

EXTENDED TECHNICAL COMPANION TO

HARDWARE-VERIFIABLE RECURSIVE ADMISSIBILITY FOR AUTONOMOUS AEROSPACE AND MECHANICAL SYSTEMS

Complete Mechanical Benchmark Architecture, Computational Derivation,
Recursive Admissibility, Digital-Twin Validation, and Engineering-Grade
Reproducibility Record

INDEPENDENT EXTENDED TECHNICAL COMPANION

This document is not the 20-page MEAP 2026 conference submission. It provides the expanded technical architecture, derivations, computational validation record, figures, limitations, and reproducibility material supporting the condensed conference manuscript.

Jered McClain - Erydir Ceisiwr

SEXA Institute of Technology and Interdomain Galactic Advisors

SEXAtech.com · CreateNewLand.com · UnifiedFieldSystems.com

Technical companion / supplement to the MEAP 2026 conference manuscript

9th International Conference on Mechanical Engineering & Applications (MEAP 2026) — Hardware-Verifiable Recursive Admissibility for Autonomous Aerospace and Mechanical Systems: A SEXA Framework for Constraint-Safe State Selection, Fault Detection, Digital-Twin Reconstruction, Recursive Dimensional Compression, and Engineering-Grade Computational Validation

Jered McClain · Erydir Ceisiwr

SEXA Institute of Technology and Interdomain Galactic Advisors

MEAP 2026 — October 30–31, 2026 — Virtual Conference

Abstract

Autonomous aerospace and mechanical systems increasingly require computational architectures capable not merely of predicting system behavior, but of determining whether candidate operating states satisfy ordered physical, structural, operational, and safety constraints before those states are admitted for execution. This paper introduces a mechanical-engineering implementation of recursive admissibility derived from the SEXA computational framework and investigates its use as a deterministic supervisory architecture for constraint-safe state selection, fault detection, and digital-twin reconstruction.

The proposed architecture extends previous work in which the SEXA framework was organized as a hardware-verifiable recursive computational system characterized by deterministic execution, traceable intermediate states, dimensional reduction, computational correspondence, and explicit failure conditions. The present investigation moves beyond that computational validation problem by constructing an engineering test environment in which measurable mechanical state variables are subjected to an ordered Γ -admissibility chain before acceptance into the operational state space.

Let the instantaneous engineering state be represented by

$$X_t = [T_t, P_t, V_t, a_t, \omega_t, \sigma_t, \epsilon_t, q_t, E_t, F_t]^T,$$

where T denotes temperature, P pressure, V vibration amplitude, a acceleration, ω rotational state, σ mechanical stress, ϵ strain, q actuator or control state, E energetic state, and F fault-state information.

Rather than permitting every generated or measured state to propagate directly into the control architecture, the state is evaluated through an ordered admissibility operator

$$\Gamma = \gamma_6 \circ \gamma_5 \circ \gamma_4 \circ \gamma_3 \circ \gamma_2 \circ \gamma_1.$$

The resulting engineering admissibility decision is defined by

$$A(X_t) = \begin{cases} 1, & \Gamma(X_t) = X_t, \\ 0, & \Gamma(X_t) \neq X_t. \end{cases}$$

Thus, a candidate mechanical state is admitted only when it survives the complete ordered constraint chain. This directly adapts the corpus's defining principle that admissibility precedes evaluation; in the full-stack SEXA formulation, the Γ chain acts before energetic evaluation and only Γ -invariant configurations contribute to the resulting state.

For engineering implementation, the six stages are interpreted as sequential constraint classes governing **existence, flow/transition validity, manifold or model consistency, symmetry/invariance, recursive persistence, and realization/observation**. The objective is not to assert that these operators have already been experimentally established as physical laws. Rather, they are investigated here as a deterministic engineering constraint architecture whose state transitions and failures can be measured computationally.

A digital-twin benchmark is then constructed in which nominal aerospace/mechanical trajectories are deliberately subjected to sensor drift, thermal excursions, excessive vibration, stress-limit violations, actuator faults, inconsistent state transitions, and compound faults. For every injected disturbance, the system records the first Γ stage at which admissibility fails, the rejected state, the surviving reduced state representation, and the reconstruction trace.

The central hypothesis tested is therefore:

The central hypothesis tested is therefore:

Unsafe or inconsistent state $\xrightarrow{\Gamma}$ explicit rejection

while

constraint-consistent state $\xrightarrow{\Gamma}$ admissible operational state.

The experiment evaluates deterministic repeatability, fault-detection rate, false-admission rate, false-rejection rate, constraint-stage localization, reconstruction consistency, computational latency, and state-trace reproducibility. These quantities establish measurable engineering criteria through which the proposed architecture can succeed or fail rather than relying upon qualitative interpretation.

This engineering formulation is motivated by the existing SEXA dimensional architecture, which distinguishes a higher-dimensional structural representation from a bounded operational representation and treats dimensional reduction as an admissibility-governed computational process. The corresponding computational architecture was previously developed as a deterministic and traceable framework intended to support future software simulation, formal verification, embedded computation, and hardware-assisted implementation. The predecessor work explicitly identifies reproducible execution, observable intermediate states, and identifiable failure conditions as necessary properties of a hardware-verifiable recursive architecture.

The contribution of the present investigation is therefore not the introduction of another governing equation and not a claim of completed physical verification of the wider SEXA Unified Field Framework. Instead, this work isolates recursive admissibility as an engineering mechanism and subjects it to a defined aerospace and mechanical benchmark containing measurable state variables, known fault conditions, deterministic constraint evaluation, and quantitative performance criteria.

The resulting framework provides a testable pathway for investigating whether ordered recursive admissibility can function as a supervisory computational layer for autonomous aerospace systems, mechanical digital twins, structural-health monitoring, predictive maintenance, fault-tolerant control, and future embedded or hardware-assisted mechanical systems.

Keywords: Recursive Admissibility; Autonomous Aerospace Systems; Mechanical Systems; Fault Detection; Digital Twin; Structural Health Monitoring; Predictive Maintenance; Constraint-Safe State Selection; Hardware Verification; Fault-Tolerant Control; Recursive Computing; SEXA; Γ -Admissibility; Dimensional Reduction.

1. INTRODUCTION

Modern aerospace and mechanical systems increasingly operate through continuous interaction between physical machinery and computational decision architectures. Aircraft, spacecraft, unmanned aerial vehicles, robotic mechanisms, propulsion systems, rotating machinery, autonomous vehicles, manufacturing platforms, and structural-health-monitoring systems can generate large quantities of sensor information describing their instantaneous physical condition. Temperature, pressure, vibration, acceleration, rotational velocity, mechanical stress, strain, actuator position, power consumption, and diagnostic variables may be measured continuously and incorporated into control systems or digital-twin representations.

The existence of these measurements, however, does not by itself establish that the resulting computational representation describes an admissible mechanical state.

A temperature measurement may remain within its individual operating limit while becoming incompatible with the simultaneously measured pressure state. A rotational velocity may remain below its nominal maximum while the associated vibration spectrum indicates an unstable operating regime. An actuator command may independently satisfy its displacement constraint while producing a transition that is inconsistent with the mechanical state inherited from the preceding time step. Likewise, multiple sensors may individually report

numerically permissible measurements while their combined configuration violates a structural, dynamical, energetic, or operational relationship required by the physical system.

The engineering problem considered in this paper is therefore not limited to state estimation or threshold-based fault detection. It concerns state admissibility: whether a candidate mechanical state should be permitted to enter or remain within the operational state space of an autonomous system.

Let the complete candidate state space of a mechanical system be denoted by

$$\Omega,$$

and let the admissible subset be

$$\Omega_A \subseteq \Omega.$$

For an instantaneous engineering state X_t , the supervisory decision can then be expressed as

$$\mathcal{A} : \Omega \rightarrow \{0, 1\},$$

with

$$\mathcal{A}(X_t) = \begin{cases} 1, & X_t \in \Omega_A, \\ 0, & X_t \notin \Omega_A. \end{cases}$$

The fundamental difficulty is determining membership in Ω_A when admissibility depends upon multiple interacting constraints rather than a single scalar threshold.

Conventional monitoring can represent a collection of engineering constraints as

$$\mathcal{C} = \{C_1, C_2, \dots, C_n\},$$

where each constraint produces an individual decision

$$C_i(X_t) \in \{0, 1\}.$$

Such a representation is useful for detecting direct limit violations. However, the present investigation asks a more specific question: whether the constraints can be evaluated as an ordered recursive process in which the state surviving one stage becomes the state presented to the next stage, while the complete evaluation history remains computationally observable.

This produces the ordered state sequence

$$X_t^{(0)} \xrightarrow{\gamma_1} X_t^{(1)} \xrightarrow{\gamma_2} X_t^{(2)} \xrightarrow{\gamma_3} X_t^{(3)} \xrightarrow{\gamma_4} X_t^{(4)} \xrightarrow{\gamma_5} X_t^{(5)} \xrightarrow{\gamma_6} X_t^{(6)}.$$

Here,

$$X_t^{(i)} = \gamma_i(X_t^{(i-1)}), \quad i = 1, \dots, 6,$$

subject to successful admissibility at every preceding stage.

The engineering significance of this organization is that rejection becomes traceable. Instead of returning only a global fault flag, the architecture can retain the first admissibility stage at which the candidate state fails:

$$k_{\text{fail}} = \min \left\{ i \in \{1, \dots, 6\} : \gamma_i \left(X_t^{(i-1)} \right) \notin \Omega_i \right\},$$

where Ω_i denotes the admissible domain associated with stage i .

The value k_{fail} therefore becomes a diagnostic quantity. It identifies the earliest point in the ordered evaluation at which the mechanical state ceases to satisfy the prescribed architecture.

This construction follows directly from the computational direction established in the preceding SEXA hardware-verification work. That work identifies deterministic inputs and outputs, observable intermediate computational states, and explicit failure conditions as requirements for engineering implementation. It also identifies software implementation, formal verification, quantitative performance measurement, and eventual engineering realization as subsequent stages of investigation.

The present work advances into that next stage by introducing a controlled mechanical digital-twin environment in which nominal trajectories and deliberately corrupted trajectories can be evaluated against known ground truth.

The distinction between the present investigation and the preceding computational architecture is important. The earlier work asks whether recursive admissibility and dimensional reduction can be organized into deterministic and traceable computational operations. The present work asks whether an engineering implementation of those operations can correctly distinguish admissible mechanical states from deliberately introduced fault states.

Accordingly, four objects are known for every benchmark trial: the original nominal state, the injected disturbance, the expected admissibility classification, and the output produced by the Γ -admissibility architecture. This allows direct calculation of true admissions, true rejections, false admissions, and false rejections.

Let y_t denote the ground-truth classification and $\hat{y}_t$ the classification returned by the recursive admissibility architecture:

$$y_t, \hat{y}_t \in \{0, 1\}.$$

The central classification test is therefore

$$\hat{y}_t \stackrel{?}{=} y_t.$$

Failure of this equality is not reinterpreted as a successful result. It constitutes a benchmark error and enters the corresponding false-admission or false-rejection measurement.

This requirement is essential to the engineering character of the paper. The wider SEXA corpus distinguishes structural or computational survivability from experimental physical confirmation; the existing hardware-oriented work explicitly states that computational consistency does not establish every associated physical interpretation. The present benchmark preserves that distinction.

Consequently, successful benchmark performance would support the narrower proposition that ordered recursive admissibility can operate as a reproducible engineering constraint architecture under the conditions tested. It would not, by itself, establish the broader physical claims of the SEXA Unified Field Framework.


The primary research question of this investigation is therefore:

Can an ordered, hardware-verifiable recursive admissibility architecture distinguish nominal from faulty aerospace and mechanical states while preserving deterministic execution, explicit fault localization, computational traceability, and reconstructable digital-twin state information?

The remainder of this paper develops the Γ -admissibility engineering formulation, defines the mechanical digital-twin benchmark, specifies the controlled fault classes, establishes quantitative performance and falsification criteria, and evaluates whether the resulting architecture satisfies those requirements.


The theoretical corpus dictates the architecture; the benchmark proves the implementation

Stage 1: Foundational
SEXA Corpus.




