum_i \sigma_i^2(\mathbf{h}) \frac{\partial^2 V}{\partial h_i^2} \right\} \quad (\text{A1})$$

with terminal condition $V(\mathbf{h}, T) = 0$.

Step 3 (Boundary Analysis): At $t = T - \epsilon$ (small $\epsilon > 0$), Taylor expansion gives $V(\mathbf{h}, T - \epsilon) \approx \epsilon \cdot \max_{\mathbf{u}} P(\mathbf{u}, \mathbf{h})$. Since performance $P(\mathbf{u}, \mathbf{h})$ is non-decreasing in each h_i (healthier components deliver higher output), we have $\frac{\partial V}{\partial h_i} > 0$ near T .

Step 4 (Monotonicity Propagation): From Equation (2), the cross-partial satisfies $\frac{\partial f_j}{\partial h_i} = -C_{ji} f_{j, \text{base}} \leq 0$ for $i \neq j$. Under unidirectional dominance ($C_{ij} > C_{ji}$), the “upstream” component’s health reduces downstream degradation more than vice versa. Differentiating the HJB equation with respect to h_i , we obtain a linear PDE for $w_i = \frac{\partial V}{\partial h_i}$ with non-negative source terms. By the maximum principle for parabolic PDEs, if $w_i > 0$ at the terminal boundary, it remains positive throughout the domain.

Step 5 (Numerical Verification): We verified monotonicity empirically by computing $V(\mathbf{h}, t)$ via backward dynamic programming on a discretized state space (50^3 grid points). For all sampled states with $h_i \in [0.1, 1.0]$ and $t \in [0, T]$, finite differences confirmed $\frac{V(\mathbf{h} + \delta \mathbf{e}_i, t) - V(\mathbf{h}, t)}{\delta} > 0$ with $\delta = 0.01$. $\square$

Remark A1. The above proof assumes classical (smooth) solutions to the HJB equation. For a fully rigorous treatment using viscosity solutions, additional verification of semicontinuity and the comparison principle would be required; see [25], Chapter IV. Our numerical verification confirms that monotonicity holds for the specific EMA system parameters.

Appendix A.2. Proof of Proposition 1 (Recursive Feasibility)

Statement: If the MPC problem is feasible at time k , it remains feasible at $k + 1$ under bounded disturbances $\|\tilde{\xi}(k)\| \leq \bar{\xi}$.

Proof. Let $\mathcal{F}(k) \subset \mathbb{R}^N$ denote the feasible health region at time k :

$$\mathcal{F}(k) = \{\mathbf{h} : \exists \mathbf{u} \in \mathcal{U}, h_i - z_\alpha \sigma_i \sqrt{\Delta t} \geq h_{req}(k), \forall i\} \quad (\text{A2})$$

where $z_\alpha = \Phi^{-1}(\alpha)$. We show that $\mathbf{h}(k+1) \in \mathcal{F}(k+1)$ with probability $\geq \alpha$.

From the dynamics, $h_i(k+1) = h_i(k) - f_i(\mathbf{u}^*, \mathbf{h}(k))\Delta t + \sigma_i \sqrt{\Delta t} \tilde{\xi}_i(k)$

The chance constraint at time k ensures the following:

$$\bar{h}_i(k+1) \geq h_{req}(k+1) + z_\alpha \sqrt{\Sigma_{ii}(k+1)} \quad (\text{A3})$$

where $\bar{h}_i$ is the mean prediction. The realized state satisfies $|h_i(k+1) - \bar{h}_i(k+1)| \leq \sigma_i \sqrt{\Delta t} \tilde{\xi}$ with probability 1.

For the set intersection, $\mathcal{F}(k) \cap \mathcal{F}(k+1) \neq \emptyset$ because the following is the case:

1. The target trajectory $h_{req}(k)$ is monotonically decreasing and continuous.
2. The minimum-degradation action $\mathbf{u}_{min} = \arg \min_{\mathbf{u}} \sum_i f_i(\mathbf{u}, \mathbf{h})$ always exists (compact $\mathcal{U}$).
3. Under $\mathbf{u}_{min}$, the health decay rate is minimized, ensuring that the state remains above $h_{req}(k+1)$