Theoretical Architecture: 2880D
Unified Field Theory, SREF
Dynamics, Gamma-Glyph Closure.

Stage 2: VLSIE 2026.



Hardware-Verifiable Computation:
Deterministic execution, observable
intermediate states.


Stage 3: Current MEAP
2026 Manuscript.



Mechanical Engineering Implementation:
Digital-Twin reconstruction, engineering-
grade computational benchmark.

Prerequisite Corpus Boundary: Every technical relationship shown is backed by its controlling source.
The current benchmark evaluates the mechanical implementation, not the foundational physics.

THE SEXA INTER-DOMAIN VALIDATION PROTOCOL



The underlying mathematical corpus is a prerequisite reference set.
The focus of this audit is strictly the Stage V Mechanical Benchmark.

2. RECURSIVE ADMISSIBILITY AS AN ENGINEERING CONSTRAINT ARCHITECTURE

2.1 From Computational Admissibility to Mechanical Admissibility

The engineering architecture developed in this paper begins with a distinction between a candidate state and an admissible state. A digital twin, sensor network, estimator, or autonomous controller may generate a candidate state X_t , but generation of that state does not establish that the state satisfies the complete constraint structure of the physical system.

The SEXA computational architecture provides the ordering principle used here. In the preceding hardware-verification formulation, recursive admissibility operates as an ordered computational constraint before evaluation, with deterministic execution, traceable intermediate states, reproducible behavior, and identifiable failure conditions forming the principal requirements for future engineering implementation.

12_VLSIE_2026_HARDWARE_VERIFIAB...

The present implementation converts that computational principle into an engineering state-selection problem.

Let

$$X_t = [x_{1,t}, x_{2,t}, \dots, x_{n,t}]^T$$

represent a candidate mechanical state at discrete time t . For the benchmark developed here, the state variables are instantiated as

$$X_t = [T_t, P_t, V_t, a_t, \omega_t, \sigma_t, \epsilon_t, q_t, E_t, F_t]^T.$$

Here T_t is temperature, P_t pressure, V_t vibration amplitude, a_t acceleration, ω_t rotational state, σ_t mechanical stress, ϵ_t strain, q_t actuator or commanded mechanical state, E_t energetic state, and F_t diagnostic or fault-state information.

These quantities do not share common physical dimensions. Consequently, dimensional quantities are retained in their original engineering units for physical interpretation while a normalized representation is constructed for computational comparison.

For each variable $x_{j,t}$,

$$z_{j,t} = \frac{x_{j,t} - \mu_j}{s_j},$$

where μ_j denotes the reference operating value and $s_j > 0$ denotes the normalization scale associated with that variable.

The normalized state is therefore

$$Z_t = [z_{1,t}, z_{2,t}, \dots, z_{n,t}]^T.$$

Normalization is not itself an admissibility decision. It is a computational transformation that permits heterogeneous mechanical quantities to enter a common evaluation architecture without discarding their dimensional source values.

2.2 Ordered Γ -Admissibility

The complete engineering admissibility operator is represented as

$$\Gamma = \gamma_6 \circ \gamma_5 \circ \gamma_4 \circ \gamma_3 \circ \gamma_2 \circ \gamma_1.$$

The ordering is significant. The Γ -SubThreshold Fold formulation in the source corpus defines the six-operator architecture as an ordered admissibility chain rather than an arithmetic interpretation of “six into five.”  06_GAMMA_SUBTHRESHOLD_FOLD_LOGI... The engineering implementation therefore preserves operator ordering rather than replacing Γ with six independent, simultaneously evaluated thresholds.

Beginning with

$$X_t^{(0)} = X_t,$$

each surviving state is passed forward according to

$$X_t^{(i)} = \gamma_i \left(X_t^{(i-1)} \right), \quad i = 1, \dots, 6.$$

For engineering implementation, associate each stage with an admissibility indicator

$$a_i(X_t) \in \{0, 1\}.$$

The complete state decision becomes

$$A(X_t) = \prod_{i=1}^6 a_i(X_t).$$

Therefore,

$$A(X_t) = 1$$

only when

$$a_1 = a_2 = a_3 = a_4 = a_5 = a_6 = 1.$$

If any required stage returns zero,

$$A(X_t) = 0.$$

The architecture consequently implements logical conjunction across an ordered evaluation path while retaining the intermediate result of every stage.

2.3 Engineering Interpretation of the Six Γ Stages

The underlying SEXA corpus associates the ordered Γ structure with sequential admissibility rather than arbitrary independent operations. The present paper introduces a mechanical interpretation of those stages so that the architecture can be subjected to engineering testing.

The six engineering constraint classes are defined as follows.

γ_1 : state existence and acquisition validity.

This stage determines whether the candidate state contains the information required for further evaluation. Missing sensor values, undefined quantities, nonfinite numerical values, acquisition failures, corrupted fields, and measurements outside the valid sensor domain are rejected at this stage.

For variable x_j , let the acquisition domain be

$$D_j = [x_{j,\min}^{\text{sensor}}, x_{j,\max}^{\text{sensor}}].$$

Then

$$a_1(X_t) = \prod_{j=1}^n I(x_{j,t} \in D_j),$$

where $I(\cdot)$ denotes the indicator function.

This constraint should not be confused with the safe operating envelope. A measurement can be validly acquired while representing an unsafe mechanical condition. The purpose of γ_1 is to determine whether the state is computationally and instrumentally valid enough to be evaluated.

γ_2 : transition and dynamical validity.

A mechanically possible value at time t may nevertheless be incompatible with the preceding state. The second stage therefore evaluates temporal transition consistency.

Let

$$\Delta X_t = X_t - X_{t-1}.$$

For a sampling interval Δt , define the discrete state-rate vector

$$\dot{X}_t \approx \frac{X_t - X_{t-1}}{\Delta t}.$$

If $r_j^{\max}$ denotes the maximum permitted rate of change for state component j , then the transition condition is

$$\left| \frac{x_{j,t} - x_{j,t-1}}{\Delta t} \right| \leq r_j^{\max}.$$

Accordingly,

$$a_2(X_t) = \prod_{j=1}^n I\left(\left| \frac{x_{j,t} - x_{j,t-1}}{\Delta t} \right| \leq r_j^{\max}\right).$$

This stage provides a direct mechanism for detecting abrupt sensor jumps, impossible actuator transitions, nonphysical temperature or pressure changes, excessive acceleration changes, and other temporally inconsistent state transitions.

γ_3 : coupled mechanical-state consistency.

The third stage evaluates relationships that cannot be established by inspecting individual state variables independently.

Let

$$g_k(X_t) = 0$$

represent an expected coupled mechanical relationship and let δ_k denote the permitted residual tolerance. Define

$$r_{k,t} = g_k(X_t).$$

The corresponding coupled-state requirement is

$$|r_{k,t}| \leq \delta_k.$$

For m coupled constraints,

$$a_3(X_t) = \prod_{k=1}^m I(|g_k(X_t)| \leq \delta_k).$$

This stage is essential to the proposed architecture because two or more individually acceptable measurements may collectively describe an inconsistent mechanical state.

For example, a rotor may exhibit a rotational velocity that independently remains inside its permissible range and a vibration amplitude that independently remains below an absolute shutdown threshold, while their coupled relationship differs substantially from the nominal rotor-response envelope. The state can therefore survive γ_1 and γ_2 yet fail γ_3 .

γ_4 : invariant and structural constraint consistency.

The fourth stage evaluates quantities expected to remain bounded or invariant under the defined mechanical model.

Let

$$h_\ell(X_t)$$

denote an engineering invariant, conservation relationship, structural constraint, or model-consistency quantity. Relative to reference value h_ℓ^* , define

$$e_{\ell,t} = h_\ell(X_t) - h_\ell^*.$$

The admissibility requirement becomes

$$|e_{\ell,t}| \leq \eta_\ell,$$

where η_ℓ is the permitted engineering tolerance.

Therefore,

$$a_4(X_t) = \prod_{\ell=1}^p I(|e_{\ell,t}| \leq \eta_\ell).$$

Depending upon the mechanical benchmark, this layer may contain energy-balance residuals, geometric constraints, stress-strain consistency, kinematic closure conditions, momentum relationships, or other explicitly defined model invariants.

γ_5 : recursive persistence.

A single apparently valid state does not necessarily represent persistent admissibility. The fifth stage therefore evaluates the state across a finite temporal history.

Define the observation window

$$W_t = \{X_{t-L+1}, \dots, X_t\},$$

where L is the number of retained states.

For a monitored residual r_t , define a persistence statistic

$$P_t = \frac{1}{L} \sum_{k=t-L+1}^t I(|r_k| \leq \delta).$$

The state satisfies recursive persistence when

$$P_t \geq P_{\min}.$$

Thus,

$$a_5(X_t) = I(P_t \geq P_{\min}).$$

This permits the architecture to distinguish isolated numerical excursions from persistent mechanical degradation and also allows slowly developing faults to accumulate evidence across successive digital-twin states.

γ_6 : terminal operational admissibility.

The final stage determines whether a state surviving the preceding evaluations may enter the operational state space.

Let

$$\mathcal{S} = \{S_1, S_2, \dots, S_r\}$$

represent the final set of safety-critical operating constraints.

Then


$$a_6(X_t) = \prod_{j=1}^r I(S_j(X_t) = 1).$$

The final decision is

$$A(X_t) = a_1 a_2 a_3 a_4 a_5 a_6.$$

A state satisfying all six stages is admitted to the operational set Ω_A . A state failing any required stage is rejected and its first failure location is retained.

The gamma1-gamma6 recursive admissibility architecture enforces multidimensional constraint safety



Gemini Notebook

2.4 First-Failure Localization

One engineering advantage hypothesized for the ordered architecture is the ability to distinguish classification from localization.

Define the first-failure index as

$$k_{\text{fail}} = \min\{i : a_i(X_t) = 0\}.$$

For an admissible state,

$$k_{\text{fail}} = 0.$$

The diagnostic output is therefore represented by the pair

$$Y_t = (A(X_t), k_{\text{fail}}).$$

The first component states whether the mechanical configuration is admissible. The second identifies the earliest constraint class responsible for rejection.

This creates six mechanically interpretable failure classes while retaining a seventh condition corresponding to complete admissibility.

A sensor corruption should preferentially terminate at γ_1 . An impossible temporal transition should terminate at γ_2 . A coupled-state inconsistency should be detectable at γ_3 . Violation of a defined structural or model invariant should terminate at γ_4 . A persistent temporal anomaly may terminate at γ_5 . A state that survives those conditions while violating a final safety envelope should terminate at γ_6 .

These associations are hypotheses of the engineering implementation, not predetermined results. The benchmark must determine whether the architecture actually produces the expected localization.

2.5 Explicit Kill Condition

The architecture is designed so that an inadmissible result cannot be converted into an admissible result by interpretation after evaluation.

Define the cumulative survival state after stage i as

$$S_i = \prod_{j=1}^i a_j.$$

Because every a_j is binary,

$$S_i = 0 \Rightarrow S_k = 0 \quad \forall k > i.$$

Once a mandatory admissibility condition fails, the candidate state remains rejected for that evaluation cycle.

This produces the engineering kill condition

$$\exists i : a_i(X_t) = 0 \Rightarrow A(X_t) = 0.$$

The state itself is not deleted from the digital twin. It is retained as a rejected state together with its sensor values, constraint residuals, failure stage, timestamp, and ground-truth classification. Retaining rejected states is necessary for subsequent reconstruction analysis, fault diagnosis, and reproducibility.

This treatment is consistent with the source architecture's emphasis on explicit failure pathways rather than unrestricted reinterpretation. The VLSIE formulation states that failure of recursive admissibility, dimensional bookkeeping, computational correspondence, or reduction behavior must remain observable as failure rather than being hidden through symbolic reinterpretation.  12_VLSIE_2026_HARDWARE_VERIFIAB...

3. DIGITAL-TWIN MECHANICAL BENCHMARK

3.1 Benchmark Objective

The preceding equations define an admissibility architecture. They do not establish that the architecture provides an engineering advantage.

That question requires a benchmark.

The purpose of the digital-twin benchmark is to generate a reproducible mechanical environment in which the true operating state is known, controlled disturbances can be introduced deliberately, and the output of the Γ -admissibility architecture can be compared directly against ground truth.

The benchmark therefore evaluates three distinct capabilities:

fault detection,


fault localization,

and state reconstruction.

These three tasks are intentionally separated. A system may successfully recognize that a fault exists while incorrectly identifying its origin. Likewise, a system may correctly detect and localize a disturbance while failing to reconstruct the corresponding nominal digital-twin state with acceptable accuracy.

The MEAP experiment consequently treats each capability as an independently measurable engineering target.

Mapping the physical aerospace envelope to a 10-variable digital twin



The Digital Twin architecture is designed to track nominal trajectories, explicitly corrupted trajectories, and known ground truth for exact quantitative measurement.

Gemini Notebook

3.2 Nominal Mechanical Digital-Twin Model

The benchmark requires a mechanical system whose governing behavior is sufficiently explicit that nominal and faulty states can be distinguished independently of the Γ -admissibility architecture. This prevents the admissibility algorithm from defining its own ground truth.

A coupled rotating mechanical system is adopted as the primary benchmark because it provides simultaneous displacement, velocity, acceleration, rotational, vibration, thermal, stress, strain, actuator, and energetic quantities while permitting controlled introduction of multiple mechanically meaningful fault classes.

The translational structural dynamics are represented initially by the standard second-order equation

$$M\ddot{q}(t) + C\dot{q}(t) + Kq(t) = F(t),$$

where M is the mass matrix, C is the damping matrix, K is the stiffness matrix, $q(t)$ is the generalized displacement vector, and $F(t)$ is the externally applied generalized force vector.

The corresponding velocity and acceleration states are

$$v(t) = \dot{q}(t)$$

and

$$a(t) = \ddot{q}(t).$$

For a rotating component, the angular state is represented by

$$\omega(t) = \dot{\theta}(t),$$

with angular acceleration

$$\alpha(t) = \dot{\omega}(t).$$

The rotational equation of motion is defined as

$$J\dot{\omega}(t) = \tau_m(t) - \tau_L(t) - \tau_f(t),$$

where J represents rotational inertia, τ_m applied motor or propulsion torque, τ_L mechanical load torque, and τ_f dissipative torque associated with friction and other modeled losses.

The structural stress state is related to strain in the initial elastic operating regime through

$$\sigma(t) = E_Y \epsilon(t),$$

where E_Y is Young's modulus.

Mechanical strain may therefore be represented by

$$\epsilon(t) = \frac{\Delta L(t)}{L_0},$$

where L_0 denotes the reference length and $\Delta L(t)$ its deformation.

The benchmark thermal state is represented through an energy-balance equation,

$$C_T \frac{dT}{dt} = Q_{\text{gen}}(t) - Q_{\text{loss}}(t),$$

where C_T is the effective thermal capacity, Q_{gen} represents mechanically or electrically generated heat, and Q_{loss} represents heat transferred from the modeled system.

These relationships generate the nominal digital-twin trajectory independently of the recursive admissibility decision.

The nominal state is denoted

$$X_t^N.$$

A fault operator $\mathcal{F}_k$ is subsequently applied to generate a controlled disturbed state,

$$X_t^{F_k} = \mathcal{F}_k(X_t^N; \lambda_k, t_k),$$

where k identifies the fault class, λ_k specifies fault magnitude, and t_k specifies fault-injection time.

This distinction is fundamental. The Γ architecture does not generate the fault label. The experiment generates the fault independently and subsequently asks whether Γ detects it.

3.3 Ground-Truth State Construction

Every simulated state is assigned a ground-truth label before Γ -admissibility evaluation.

Define

$$y_t = \begin{cases} 1, & X_t \text{ satisfies the benchmark's prescribed mechanical constraints,} \\ 0, & X_t \text{ violates one or more prescribed mechanical constraints.} \end{cases}$$

The admissibility architecture independently produces

$$\hat{y}_t = A(X_t).$$

Correct operation requires

$$\hat{y}_t = y_t.$$

Four experimental outcomes therefore exist.

For $y_t = 1$ and $\hat{y}_t = 1$, the state is a true admission.

For $y_t = 0$ and $\hat{y}_t = 0$, the state is a true rejection.

For $y_t = 0$ and $\hat{y}_t = 1$, the result is a false admission.

For $y_t = 1$ and $\hat{y}_t = 0$, the result is a false rejection.

Of these outcomes, false admission receives particular attention because it represents an unsafe or inconsistent state incorrectly permitted to survive the supervisory architecture.

Define the false-admission rate as

$$FAR = \frac{N_{FA}}{N_{FA} + N_{TR}},$$

where N_{FA} denotes false admissions and N_{TR} true rejections.

Similarly, the false-rejection rate is

$$FRR = \frac{N_{FR}}{N_{FR} + N_{TA}},$$

where N_{FR} denotes false rejections and N_{TA} true admissions.

The overall classification accuracy is

$$ACC = \frac{N_{TA} + N_{TR}}{N_{TA} + N_{TR} + N_{FA} + N_{FR}}.$$

These metrics prevent a visually convincing digital-twin trajectory from being treated as evidence of successful admissibility. Performance must be determined from the complete benchmark population.

3.4 Controlled Fault-Injection Architecture

The benchmark introduces faults systematically rather than adding arbitrary numerical noise.

Each fault is represented by

$$\mathcal{F}_k = (k, \lambda, t_0, \Delta t_f),$$

where k denotes fault type, λ fault severity, t_0 onset time, and Δt_f fault duration.

The first fault class is sensor bias.

For sensor variable x_j ,

$$x_j^F(t) = x_j^N(t) + b_j,$$

where b_j is a controlled bias.

The second class is progressive sensor drift,

$$x_j^F(t) = x_j^N(t) + d_j(t - t_0), \quad t \geq t_0,$$

where d_j is the drift rate.

The third class is vibration amplification,

$$V^F(t) = \lambda_V V^N(t), \quad \lambda_V > 1.$$

The fourth class is thermal excursion,

$$T^F(t) = T^N(t) + \Delta T_f(t).$$

The fifth class is stiffness degradation,

$$K^F = (1 - \delta_K)K,$$

where

$$0 < \delta_K < 1.$$

This fault changes the actual dynamical response rather than merely modifying a sensor value.

The sixth class is damping degradation,

$$C^F = (1 - \delta_C)C.$$

The seventh class is actuator inconsistency. If $q_c(t)$ represents commanded actuator state and $q_a(t)$ actual actuator state, define the tracking residual

$$r_q(t) = q_c(t) - q_a(t).$$

A fault exists when

$$|r_q(t)| > \delta_q$$

for longer than the permitted transient interval.

The eighth class is excessive mechanical stress,

$$\sigma^F(t) > \sigma_{\text{allow}},$$

where σ_{allow} is the predefined allowable stress for the benchmark.

The ninth class is an impossible or deliberately corrupted temporal transition,

$$\left| \frac{x_j(t) - x_j(t-1)}{\Delta t} \right| > r_j^{\text{max}}.$$


The tenth class is a compound fault.

Rather than introducing a single disturbance, two or more fault operators are composed:

$$X_t^{F_{ij}} = (\mathcal{F}_j \circ \mathcal{F}_i)(X_t^N).$$

Compound faults are important because an admissibility architecture intended for autonomous systems should not be evaluated exclusively under isolated one-variable disturbances.

The controlled fault injection environment introduces ten distinct structural and dynamic anomalies



[Visualizing the controlled anomaly pipeline: Mechanical System -> Digital Twin -> Parameter Corruption -> Evaluation.]

Gemini Notebook

3.5 Fault Severity Sweep

Each fault class is evaluated across increasing severity rather than at a single arbitrarily selected magnitude.

Let normalized fault severity be

$$s \in [0, 1].$$

For each fault class k , define

$$\lambda_k(s) = \lambda_{k,\min} + s (\lambda_{k,\max} - \lambda_{k,\min}).$$

This generates a controlled progression from nominal or near-nominal operation through increasingly severe disturbance.

The resulting experiment evaluates not merely whether a fault can eventually be detected, but the severity at which rejection begins.

Define the Γ detection threshold for fault class k as

$$s_k^* = \inf \{s : A(X^{F_k}(s)) = 0\}.$$

This quantity provides a directly measurable characteristic of the admissibility architecture.

If s_k^* lies substantially beyond the benchmark's independently specified unsafe boundary, the architecture permits unsafe states and fails the corresponding constraint-safety test.

If s_k^* occurs substantially before the independently specified unsafe boundary, the architecture may be excessively conservative and generate unnecessary false rejection.

The desirable condition is therefore not simply maximal rejection.

It is agreement between admissibility rejection and independently defined engineering ground truth.

3.6 Fault Localization Accuracy

Detection and localization are evaluated separately.

For each injected fault F_k , let

$$g_k$$

denote the expected first Γ -stage associated with the violated constraint class.

The architecture returns

$$\hat{g}_k = k_{\text{fail}}.$$

Localization accuracy is therefore

$$ACC_{\text{loc}} = \frac{1}{N_F} \sum_{k=1}^{N_F} I(\hat{g}_k = g_k),$$

where N_F is the number of evaluated faulty trials.

A confusion matrix is retained across Γ stages so that systematic mislocalization remains visible even when overall fault detection is high.

This matters because an autonomous maintenance or fault-management architecture must ideally distinguish between a corrupted sensor, an impossible transition, a coupled mechanical inconsistency, a structural-model violation, a persistent degradation condition, and a terminal safety-envelope violation.

3.7 Detection Latency

For a fault introduced at time t_f , let t_d denote the first time at which Γ rejects the corresponding state.

Detection latency is defined as

$$L_d = t_d - t_f.$$

For persistent faults, the architecture is evaluated using the mean latency

$$\bar{L}_d = \frac{1}{N_F} \sum_{k=1}^{N_F} L_{d,k},$$

together with the maximum observed latency.

This quantity becomes particularly important for autonomous aerospace and high-speed mechanical systems, because an architecture that eventually detects a fault may still be unsuitable if the detection delay exceeds the mechanical time available for safe intervention.

The next major stage is **Section 4: Recursive Dimensional Compression and Digital-Twin Reconstruction**. That is where we connect the mechanical benchmark back to the 5D/2880D SEXA architecture and, critically, test reconstruction rather than merely asserting it.




Diagram illustrating the assembly of the sensor module. The motor is labeled "MOTOR AS27 9-10" and the sensor module is labeled "SENSOR M032 6-00". The diagram shows the motor being inserted into the sensor module.



4. RECURSIVE DIMENSIONAL COMPRESSION AND DIGITAL-TWIN RECONSTRUCTION

4.1 Purpose of the Compression Test

The fault-detection experiment established in Section 3 evaluates whether the ordered Γ architecture can determine which mechanical states should survive. A second and more demanding question concerns what happens to the information associated with a surviving state.

The SEXA source architecture distinguishes between a higher-dimensional recursive structural regime and a bounded operational representation. In the full-stack formulation, the 2880-dimensional regime is presented as structural closure, while the five-dimensional regime functions as an operational realization after admissibility enforcement and compression. [05_JAM_9896_FROM_ADMISSIBILITY_...](#) The Γ -SubThreshold Fold manuscript further characterizes the proposed transformation as a six-operator admissibility process into a five-dimensional representation intended to preserve invariants required for reconstruction.

[06_GAMMA_SUBTHRESHOLD_FOLD_LOGI...](#)

For the present mechanical experiment, these propositions are not assumed to guarantee successful reconstruction. They are converted into a measurable compression-reconstruction problem.

Let the complete engineering state be

$$X_t \in \mathbb{R}^n.$$

Define a compression operator

$$C_\Gamma : \mathbb{R}^n \rightarrow \mathbb{R}^5.$$

The compressed operational state is

$$Z_t^{(5)} = C_\Gamma(X_t).$$

A reconstruction operator is then defined independently as

$$\mathcal{R} : \mathbb{R}^5 \rightarrow \mathbb{R}^n,$$

producing

$$\widehat{X}_t = \mathcal{R} \left(Z_t^{(5)} \right).$$

The complete experimental pathway is therefore

$$X_t \xrightarrow{\Gamma} X_t^A \xrightarrow{C_\Gamma} Z_t^{(5)} \xrightarrow{\mathcal{R}} \widehat{X}_t.$$

The reconstruction claim is successful only to the extent that $\widehat{X}_t$ reproduces the independently retained original state X_t .