Thus, feasibility is maintained inductively. $\square$

References

1. Cao, W.; Mecrow, B.C.; Atkinson, G.J.; Bennett, J.W.; Atkinson, D.J. Overview of Electric Motor Technologies Used for More Electric Aircraft (MEA). *IEEE Trans. Ind. Electron.* **2012**, *59*, 3523–3531. <https://doi.org/10.1109/TIE.2011.2165453>.
2. Hussain, Y.M.; Burrow, S.G.; Keogh, P.S. A Review of Techniques to Mitigate Jamming in Electromechanical Actuators for Safety Critical Applications. *Int. J. Progn. Health Manag.* **2018**, *9*, 1–13. <https://doi.org/10.36001/IJPHM.2018.V9I3.2749>.
3. Salazar, J.C.; Weber, P.; Nejari, F.; Sarrate, R.; Theilliol, D. System Reliability Aware Model Predictive Control Framework. *Reliab. Eng. Syst. Saf.* **2017**, *167*, 663–672.
4. Yin, Z.; Hu, N.; Chen, J.; Yang, Y.; Shen, G. A review of fault diagnosis, prognosis and health management for aircraft electromechanical actuators. *IET Electr. Power Appl.* **2022**, *16*, 1249–1272. <https://doi.org/10.1049/elp2.12225>.
5. Yang, N.; Gu, C.; Huang, Y.; Wen, P.; Zhao, S.; Feng, S.; Chen, S. Remaining Useful Life Prediction of Electromechanical Equipment based on Particle Filter and LSTM. In *2022 IEEE International Conference on Sensing, Diagnostics, Prognostics, and Control (SDPC)*; IEEE: New York, NY, USA, 2022; pp. 87–91. <https://doi.org/10.1109/SDPC55702.2022.9915827>.
6. Raja, M.S.K.; Ali, Q. Recent advances in active fault tolerant flight control systems. *Proc. Inst. Mech. Eng. Part G J. Aerosp. Eng.* **2022**, *236*, 2151–2161. <https://doi.org/10.1177/09544100211062812>.
7. Zhao, Y.; Yang, J.; Wang, W.; Yang, H.; Niyato, D. TranDRL: A Transformer-Driven Deep Reinforcement Learning Enabled Prescriptive Maintenance Framework. *IEEE Internet Things J.* **2023**, *11*, 35432–35444. <https://doi.org/10.1109/JIOT.2024.3436110>.
8. Luettig, B.; Akhiat, Y.; Daw, Z. ML meets aerospace: Challenges of certifying airborne AI. *Front. Aerosp. Eng.* **2024**, *3*, 1475139. <https://doi.org/10.3389/fpace.2024.1475139>.
9. Haghighat, S.; Liu, H.H.T.; Martins, J.R.R.A. Model-Predictive Gust Load Alleviation Controller for a Highly Flexible Aircraft. *J. Guid. Control. Dyn.* **2012**, *35*, 1751–1766.
10. Mayne, D.Q.; Seron, M.M.; Raković, S.V. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica* **2005**, *41*, 219–224. <https://doi.org/10.1016/j.automatica.2004.08.019>.
11. Garone, E.; Nicotra, M.M. Explicit Reference Governor for Constrained Nonlinear Systems. *IEEE Trans. Autom. Control* **2016**, *61*, 1379–1384. <https://doi.org/10.1109/TAC.2015.2476195>.
12. Mammarella, M.; Lorenzen, M.; Capello, E.; Park, H.; Dabbene, F.; Guglieri, G.; Romano, M.; Allgöwer, F. An Offline-Sampling SMPC Framework With Application to Autonomous Space Maneuvers. *IEEE Trans. Control Syst. Technol.* **2020**, *28*, 388–402. <https://doi.org/10.1109/TCST.2018.2879938>.
13. Tedesco, F.; Akram, W.; Forouzanfar, M.; Parisini, T. Predictive maintenance of actuators in linear systems: A receding horizon set-theoretic approach. *Int. J. Robust Nonlinear Control* **2022**, *32*, 6395–6411. <https://doi.org/10.1002/rnc.6148>.

14. Forouzanfar, M.; Gagliardi, G.; Tedesco, F.; Casavola, A. Integrated Model-Based Control Allocation Strategies Oriented to Predictive Maintenance of Saturated Actuators. *IEEE Trans. Autom. Sci. Eng.* **2025**, *22*, 1045–1056.
15. Liu, X.; Wu, Z.; Yuan, Y.; Sun, W.; Gao, J. Reliability Modeling of Wind Turbine Gearbox System Considering Failure Correlation Under Shock–Degradation. *Sensors* **2025**, *25*, 4425. <https://doi.org/10.3390/s25144425>.
16. Mazzoleni, M.; Di Rito, G.; Previdi, F. Fault Diagnosis and Condition Monitoring of Aircraft Electro-Mechanical Actuators. In *Electro-Mechanical Actuators for the More Electric Aircraft*; Bilal, G.A., Ed.; Springer International Publishing: Cham, Switzerland, 2021; pp. 119–224. https://doi.org/10.1007/978-3-030-61799-8_4.
17. Shi, R.; Lü, Y.; Zhang, Y.; Meng, J.; Chen, R.; Hu, R. Remaining Useful Life Estimation with Coupling Dual Random Effects Based on Nonlinear Wiener Process. *Eksplot. Niezawodn. Maint. Reliab.* **2025**, *27*, 4. <https://doi.org/10.17531/ein/202908>.
18. Mrugalska, B.; Stetter, R. Health-Aware Model-Predictive Control of a Cooperative AGV-Based Production System. *Sensors* **2019**, *19*, 532. <https://doi.org/10.3390/s19030532>.
19. Salazar, J.C.; Nejari, F.; Sarrate, R. Reliable control of a twin rotor MIMO system using actuator health monitoring. In *22nd Mediterranean Conference on Control and Automation*; IEEE: Palermo, Italy, 2014; pp. 481–486.
20. Ray, A.; Tangirala, S.; Phoha, S. Stochastic modeling of fatigue crack propagation. In *Proceedings of the 1997 American Control Conference*; IEEE: Albuquerque, NM, USA, 1997; Volume 1, pp. 425–429.
21. Fouvry, S.; Liskiewicz, T.; Kapsa, P. An Energy Description of Wear Mechanisms and Its Applications to Oscillating Sliding Contacts. *Wear* **2003**, *255*, 287–298.
22. Paris, P.C.; Erdogan, F. A Critical Analysis of Crack Propagation Laws. *J. Basic Eng.* **1963**, *85*, 528–533. <https://doi.org/10.1115/1.3656900>.
23. Archard, J.F. Contact and Rubbing of Flat Surfaces. *J. Appl. Phys.* **1953**, *24*, 981–988. <https://doi.org/10.1063/1.1721448>.
24. Lu, Y.; Morris, K.; Frechette, S. *Current Standards Landscape for Smart Manufacturing Systems*; NISTIR 8107; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2016. <https://doi.org/10.6028/nist.ir.8107>.
25. Fleming, W.H.; Soner, H.M. *Controlled Markov Processes and Viscosity Solutions*, 2nd ed.; Springer: New York, NY, USA, 2006. <https://doi.org/10.1007/0-387-31071-1>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.