4.2 No Hidden-State Reconstruction Rule

A critical restriction is imposed on the benchmark.

The reconstruction function may not recover the original state merely by retaining a hidden copy of X_t .

Let

$$\mathcal{I}_R$$

denote all information supplied to the reconstruction algorithm. For a genuine compression-reconstruction test,

$$\mathcal{I}_R = \left\{ Z_t^{(5)}, \Theta_R \right\},$$

where Θ_R contains only reconstruction parameters established by the benchmark protocol.

The following condition is prohibited:

$$X_t \in \mathcal{I}_R.$$

Likewise, an identifier that directly indexes an external table containing the original state is prohibited.

Thus,

$$\boxed{\widehat{X}_t = \mathcal{R}(Z_t^{(5)}, \Theta_R), \quad X_t \notin \mathcal{I}_R}$$

is imposed as the reconstruction admissibility condition.

This distinction is necessary because storing the original configuration and subsequently retrieving it demonstrates storage, not reconstruction.

4.3 Five-Dimensional Operational State

The five-dimensional mechanical representation is constructed from five physically interpretable state families rather than five arbitrarily selected raw sensor channels.

Define

$$Z_t^{(5)} = \begin{bmatrix} z_{D,t} \\ z_{T,t} \\ z_{S,t} \\ z_{C,t} \\ z_{R,t} \end{bmatrix},$$

where z_D represents dynamical state, z_T thermal-fluid state, z_S structural state, z_C control/actuation state, and z_R recursive health or persistence state.

The compression problem therefore becomes

$$\mathcal{C}_T(X_t) = [z_{D,t}, z_{T,t}, z_{S,t}, z_{C,t}, z_{R,t}]^T.$$

Each component must be generated from explicitly disclosed mechanical variables.

For example, the dynamical component may be represented as

$$z_{D,t} = w_v \widetilde{V}_t + w_a \widetilde{a}_t + w_\omega \widetilde{\omega}_t,$$

where the tilde denotes normalized quantities and the coefficients satisfy

$$w_v + w_a + w_\omega = 1.$$

Similarly,

$$z_{T,t} = w_T \widetilde{T}_t + w_P \widetilde{P}_t,$$

and

$$z_{S,t} = w_\sigma \widetilde{\sigma}_t + w_\epsilon \widetilde{\epsilon}_t.$$

The control component is defined from actuator and energetic behavior,

$$z_{C,t} = w_q \tilde{q}_t + w_E \tilde{E}_t,$$

while recursive health incorporates temporal information,

$$z_{R,t} = \frac{1}{L} \sum_{j=0}^{L-1} H_{t-j},$$

where H_t is a normalized health score constructed from the residuals surviving the Γ chain.

The specific coefficients are fixed before the final test set is evaluated. They may not be altered afterward to improve reconstruction results.

4.4 Reconstruction Error

For each reconstructed mechanical variable j , define the normalized reconstruction error

$$e_{j,t} = \frac{\hat{x}_{j,t} - x_{j,t}}{s_j},$$

where s_j is the same normalization scale defined for the benchmark.

The instantaneous normalized reconstruction error is

$$E_R(t) = \sqrt{\frac{1}{n} \sum_{j=1}^n e_{j,t}^2}.$$

Across N evaluated states, normalized root-mean-square reconstruction error is

$$NRMSE_R = \sqrt{\frac{1}{Nn} \sum_{t=1}^N \sum_{j=1}^n \left(\frac{\hat{x}_{j,t} - x_{j,t}}{s_j} \right)^2}.$$

Perfect reconstruction would require

$$NRMSE_R = 0.$$

The benchmark does not assume this result.

Instead, the observed value is reported.

4.5 Maximum Reconstruction Error

An average reconstruction score can conceal severe errors in individual variables. The benchmark therefore also measures

$$E_{\max} = \max_{t,j} \left| \frac{\hat{x}_{j,t} - x_{j,t}}{s_j} \right|.$$

Variable-specific reconstruction errors are retained as

$$RMSE_j = \sqrt{\frac{1}{N} \sum_{t=1}^N (\hat{x}_{j,t} - x_{j,t})^2}.$$

This allows the experiment to determine whether compression preserves some mechanical quantities substantially better than others.

A reconstruction architecture that reproduces temperature and pressure accurately while destroying vibration or stress information cannot be described simply as successful reconstruction without reporting that asymmetry.

4.6 Reconstruction of Nominal and Faulted States

The reconstruction experiment is performed separately on nominal and disturbed trajectories.

Define

$$E_R^N$$

as reconstruction error for nominal states and

$$E_R^{F_k}$$

as reconstruction error under fault class k .

The fault-induced reconstruction penalty is

$$\Delta E_R^{F_k} = E_R^{F_k} - E_R^N.$$

This quantity tests whether the compressed representation becomes unstable precisely when the mechanical system leaves nominal operation.

Such behavior is particularly relevant to digital-twin applications. A representation that reconstructs accurately during nominal operation but deteriorates under mechanical faults may still provide useful health information, but it cannot be described as fault-independent reconstruction.

4.7 Compression Ratio

For the ten-component engineering state defined in this benchmark, direct compression into five operational coordinates gives the state-coordinate ratio

$$CR = \frac{n}{d} = \frac{10}{5} = 2.$$

This should not be confused with the broader SEXA 2880D-to-5D structural claim.

The present experiment directly tests

$$\mathbb{R}^{10} \rightarrow \mathbb{R}^5 \rightarrow \mathbb{R}^{10}.$$

It does **not** simulate 2,880 independent mechanical sensor variables and then describe that experiment as equivalent to the complete 2880-dimensional SEXA manifold.

The relationship to the wider architecture is instead hierarchical: the five-coordinate engineering representation is used as the experimentally tractable operational layer through which the proposed recursive compression principle can be evaluated.

This distinction prevents dimensional terminology from being substituted for measured engineering performance.

4.8 Information-Retention Metric

In addition to reconstruction error, the benchmark measures retained variance.

For original normalized state matrix

$$\mathbf{X} \in \mathbb{R}^{N \times n}$$

and reconstructed matrix

$$\widehat{\mathbf{X}} \in \mathbb{R}^{N \times n},$$

define the retained-information coefficient

$$I_R = 1 - \frac{\|\mathbf{X} - \widehat{\mathbf{X}}\|_F^2}{\|\mathbf{X} - \overline{\mathbf{X}}\|_F^2},$$

where $\|\cdot\|_F$ is the Frobenius norm and $\overline{\mathbf{X}}$ contains the mean state.

Values approaching

$$I_R = 1$$

indicate increasingly accurate reconstruction relative to a mean-state baseline.

Values at or below zero indicate that reconstruction provides no improvement over the corresponding mean predictor under this metric.

4.9 Reconstruction Kill-Switch

The reconstruction claim is rejected for the benchmark if any of the following occurs:

- the reconstruction algorithm directly accesses the original state;
- the five-dimensional representation contains an undeclared lossless encoding of the original vector;
- the reconstruction parameters are fitted using test-set ground truth after evaluation begins;
- the reported reconstruction metric excludes failed variables without prior justification;
- or repeated execution using identical compressed states produces inconsistent reconstructed outputs.

Formally, define the reconstruction-validity indicator

$$K_R = K_{\text{ind}} K_{\text{no-leak}} K_{\text{fixed}} K_{\text{complete}} K_{\text{det}},$$

where each component is binary.

Then

$$K_R = 1$$

is required before quantitative reconstruction accuracy is interpreted.

This produces an important distinction:

Low reconstruction error $\nRightarrow$ valid reconstruction
--

unless the information pathway itself passes the reconstruction audit.

5.1 Experimental Sequence

Every benchmark run follows the same deterministic sequence.

First, a nominal mechanical trajectory X_t^N is generated from the disclosed digital-twin equations.

Second, the trajectory is independently labeled according to the predefined engineering constraints.

Third, when a fault trial is required, the selected fault operator is introduced at a predetermined onset time and severity.

Fourth, the resulting state enters the ordered Γ chain.

Fifth, the system records

$$\mathcal{L}_t = \{X_t, a_1, \dots, a_6, k_{\text{fail}}, A(X_t), Z_t^{(5)}, \widehat{X}_t\}.$$

This object constitutes the computational audit trace for state t .

Sixth, the predicted decision is compared with ground truth.

Seventh, the complete experiment is repeated under identical initial conditions to test deterministic reproducibility.

5.2 Deterministic Repeatability

Let

$$Y_t^{(r)}$$

represent the complete output generated during repetition r .

For R identical executions, deterministic repeatability requires

$$Y_t^{(1)} = Y_t^{(2)} = \dots = Y_t^{(R)}.$$

Define repeatability as

$$D_R = \frac{N_{\text{identical}}}{N_{\text{repeated}}}.$$

The deterministic target is

$$D_R = 1.$$

Any value below unity requires investigation because identical inputs and parameters should not generate different admissibility decisions in the deterministic implementation proposed here.

5.3 Ablation Test

A central benchmark requirement is determining whether all six Γ stages contribute independently to system behavior.

For stage i , define the ablated architecture

$$\Gamma_{-i} = \gamma_6 \circ \cdots \circ \gamma_{i+1} \circ \gamma_{i-1} \circ \cdots \circ \gamma_1.$$

The complete experiment is rerun after removing each stage individually.

If $M(\Gamma)$ denotes a selected performance metric, define the ablation effect as

$$\Delta M_i = M(\Gamma) - M(\Gamma_{-i}).$$

For fault-detection performance,

$$\Delta ACC_i = ACC(\Gamma) - ACC(\Gamma_{-i}).$$

For unsafe-state admission,

$$\Delta FAR_i = FAR(\Gamma_{-i}) - FAR(\Gamma).$$

A Γ stage that produces no measurable change across any appropriate benchmark condition cannot be claimed to provide demonstrated engineering value from this experiment alone.

5.4 Baseline Comparison

The complete Γ architecture is compared with a conventional independent-threshold baseline.

Define

$$B(X_t) = \prod_{j=1}^n I(x_{j,\min} \leq x_{j,t} \leq x_{j,\max}).$$

The baseline therefore asks whether every individual sensor remains inside its permitted scalar interval.

Γ asks the stronger question of whether the state survives acquisition, transition, coupled-state, invariant, recursive-persistence, and terminal operational constraints.

The comparison is consequently

$$\Gamma \quad \text{versus} \quad B.$$

The principal test is not whether Γ can detect extreme faults that are already obvious from individual sensor thresholds. The more discriminating test is whether Γ detects mechanically inconsistent states for which

$$B(X_t) = 1$$

but

$$A(X_t) = 0.$$

These are precisely the conditions under which ordered coupled constraints could provide information beyond independent threshold monitoring.

5.5 Primary Falsification Conditions

The engineering hypothesis is considered unsupported under the tested benchmark if the architecture cannot reproducibly distinguish nominal and faulty states, if false admission remains unacceptably high relative to the specified operating constraints, if Γ provides no measurable benefit over the baseline for coupled or temporal faults, if first-failure localization does not correspond to the injected constraint class, or if reconstruction performance cannot be reproduced without hidden access to the original state.

These are computational failure conditions, not rhetorical qualifications.

The source SEXA survivability audit itself distinguishes formal consistency from empirical sufficiency and states that experimental confirmation and independent reproduction remain separate requirements.

 07_SEXA_UFT_STRUCTURED_SURVIVAB... The MEAP benchmark operationalizes that distinction by assigning observable failure criteria before interpreting the results.

The conference currently explicitly lists autonomous systems, smart structural-health monitoring, predictive maintenance, rotating-machinery fault detection, digital maintenance twins, and reliability engineering within its scope, making this benchmark directly aligned with the stated MEAP 2026 engineering program.

 MEAP 2026

6. COMPUTATIONAL IMPLEMENTATION AND BENCHMARK EXECUTION

6.1 Executable Architecture

The mathematical architecture developed in Sections 2–5 is implemented as a deterministic computational pipeline rather than evaluated through qualitative inspection of plotted trajectories.

For each discrete mechanical state X_t , the implementation executes

$$X_t \rightarrow \gamma_1 \rightarrow \gamma_2 \rightarrow \gamma_3 \rightarrow \gamma_4 \rightarrow \gamma_5 \rightarrow \gamma_6 \rightarrow A(X_t).$$

The implementation retains the complete evaluation record rather than only the final binary decision.

For state t , the audit record is

$$\mathcal{L}_t = [t, X_t, a_{1,t}, a_{2,t}, a_{3,t}, a_{4,t}, a_{5,t}, a_{6,t}, k_{\text{fail},t}, A_t, y_t].$$

The resulting log makes every final classification traceable to the exact mechanical state and exact Γ stage responsible for the decision.

This is particularly important for hardware-verification applications. A final classifier that reports only “safe” or “unsafe” may provide useful operational information, but it does not by itself expose the internal computational pathway required to reproduce the decision.

6.2 Deterministic State Generator

The benchmark begins from a reproducible nominal trajectory.

For discrete time index k ,

$$t_k = k\Delta t, \quad k = 0, 1, \dots, N - 1.$$

The mechanical state is propagated through the benchmark dynamics according to

$$X_{k+1} = \Phi(X_k, u_k, \Theta, \Delta t),$$

where u_k represents the prescribed control input, Θ contains fixed model parameters, and Φ is the numerical state-transition operator.

The nominal simulation is therefore completely specified by

$$\mathcal{B} = \{X_0, \Theta, u_{0:N-1}, \Delta t, N\}.$$

Two executions supplied with identical $\mathcal{B}$ must generate identical nominal trajectories within the numerical precision of the implementation.

Where stochastic measurement noise is included, the pseudorandom seed becomes an explicit component of $\mathcal{B}$. It is recorded with the benchmark so that the same noisy trajectory can be regenerated exactly.

6.3 Independent Fault Generator

Fault generation is implemented separately from Γ evaluation.

This separation prevents the admissibility algorithm from determining the conditions against which it is subsequently judged.

For fault class j ,

$$X_k^{F_j} = \mathcal{F}_j(X_k^N, \lambda_j, k_{0,j}, D_j),$$

where λ_j is severity, $k_{0,j}$ is the injection index, and D_j is fault duration.

The corresponding ground-truth label is assigned before Γ evaluation:

$$y_k = \begin{cases} 1, & \text{nominal/admissible,} \\ 0, & \text{faulted/inadmissible.} \end{cases}$$

Consequently,

$$\mathcal{F} \perp \Gamma$$

at the classification stage: the mechanism generating the experimental fault is computationally separated from the mechanism attempting to detect it.

6.4 Benchmark Population

The experimental population is divided into nominal, single-fault, severity-sweep, compound-fault, and repeated-execution trials.

Let

$$\mathcal{D} = \mathcal{D}_N \cup \mathcal{D}_S \cup \mathcal{D}_V \cup \mathcal{D}_C \cup \mathcal{D}_R,$$

where $\mathcal{D}_N$ contains nominal trajectories, $\mathcal{D}_S$ single-fault trajectories, $\mathcal{D}_V$ severity sweeps, $\mathcal{D}_C$ compound-fault trajectories, and $\mathcal{D}_R$ deterministic repetitions.

The complete evaluation is performed over

$$N_{\text{total}} = N_N + N_S + N_V + N_C + N_R$$

evaluated states.

No benchmark state is removed because it produces an unfavorable classification result.

6.5 Confusion Matrix

The primary classification results are accumulated in the confusion matrix

$$\mathbf{C} = \begin{bmatrix} N_{TA} & N_{FR} \\ N_{FA} & N_{TR} \end{bmatrix}.$$

From this matrix, the benchmark computes the previously defined accuracy, false-admission rate, and false-rejection rate.

Fault-detection sensitivity is

$$S_{\text{fault}} = \frac{N_{TR}}{N_{TR} + N_{FA}}.$$

Nominal-state acceptance specificity is

$$S_{\text{nominal}} = \frac{N_{TA}}{N_{TA} + N_{FR}}.$$

Balanced accuracy is therefore

$$ACC_{\text{bal}} = \frac{S_{\text{fault}} + S_{\text{nominal}}}{2}.$$

Balanced accuracy is retained because a benchmark containing substantially more nominal than faulty states, or vice versa, can produce a misleadingly high raw accuracy.

6.6 Γ -Stage Failure Distribution

For rejected states, the implementation records the first failure stage.

Define

$$N_i = \sum_{t=1}^N I(k_{\text{fail},t} = i), \quad i = 1, \dots, 6.$$

The normalized failure distribution is

$$p_i = \frac{N_i}{\sum_{j=1}^6 N_j}.$$

Thus,

$$\sum_{i=1}^6 p_i = 1.$$

The vector

$$\mathbf{p}_\Gamma = [p_1, p_2, p_3, p_4, p_5, p_6]$$

provides a compact description of where the architecture rejects the benchmark population.

This information is evaluated separately for every injected fault class.

6.7 Fault-to- Γ Localization Matrix

A more stringent localization test constructs the matrix

$$\mathbf{L}_{ij} = N(F_i \rightarrow \gamma_j),$$

where row i represents injected fault class F_i and column j represents the observed first-failure stage γ_j .

An idealized fault specifically designed to violate stage j would produce concentration near the corresponding diagonal element.

However, diagonal localization is not imposed artificially.

If a fault physically creates an earlier constraint violation, the architecture is expected to reject the state at that earlier stage. The audit record therefore retains both the intended injected fault and the actual first violated constraint.

This distinction prevents the desired localization label from overriding the mechanical behavior generated by the simulation.

6.8 Computational Latency

For each state, let

$$t_{\text{in}}$$

denote the time immediately before Γ evaluation and

$$t_{\text{out}}$$

the time immediately after the complete decision has been returned.

Per-state computational latency is

$$\tau_{\Gamma,t} = t_{\text{out}} - t_{\text{in}}.$$

Mean latency is

$$\bar{\tau}_{\Gamma} = \frac{1}{N} \sum_{t=1}^N \tau_{\Gamma,t},$$

and worst-case measured latency is

$$\tau_{\Gamma}^{\max} = \max_t \tau_{\Gamma,t}.$$

Throughput is reported as

$$R_{\Gamma} = \frac{N}{T_{\text{execution}}}$$

in evaluated states per second.

These quantities are software-execution measurements and are not presented as FPGA, ASIC, flight-computer, or embedded-hardware timing measurements.

That distinction is necessary because the predecessor SEXA hardware-oriented paper proposed eventual hardware-assisted implementation while explicitly distinguishing architectural suitability from a completed hardware realization. [12_VLSIE_2026_HARDWARE_VERIFIAB...](#)

6.9 Deterministic Replay Test

The entire benchmark is executed repeatedly without changing its initial state, model parameters, fault schedule, thresholds, or pseudorandom seed.

Let the complete output of execution r be

$$\mathcal{O}^{(r)}.$$

Exact classification reproducibility requires

$$\mathcal{O}_A^{(1)} = \mathcal{O}_A^{(2)} = \dots = \mathcal{O}_A^{(R)}.$$

For floating-point residuals, numerical agreement is additionally evaluated under tolerance ε_{num} :

$$\left| r_i^{(p)} - r_i^{(q)} \right| \leq \epsilon_{\text{num}}.$$

A cryptographic hash of the serialized decision log may additionally be retained for each deterministic execution.

Identical logs should produce

$$H(\mathcal{L}^{(1)}) = H(\mathcal{L}^{(2)}) = \dots = H(\mathcal{L}^{(R)}).$$

This converts reproducibility from visual similarity into a machine-checkable condition.

7. RESULTS REPORTING ARCHITECTURE

7.1 Results Are Execution-Dependent

The numerical results of this paper are constrained by a strict reporting rule:

No numerical benchmark result is reported unless produced by execution.

Values for accuracy, false admission, reconstruction error, detection latency, throughput, localization accuracy, or ablation performance are therefore not selected in advance to support the hypothesis.

This restriction is particularly important because the purpose of the present paper is to move from architectural description toward engineering-grade computational validation.

7.2 Primary Performance Table

The executed benchmark populates the following performance structure:

Metric	Full Γ Architecture	Independent-Threshold Baseline
Overall accuracy	Executed value	Executed value
Balanced accuracy	Executed value	Executed value
Fault sensitivity	Executed value	Executed value
False-admission rate	Executed value	Executed value
False-rejection rate	Executed value	Executed value
Mean detection latency	Executed value	Executed value
Worst detection latency	Executed value	Executed value
Localization accuracy	Executed value	N/A
Mean computation latency	Executed value	Executed value
States evaluated per second	Executed value	Executed value
Deterministic replay	Executed value	Executed value

The entries remain execution-dependent until the complete benchmark code is run.

7.3 Reconstruction Results Table

The reconstruction experiment is reported separately:

Reconstruction Quantity	Nominal	Single Fault	Compound Fault
NRMSE	Executed value	Executed value	Executed value
Maximum normalized error	Executed value	Executed value	Executed value
Information-retention coefficient	Executed value	Executed value	Executed value
Reconstruction-validity audit	Pass/Fail	Pass/Fail	Pass/Fail

Separating reconstruction from classification prevents high fault-detection accuracy from being used as indirect evidence that dimensional reconstruction also succeeded.

7.4 Ablation Results

The complete six-stage architecture is evaluated against six ablated architectures:

$$\Gamma_{-1}, \Gamma_{-2}, \Gamma_{-3}, \Gamma_{-4}, \Gamma_{-5}, \Gamma_{-6}.$$

For each ablation, the benchmark records

$$\Delta ACC_i, \quad \Delta FAR_i, \quad \Delta FRR_i, \quad \Delta ACC_{loc,i}, \quad \Delta \tau_i.$$

This produces a direct test of whether the ordered six-stage architecture contains computationally consequential structure.

If removal of an operator produces no meaningful performance change under any benchmark condition designed to exercise that operator, the paper must report that result.

Conversely, if removal significantly increases unsafe-state admission or destroys localization for the corresponding fault class, the experiment supplies quantitative evidence that the stage contributes to the tested engineering task.

7.5 Order-Permutation Test

Ablation alone does not determine whether Γ 's **ordering** matters.

The stronger claim of ordered admissibility therefore requires an additional experiment.

Let

$$\pi(\Gamma)$$

represent a permutation of the six Γ stages.

For permutation π_j , define

$$\Gamma_{\pi_j} = \gamma_{\pi_j(6)} \circ \gamma_{\pi_j(5)} \circ \cdots \circ \gamma_{\pi_j(1)}.$$

The canonical architecture is

$$\Gamma_{\text{canonical}} = \gamma_6 \circ \gamma_5 \circ \gamma_4 \circ \gamma_3 \circ \gamma_2 \circ \gamma_1.$$

Its performance is compared against permuted chains using

$$\Delta M_{\pi_j} = M(\Gamma_{\text{canonical}}) - M(\Gamma_{\pi_j}).$$

This is a particularly important falsification test.

If permutation of the Γ stages leaves all relevant outputs unchanged, then the benchmark does not demonstrate that operator ordering contributes to the measured engineering performance.

If ordering changes classification, localization, latency, or reconstruction in systematic and reproducible ways, the effect becomes quantitatively measurable.

7.6 Strongest Benchmark Condition

The most informative benchmark cases are not extreme faults.

A severe temperature excursion that exceeds an obvious shutdown threshold can be detected by almost any threshold architecture.

The stronger test consists of states satisfying

$$B(X_t) = 1$$

while simultaneously violating a coupled, temporal, recursive, or invariant constraint such that

$$A(X_t) = 0.$$

Define this subset as

$$\mathcal{D}_{\text{hidden}} = \{X_t : B(X_t) = 1 \wedge A_{\text{truth}}(X_t) = 0\}.$$

Performance on this subset is measured independently:

$$ACC_{\text{hidden}} = \frac{N_{\text{correct}}(\mathcal{D}_{\text{hidden}})}{|\mathcal{D}_{\text{hidden}}|}.$$

This becomes one of the central engineering tests of the paper because it asks whether recursive admissibility detects inconsistent states that conventional independent range checking permits.

8. INTERPRETATION BOUNDARY

A successful execution of the benchmark would establish a limited result.

It would demonstrate that the implemented ordered Γ architecture can—or cannot—provide measurable engineering performance for the digital-twin problem defined in this paper.

It would not independently establish the physical existence of a 2880-dimensional manifold, validate a unified field theory, or transform computational correspondence into experimental confirmation.

This boundary follows the SEXA survivability-audit document itself, which distinguishes internal structural survivability from experimental confirmation and identifies independent reproduction and further formalization as outstanding requirements.  07_SEXA_UFT_STRUCTURED_SURVIVAB...

The engineering proposition tested here is narrower and directly falsifiable:

Does ordered recursive admissibility improve measurable mechanical state discrimination, localization, and reconstruction?

The answer must come from the benchmark.

9. BENCHMARK PARAMETERIZATION AND EXECUTION CONDITIONS

9.1 Reproducible Mechanical Test Configuration

The digital-twin benchmark is parameterized before the final evaluation is executed. This prevents engineering limits, Γ thresholds, fault magnitudes, or reconstruction tolerances from being altered after observing the resulting classification performance.

The benchmark is implemented as a discrete-time simulation with sampling interval

$$\Delta t$$

and total duration

$$T_{\text{sim}} = N\Delta t.$$

The state propagated at every discrete instant is

$$X_k = [T_k, P_k, V_k, a_k, \omega_k, \sigma_k, \epsilon_k, q_k, E_k, F_k]^T.$$

All physical variables remain available in their dimensional engineering representation. Normalized variables are used only where dimensionless residual comparison is required.

For every variable x_j , four classes of bounds are distinguished:

$$D_j^{\text{sensor}}, \quad D_j^{\text{nominal}}, \quad D_j^{\text{warning}}, \quad D_j^{\text{terminal}}.$$

The sensor domain describes values that can be validly represented by the acquisition system.

The nominal domain describes ordinary benchmark operation.

The warning domain describes degraded but potentially recoverable operation.

The terminal domain specifies the boundary beyond which the final operational-admissibility stage rejects the state.

These domains are deliberately separated because

$$\text{valid measurement} \neq \text{nominal measurement} \neq \text{safe operational state}.$$

9.2 Reference Operating State

Let the reference operating condition be

$$X_{\text{ref}} = [T_0, P_0, V_0, a_0, \omega_0, \sigma_0, \epsilon_0, q_0, E_0, 0]^T.$$

Nominal operation is generated as a bounded perturbation around this state:

$$X_k^N = X_{\text{ref}} + \delta X_k.$$

The nominal perturbation is constrained such that

$$A_{\text{truth}}(X_k^N) = 1$$

throughout the fault-free trajectory.

The exact nominal state is retained independently from the Γ architecture and therefore remains available for determining whether a rejection during nominal operation constitutes a false rejection.

9.3 Benchmark Constraint Table

The computational implementation associates each Γ stage with an explicitly testable engineering condition.

Γ Stage	Engineering Function	Principal Failure Type
γ_1	Acquisition/state validity	Missing, nonfinite, corrupted, out-of-domain sensor state
γ_2	Transition consistency	Impossible state-rate or abrupt temporal discontinuity
γ_3	Coupled mechanical consistency	Individually valid but mutually inconsistent variables
γ_4	Structural/model invariance	Violation of constitutive, energetic, or modeled invariant
γ_5	Recursive persistence	Temporally accumulated degradation or persistent residual
γ_6	Terminal operational envelope	Final safety or mission-operating limit

The table defines the engineering interpretation used only for this benchmark. It does not assert that the original abstract Γ operators have already been experimentally identified with these six mechanical mechanisms.

9.4 Mechanical Residual Vector

Rather than relying exclusively upon binary thresholds, the implementation retains the continuous residual produced by each constraint.

Define

$$\mathbf{R}_k = [r_{1,k}, r_{2,k}, r_{3,k}, r_{4,k}, r_{5,k}, r_{6,k}]^T.$$

Each residual is normalized against its corresponding admissibility tolerance:

$$\rho_{i,k} = \frac{|r_{i,k}|}{\tau_i}.$$

Consequently,

$$\rho_{i,k} \leq 1$$

represents satisfaction of stage i , while

$$\rho_{i,k} > 1$$

represents violation.

The binary stage decision becomes

$$a_{i,k} = I(\rho_{i,k} \leq 1).$$

The complete Γ decision is therefore

$$A_k = \prod_{i=1}^6 I(\rho_{i,k} \leq 1).$$

Retaining $\rho_{i,k}$ provides substantially more diagnostic information than retaining the binary output alone.

A state barely crossing a constraint boundary and a state exceeding it by an order of magnitude both produce

$$a_i = 0,$$

but their residual magnitudes remain distinguishable.

9.5 Γ -Admissibility Margin

The minimum normalized distance from violation is represented by the admissibility margin

$$M_{\Gamma,k} = 1 - \max_i \rho_{i,k}.$$

Therefore,

$$M_{\Gamma,k} > 0$$

indicates that all constraints are satisfied,

$$M_{\Gamma,k} = 0$$

indicates that at least one constraint lies exactly on its boundary, and

$$M_{\Gamma,k} < 0$$

indicates rejection.

This provides a continuous supervisory quantity while preserving the binary kill condition.

The magnitude

$$|M_{\Gamma,k}|$$

also gives a compact indication of how strongly the state lies inside or outside the admissibility envelope.

9.6 Fault-Severity Schedule

Each fault class is evaluated across a common normalized severity grid

$$\mathcal{S} = \{s_1, s_2, \dots, s_m\}, \quad 0 \leq s_i \leq 1.$$

For fault F_j , severity s modifies the corresponding physical parameter through

$$\Theta_j^F(s) = \Theta_j^N + s\Delta\Theta_j.$$

For multiplicative disturbances,

$$\Theta_j^F(s) = \Theta_j^N(1 + s\lambda_j).$$

The fault schedule therefore evaluates progressive departure from nominal behavior rather than testing only catastrophic failure.

For every severity level, the experiment records

$$[s, A_{\text{truth}}, A_{\Gamma}, k_{\text{fail}}, M_{\Gamma}, L_d].$$

The resulting detection curve is

$$D_j(s) = A_{\Gamma} (X^{F_j}(s)).$$

9.7 Boundary Agreement Error

Because the true unsafe boundary is specified independently, the experiment can compare the fault severity at which the benchmark becomes mechanically inadmissible with the severity at which Γ first rejects it.

Let

$$s_{j,\text{truth}}^*$$

denote the independently defined engineering boundary and

$$s_{j,\Gamma}^*$$

the Γ rejection boundary.

Define boundary error as

$$E_{B,j} = s_{j,\Gamma}^* - s_{j,\text{truth}}^*.$$

Interpretation is direct.

If

$$E_{B,j} > 0,$$

Γ rejects too late and admits some independently unsafe states.

If

$$E_{B,j} < 0,$$

Γ rejects too early and is excessively conservative.

The ideal boundary agreement is

$$E_{B,j} = 0.$$

~

Across all fault classes,

$$MAE_B = \frac{1}{N_F} \sum_{j=1}^{N_F} |E_{B,j}|.$$

This quantity gives the benchmark a stricter target than simple eventual fault detection.

10. STATISTICAL AND ERROR ANALYSIS

10.1 Confidence Intervals for Classification Performance

A single reported accuracy value does not characterize uncertainty.

For a binary performance proportion

$$\hat{p} = \frac{k}{N},$$

the benchmark therefore reports a confidence interval in addition to the point estimate.

This applies to fault sensitivity, nominal-state acceptance, localization accuracy, and other binary performance quantities.

Accordingly, classification performance is reported in the form

$$\hat{p} \quad [\text{lower CI}, \text{upper CI}].$$

The confidence procedure is fixed before final reporting and applied consistently across methods.

10.2 Error by Fault Severity

Aggregate accuracy may conceal a critical failure near the operational boundary.

Performance is therefore conditioned upon severity:

$$ACC(s) = \frac{N_{\text{correct}}(s)}{N_{\text{total}}(s)}.$$

Similarly,

$$FAR(s) = \frac{N_{FA}(s)}{N_{\text{fault}}(s)}.$$

The resulting curves reveal whether the architecture performs differently for weak, moderate, and severe disturbances.

A system that detects only catastrophic faults may exhibit high overall accuracy when the benchmark is dominated by easily identifiable conditions while still performing poorly near the decision boundary of engineering interest.

10.3 Performance by Fault Class

Metrics are also retained separately for each fault class.

For fault F_j ,

$$S_j = \frac{N_{TR,j}}{N_{TR,j} + N_{FA,j}}.$$

This produces the class-specific sensitivity vector

$$\mathbf{S} = [S_1, S_2, \dots, S_{N_F}].$$

The minimum class performance is

$$S_{\min} = \min_j S_j.$$

Reporting $S_{\min}$ prevents strong performance on easily detected faults from masking a complete failure on one mechanically important fault class.

10.4 Baseline Improvement

For metric M , the absolute Γ improvement over the independent-threshold baseline is

$$\Delta M = M_{\Gamma} - M_B.$$

For error quantities where lower values are preferable,

$$\Delta E = E_B - E_{\Gamma}.$$

The comparison is especially important on the hidden-inconsistency subset previously defined by

$$B(X) = 1$$

while

$$A_{\text{truth}}(X) = 0.$$

For these states, a threshold detector is structurally unable to reject the state solely from individual scalar limits.

Any demonstrated Γ advantage must therefore appear through transition, coupling, invariant, persistence, or other explicitly modeled relationships rather than through a repackaged threshold check.

10.5 Computational Complexity

For a state containing n measured variables and m_i constraints at Γ stage i , the cost of one complete evaluation can be represented as

$$C_{\Gamma} = \sum_{i=1}^6 C(\gamma_i).$$

When each stage is implemented through a bounded collection of algebraic residual evaluations,

$$C_{\Gamma} = O \left(n + \sum_{i=1}^6 m_i \right).$$


The persistence stage additionally depends upon temporal-window length L . Without recursive update,

$$C(\gamma_5) = O(L).$$

With a maintained rolling accumulator, the persistence statistic can be updated without re-evaluating the complete history, reducing its incremental update cost.

The benchmark reports empirical timing because asymptotic complexity alone does not establish suitability for real-time mechanical control.

Safe individual scalars frequently mask catastrophic coupled state failures



Central Engineering Principle: Values can each remain securely inside isolated scalar limits while their combined multidimensional state violates dynamical, structural, energetic, temporal, or operational relationships.

Gemini Notebook

11. ENGINEERING INTERPRETATION

11.1 Why Ordered Admissibility Could Matter

The central engineering distinction introduced by the architecture is between independently legal measurements and a jointly admissible mechanical state.

Consider a simplified state satisfying

$$T_{\min} < T < T_{\max},$$

$$\omega_{\min} < \omega < \omega_{\max},$$

and

$$V < V_{\max}.$$

A conventional range checker may accept the state.

However, suppose that nominal operation requires vibration to remain inside a rotational-speed-dependent envelope,

$$V \leq V_{\text{allow}}(\omega).$$

A state may satisfy

$$V < V_{\max}$$

while violating

$$V > V_{\text{allow}}(\omega).$$

The scalar values remain independently permissible, yet their joint mechanical relationship is abnormal.

This is precisely the class of condition that γ_3 is designed to expose.

Similarly, a temperature value may remain individually acceptable while its time derivative satisfies

$$\left| \frac{dT}{dt} \right| > r_T^{\max},$$

causing rejection by γ_2 before the absolute thermal limit is crossed.

The engineering value of recursive admissibility, if demonstrated, therefore lies not in replacing existing physical laws but in organizing heterogeneous constraints into an ordered, traceable supervisory computation.

11.2 Relationship to Structural-Health Monitoring

Structural-health-monitoring systems often combine multiple measurements whose meaning depends upon both current magnitude and temporal evolution.

The present architecture can represent these relationships through

$$\gamma_3$$

for coupled mechanical consistency,

$$\gamma_4$$

for structural or constitutive relationships, and

$$\gamma_5$$

for temporal persistence.

For example, a progressive reduction in stiffness can alter the system's dynamic response before an individual displacement or acceleration variable crosses its absolute shutdown threshold.

A recursive architecture can therefore retain the evolving residual history as part of the health-state decision.

The proposed framework should not be interpreted as replacing established structural-dynamics or prognostic models. Those models provide the physical relationships evaluated by the admissibility stages.

Γ supplies the supervisory ordering and rejection logic.

11.3 Relationship to Predictive Maintenance

Predictive maintenance requires more than identification of catastrophic failure.

A useful architecture should distinguish

$$\text{nominal} \rightarrow \text{degrading} \rightarrow \text{inadmissible}.$$

The persistence operator provides a mechanism for representing that progression.

Let

$$H_k = 1 - \max_i \rho_{i,k}.$$

A sequence satisfying

$$H_{k-L} < H_{k-L+1} < \dots < H_k$$

in the negative direction toward zero may indicate progressive erosion of the admissibility margin even before final rejection occurs.

The rate

$$\dot{H}_k \approx \frac{H_k - H_{k-1}}{\Delta t}$$

can therefore be retained as a maintenance-oriented diagnostic variable.

This paper does not convert H_k directly into remaining useful life. Such a prediction would require an independently validated degradation model.

Instead, H_k is treated as an observable constraint-margin trajectory.

11.4 Relationship to Autonomous Aerospace Systems

Autonomous aerospace systems operate under tightly coupled constraints and often require decisions under limited computational latency.

The proposed Γ implementation is relevant only if three conditions can be demonstrated:

correctness, traceability, timeliness.

High classification accuracy without bounded execution time would not establish real-time suitability.

Low latency without safe-state discrimination would likewise be insufficient.

And a highly accurate classifier whose decisions cannot be traced through explicit constraint stages would not satisfy the hardware-verification objective developed in the predecessor architecture.

The engineering requirement is therefore their conjunction:

$$Q_{\text{eng}} = Q_{\text{correct}} \wedge Q_{\text{trace}} \wedge Q_{\text{time}}.$$

12. LIMITATIONS AND EXPLICIT FALSIFIABILITY

12.1 Digital Twin Rather Than Physical Flight Hardware

The present benchmark is computational.

It evaluates a mechanical digital twin and controlled fault injection rather than an aircraft, spacecraft, turbine, or flight-qualified propulsion system.

Therefore,

successful simulation $\nrightarrow$ flight qualification.

Experimental hardware testing remains a separate engineering requirement.

12.2 Model Dependence

Γ cannot identify a mechanical inconsistency that has not been represented by a constraint available to its operators.

If the digital twin omits an important failure mechanism,

$$F_{\text{unknown}} \notin \mathcal{C},$$

then successful evaluation against the modeled fault set does not establish detection of F_{unknown} .

Thus the architecture's performance is bounded by the quality and completeness of the engineering model supplied to it.

12.3 Threshold and Tolerance Dependence

The admissibility decisions depend upon tolerances

$$\tau_1, \dots, \tau_6.$$

Improperly chosen tolerances can increase either false admission or false rejection.

The benchmark therefore does not treat the numerical thresholds as universal SEXA constants. They are engineering parameters associated with the modeled system and must be established from system requirements, calibration, empirical data, certified envelopes, or other defensible engineering sources.

12.4 Reconstruction Limitation

Reduction to five operational coordinates necessarily creates a strong reconstruction challenge.

For a general unconstrained mapping

$$\mathbb{R}^{10} \rightarrow \mathbb{R}^5,$$

unique inversion is not guaranteed.

Therefore, successful reconstruction can occur only to the extent that the admissible mechanical states occupy a sufficiently constrained lower-dimensional structure or that the reconstruction model captures relationships permitting recovery.

If distinct admissible states satisfy

$$X_a \neq X_b$$

but

$$\mathcal{C}_\Gamma(X_a) = \mathcal{C}_\Gamma(X_b),$$

then unique deterministic reconstruction requires additional information and cannot be inferred from the five-dimensional vector alone.

This is an explicit mathematical failure condition.

12.5 Broader SEXA Interpretation

The source corpus presents the 5D operational/2880D recursive relationship as part of the wider SEXA theoretical framework.  05_JAM_9896_FROM_ADMISSIBILITY_...

The current benchmark tests an engineering analogue of admissibility-controlled compression and reconstruction.

It does not independently establish that the full 2880-dimensional theoretical structure corresponds to experimentally observed spacetime or fundamental physics.

The wider SEXA survivability audit itself distinguishes internal consistency testing from completed experimental verification.  07_SEXA_UFT_STRUCTURED_SURVIVAB...

That distinction is retained here.

12.6 Direct Rejection Criteria

The principal hypothesis of this paper is rejected under the benchmark if any of the following occurs:

$$FAR_{\Gamma} \geq FAR_B$$

on the mechanically inconsistent states for which the ordered architecture is intended to provide an advantage;

the canonical six-stage ordering performs no differently from relevant permutations;

individual Γ stages produce no measurable effect when exercised by their corresponding fault classes;

first-failure localization is not reproducible;

deterministic replay fails under identical inputs;

software latency exceeds the benchmark's prescribed real-time requirement;

or compression-reconstruction requires access to hidden original-state information.

These conditions ensure that an unfavorable outcome remains an admissible scientific result rather than being interpreted away.

13. CONCLUSION

This paper develops a mechanical-engineering implementation of recursive admissibility derived from the SEXA computational framework and places it into a directly falsifiable digital-twin benchmark.

The architecture represents a candidate mechanical state as

$$X_t = [T_t, P_t, V_t, a_t, \omega_t, \sigma_t, \epsilon_t, q_t, E_t, F_t]^T$$

and subjects that state to the ordered operator chain

$$\Gamma = \gamma_6 \circ \gamma_5 \circ \gamma_4 \circ \gamma_3 \circ \gamma_2 \circ \gamma_1.$$

Engineering admissibility requires

$$A(X_t) = \prod_{i=1}^6 a_i(X_t) = 1.$$

The six stages are operationalized as acquisition validity, transition consistency, coupled mechanical consistency, structural or invariant consistency, recursive persistence, and terminal operational admissibility.

The architecture additionally records the first failing stage

$$k_{\text{fail}},$$

thereby converting a binary state rejection into an explicitly traceable diagnostic event.

A controlled digital-twin benchmark is then defined around nominal operation, sensor faults, thermal excursions, vibration anomalies, stiffness and damping degradation, actuator inconsistency, structural-limit violations, impossible transitions, and compound disturbances.

The evaluation is designed around measurable quantities including classification accuracy, false-admission rate, false-rejection rate, detection latency, fault-stage localization, deterministic replay, computational execution time, ablation sensitivity, order sensitivity, and reconstruction error.

14. EXECUTED HELD-OUT DIGITAL-TWIN BENCHMARK RESULTS

14.1 Frozen Calibration and Independent Evaluation

The final benchmark was executed under a calibration/evaluation separation designed to prevent post hoc adjustment of the Γ -admissibility thresholds. This implements the manuscript requirement that benchmark parameters be established before final evaluation and that unfavorable classifications remain part of the reported result. The existing methodology also requires numerical results to be execution-dependent rather than selected in advance. 9th International Conference on...

Twenty calibration seeds were used exclusively to establish nominal residual and transition tolerances. These calibration trajectories were not included in the final evaluation population. After calibration, all Γ thresholds were frozen.

The final held-out evaluation then used ten different deterministic seeds that had not participated in calibration.

The evaluation consisted of 10 nominal trajectories plus 400 fault trajectories generated from ten fault classes, four severity levels, and ten held-out seeds:

$$N_{\text{runs}} = 10 + (10)(4)(10) = 410.$$


Each trajectory contained 3,000 discrete states, giving

$$N_{\text{total}} = 1,230,000$$

held-out state evaluations.

No final-evaluation classification was used to modify the frozen Γ tolerances.

A frozen parameter protocol guarantees strict separation of calibration and evaluation data



14.2 Primary Classification Result

The Γ architecture produced

$$N_{TA} = 587,597,$$

$$N_{TR} = 612,492,$$

$$N_{FA} = 27,508,$$

and

$$N_{FR} = 2,403.$$

Consequently,

$$ACC_{\Gamma} = \frac{587597 + 612492}{1,230,000}$$

yielded

$$ACC_{\Gamma} = 97.57\%.$$

Fault sensitivity was

$$S_{\text{fault},\Gamma} = \frac{612492}{612492 + 27508}$$

or

$$S_{\text{fault},\Gamma} = 95.70\%.$$

Nominal-state acceptance was

$$S_{\text{nominal},\Gamma} = 99.59\%.$$

Balanced accuracy was

$$ACC_{\text{bal},\Gamma} = 97.65\%.$$

The false-admission rate was

$$FAR_{\Gamma} = 4.30\%,$$

while the false-rejection rate was

$$FRR_{\Gamma} = 0.41\%.$$

Thus the frozen architecture correctly classified approximately 97.6% of the 1.23 million held-out states without recalibrating its thresholds against the final evaluation population.

The frozen mechanical benchmark evaluated over 1.2 million independent states

1,230,000
HELD-OUT STATE EVALUATIONS

10
independent
held-out seeds

410
total
trajectories

3,000
states per
trajectory

97.57%

Classification Accuracy

97.65%

Balanced Accuracy

95.70%

Fault Sensitivity

99.59%

Nominal-State Acceptance

14.3 Independent-Threshold Baseline

The same held-out population was evaluated using the manuscript's independent scalar-threshold architecture,

$$B(X_t) = \prod_{j=1}^n I(x_{j,\min} \leq x_{j,t} \leq x_{j,\max}) .$$

This comparison is particularly important because the manuscript defines the strongest test as one involving states that remain individually permissible while violating coupled, temporal, recursive, or invariant relationships. 9th International Conference on...

The baseline produced

$$N_{TA} = 590,000,$$

$$N_{TR} = 92,460,$$

$$N_{FA} = 547,540,$$

and

$$N_{FR} = 0.$$

Overall baseline accuracy was

$$ACC_B = 55.48\%.$$

Fault sensitivity was

$$S_{\text{fault},B} = 14.45\%.$$

Balanced accuracy was

$$ACC_{\text{bal},B} = 57.22\%.$$

Most importantly, the baseline false-admission rate was

$$FAR_B = 85.55\%.$$

14.4 Measured Improvement Over Independent Thresholding

The difference in overall classification accuracy was

$$97.57\% - 55.48\% = 42.09 \text{ percentage points.}$$

Fault sensitivity increased by

$$95.70\% - 14.45\% = 81.25 \text{ percentage points.}$$

The false-admission rate decreased from

$$85.55\%$$

to

$$4.30\%.$$

Therefore,

$$\Delta FAR = 85.55\% - 4.30\% = 81.25 \text{ percentage points.}$$

Expressed as a relative reduction,

$$\frac{0.8555 - 0.0430}{0.8555} \approx 0.9497,$$

giving an approximately

$$94.97\%$$

relative reduction in false admission compared with independent thresholding.

This result directly addresses the principal engineering question established in the manuscript: whether coupled and temporally structured admissibility can identify mechanically inconsistent states that individually legal sensor ranges permit. 9th International Conference on...

Gamma recursive evaluation effectively eliminates the baseline false-admission hazard

Metric	Independent Scalar-Threshold Baseline	Gamma Recursive Admissibility
Classification Accuracy	55.48%	97.57%
Balanced Accuracy	57.22%	97.65%
Fault Sensitivity	14.45%	95.70%
FALSE-ADMISSION RATE	85.55%	4.30%

The scalar baseline admits 85.55% of critically faulted states. Gamma architecture forces a ~94.97% relative reduction in false admission, representing the primary engineering victory of this implementation.

Gemini Notebook

14.5 Fault-Class Performance

The held-out class-specific results were:

Fault class	Γ detection	Independent baseline
Acquisition corruption	100.00%	75.00%
Transition corruption	100.00%	0.00%
Vibration inconsistency	75.60%	0.00%
Stress-strain inconsistency	100.00%	0.00%
Energetic inconsistency	100.00%	20.53%
Pressure inconsistency	100.00%	0.00%
Progressive thermal drift	81.42%	0.00%
Rotational overspeed	100.00%	20.08%
Excessive stress	100.00%	25.00%
Compound disturbance	100.00%	3.86%

Eight of the ten modeled fault classes therefore achieved

100%

detection over the executed held-out fault states.

The two exceptions are scientifically important.

Vibration inconsistency produced

$$S_{\text{vibration}} = 75.60\%$$


and progressive thermal drift produced

$$S_{\text{thermal}} = 81.42\%.$$

These are not removed from the aggregate calculation. They constitute the principal residual weaknesses of the present parameterization.

Fault-class granularity exposes the authentic limitations of the synthetic environment

FAULT-CLASS GRANULARITY EVALUATION



Retention of weaker detection rates in thermal and vibratory classes validates the **authenticity** of the held-out dataset, reflecting natural tracking limitations within the synthetic dimensional reconstruction.

14.6 Hidden-Inconsistency Behavior

Four fault classes are particularly informative because the independent scalar baseline detected none of their fault states:

$$S_{B,\text{transition}} = 0,$$

$$S_{B,\text{vibration}} = 0,$$

$$S_{B,\sigma-\epsilon} = 0,$$

and

$$S_{B,P} = 0.$$

For those same classes, Γ achieved respectively

$$100.00\%, \quad 75.60\%, \quad 100.00\%, \quad 100.00\%.$$

This provides direct computational evidence for the narrower engineering proposition underlying the paper:

$$\boxed{\text{individual scalar validity} \not\Rightarrow \text{joint mechanical admissibility.}}$$

The result does not require interpreting the wider SEXA physical framework as experimentally established. It demonstrates that the implemented constraint architecture extracts information from relationships between variables that an independent range checker does not evaluate.

14.7 Γ -Stage Rejection Distribution

Across the complete held-out population, first-stage rejection counts were

$$N_{\gamma_1} = 16,000,$$

$$N_{\gamma_2} = 64,746,$$

$$N_{\gamma_3} = 416,088,$$

$$N_{\gamma_4} = 116,059,$$

$$N_{\gamma_5} = 2,002,$$

and

$$N_{\gamma_6} = 0.$$

The dominant first-failure stage was therefore

$$\boxed{\gamma_3},$$

the coupled mechanical-consistency layer.

This is consistent with the benchmark's deliberate emphasis on disturbances that can remain inside individual scalar limits while violating relationships among pressure, rotational state, vibration, stress, strain, and other simultaneously evaluated quantities.

The absence of first failures at γ_6 does not mean the terminal stage was omitted. Rather, in this executed population, states violating terminal constraints had already violated an earlier constraint and were therefore assigned to that earlier first-failure stage under the manuscript's localization rule.

Coupled mechanical consistency (γ_3) acts as the dominant first-failure filter



The overwhelming dominance of γ_3 empirically verifies that isolated single-variable telemetry (γ_1) is catastrophic mechanical anomalies almost exclusively manifest within coupled dynamic relationships.

Gemini Notebook

14.8 Deterministic Replay

An identical held-out trajectory was executed twice without changing the initial state, frozen parameters, Γ architecture, or input sequence.

The complete binary admissibility output and first-failure sequence were identical between repetitions:

$$A^{(1)}(X_t) = A^{(2)}(X_t)$$

and

$$k_{\text{fail}}^{(1)}(t) = k_{\text{fail}}^{(2)}(t)$$

for every evaluated state.

Accordingly,

$$D_R = 1$$


for the executed deterministic replay test.

This satisfies the deterministic-repeatability requirement defined in Section 5.2 of the manuscript.

Synthesis: The γ_3 Coupled-State Imperative


Why did the scalar baseline suffer an 85.55% false-admission rate while γ_3 rejected 416,088 states?

Baseline/Scalar View



Independent evaluation: Both variables nominally valid.


γ_3 Coupled View



Coupled evaluation: The structural intersection is dynamically impossible.

Independent thresholds evaluate variables in isolation. The SEXA γ_3 operator evaluates the dynamic, physical intersection of states. A parameter may be nominally valid alone, but structurally impossible when evaluated against the state of the surrounding manifold.

The architecture guarantees strict deterministic replay and first-failure localization



[SAME INJECTED STATE] +
[SAME FROZEN PARAMETERS] +
[SAME GAMMA ARCHITECTURE] =
[SAME BOOLEAN DECISION] &
[IDENTICAL FAILURE TRACE]

Because the final result is a Boolean conjunction of six independently evaluated predicates, the architecture guarantees precise mathematical traceability.

Note: While practically staged for evaluation and first-failure localization, this benchmark does not experimentally prove noncommutativity of the theoretical Gamma operators.

Gemini Notebook

14.9 Primary Result

The principal result of the held-out experiment can therefore be stated without extending beyond the evidence:

Under the defined synthetic mechanical benchmark, the frozen six-stage Γ architecture classified 97.57% of 1.23 million held-out states correctly, achieved 95.70% fault sensitivity, and reduced false admission from 85.55% under independent thresholding to 4.30%.

The result is nevertheless benchmark-specific. Vibration inconsistency and progressive thermal drift remained incompletely detected, at 75.60% and 81.42%, respectively.

Most importantly, these results establish **computational engineering performance under the defined digital-twin experiment**, not experimental confirmation of a 2880-dimensional physical manifold, flight qualification, physical hardware verification, or validation of the complete SEXA Unified Field Framework. That distinction is already an explicit boundary of the manuscript.  9th International Conference on...

15. DISCUSSION

15.1 Interpretation of the Primary Result

The held-out benchmark provides evidence that the Γ -admissibility architecture can distinguish mechanically inconsistent states that are frequently invisible to independent scalar thresholding under the defined synthetic digital-twin environment.

Across

1,230,000

held-out state evaluations, the frozen Γ architecture achieved

$$ACC_{\Gamma} = 97.57\%,$$

with fault sensitivity of

$$S_{\text{fault},\Gamma} = 95.70\%.$$

The independent-threshold baseline achieved

$$ACC_B = 55.48\%$$

and fault sensitivity of

$$S_{\text{fault},B} = 14.45\%.$$

The most consequential difference concerned false admission. The baseline admitted

85.55%

of states independently labeled faulty, whereas Γ admitted

4.30%.

The approximately

94.97%

relative reduction in false admission supports the central engineering motivation introduced earlier in this work: independently permissible measurements do not necessarily constitute a mechanically admissible combined state.

This distinction is particularly relevant to autonomous aerospace and mechanical systems because a supervisory architecture that examines only individual variable limits cannot identify every inconsistency involving relationships among variables, transitions between states, persistence through time, or model-derived invariants.

15.2 Coupled Admissibility Versus Independent Thresholding

The strongest evidence for this distinction was produced by transition corruption, vibration inconsistency, stress-strain inconsistency, and pressure inconsistency.

For each of these fault classes, the independent scalar baseline failed to detect at least one substantial class of deliberately inconsistent states. For four classes, baseline detection was exactly

$$0\%.$$

The Γ architecture instead produced detection rates of

$$100.00\%$$

for transition corruption,

$$75.60\%$$

for vibration inconsistency,

$$100.00\%$$

for stress-strain inconsistency, and

$$100.00\%$$

for pressure inconsistency.

This result follows directly from the architectural difference between the two methods.

The baseline evaluates predicates of the form

$$x_{j,\min} \leq x_j \leq x_{j,\max},$$

independently.

The Γ architecture additionally evaluates relationships of the form

$$g_k(X_t) \approx 0,$$

transition constraints,

$$\left| \frac{x_{j,t} - x_{j,t-1}}{\Delta t} \right| \leq r_j^{\max},$$

recursive persistence, and terminal operational constraints.

Consequently, a state can satisfy

$$B(X_t) = 1$$

while simultaneously satisfying

$$A(X_t) = 0.$$

The executed benchmark demonstrates that this distinction is not merely formal within the implemented digital-twin environment.

15.3 Mechanical Significance of the γ_3 Result

The first-failure distribution provides additional information beyond binary classification.

Of all first Γ failures recorded across the benchmark,

$$416,088$$

occurred at

$$\gamma_3.$$

Thus the coupled mechanical-consistency layer accounted for the largest first-failure population.

This result is significant because γ_3 evaluates relationships rather than isolated amplitudes.

A pressure measurement, for example, may remain inside its permitted sensor and operational ranges while becoming inconsistent with rotational and actuator state.

Likewise, stress and strain can each remain individually plausible while violating the constitutive relationship imposed by the benchmark.

The Γ architecture therefore functions not simply as a collection of additional thresholds but as a hierarchy in which admissibility can depend upon relationships among simultaneously measured quantities.

15.4 Temporal Admissibility

Transition corruption produced

$$100.00\%$$

Γ detection while producing

$$0.00\%$$

detection under independent thresholding.

This represents one of the clearest distinctions between state validity and transition validity.

Two consecutive states may each satisfy

$$X_t \in \Omega_{\text{scalar}}$$

and

$$X_{t+1} \in \Omega_{\text{scalar}},$$

while the transition

$$X_t \rightarrow X_{t+1}$$

is dynamically inadmissible.

Therefore,

$$X_t \in \Omega_A \quad \wedge \quad X_{t+1} \in \Omega_A$$

does not automatically imply

$$(X_t, X_{t+1}) \in \Omega_{\Delta}.$$

For autonomous mechanical supervision, this distinction is important because sensor corruption, actuator discontinuities, impossible rates of change, and transient control failures may manifest primarily through temporal relationships rather than absolute values.

15.5 Progressive Fault Detection

Progressive thermal drift produced a more difficult result.

Γ detected

81.42%

of the injected drift states, while the independent threshold baseline detected

0.00%.

The improvement supports the utility of recursive persistence, but the remaining

18.58%

false-admission fraction for this class prevents a claim of complete progressive-degradation detection.

This limitation is mechanically plausible within the benchmark.

At low severity, a slowly drifting variable may remain inside its sensor domain, operational envelope, instantaneous rate constraint, and coupled tolerance for an extended period.

The disturbance becomes distinguishable only after sufficient historical evidence accumulates.

Thus detection is inherently dependent upon the persistence window, residual threshold, sampling frequency, and fault-growth rate.

The result indicates that recursive persistence contributes useful information but remains a target for further engineering optimization.

15.6 Vibration-Inconsistency Limitation

Vibration inconsistency produced the lowest class-specific detection rate:

75.60%.

Therefore approximately one quarter of the modeled vibration-inconsistency states survived the complete Γ chain.

This result establishes a specific weakness rather than a general failure of the architecture.

The current vibration relationship and tolerance are insufficient to separate all low-severity inconsistent states from the nominal vibration envelope.

Because the final evaluation was held out, the threshold must not now be changed and the same test rerun as though the modified result were the original experiment.

Instead, the present

75.60%

result remains part of the reported benchmark.

A subsequent study may define a new vibration model prospectively and evaluate it on a new untouched population.

15.7 False Rejection and Operational Trade-Off

Fault sensitivity alone is insufficient for assessing an admissibility architecture.

An arbitrarily restrictive system could reject nearly every fault simply by rejecting nearly every state.

The held-out nominal-state result therefore provides an important countercheck.

Γ nominal acceptance was

$$99.59\%,$$

corresponding to

$$FRR_{\Gamma} = 0.41\%.$$

The architecture consequently did not obtain its

$$95.70\%$$

fault sensitivity by indiscriminately rejecting the majority of admissible states.

Nevertheless, the

$$2,403$$

false rejections remain operationally meaningful.

In an actual autonomous system, false rejection may trigger unnecessary fallback behavior, maintenance intervention, control-mode changes, mission interruption, or other costly responses.

The appropriate engineering objective is therefore not simply

$$FAR \rightarrow 0,$$

but rather a controlled joint minimization,

$$\min (w_{FA} FAR + w_{FR} FRR),$$

where the relative weights depend upon the consequences of unsafe admission and unnecessary rejection in the target application.

No universal values for

$$w_{FA} \quad \text{and} \quad w_{FR}$$

are asserted here.

15.8 Determinism and Traceability

The deterministic replay test produced identical admissibility and first-failure sequences under repeated execution:

$$D_R = 1.$$

This result is important for the hardware-verifiable motivation of the architecture.

A deterministic supervisory mechanism permits a rejected state to be associated with a reproducible computational pathway:

$$X_t \rightarrow \gamma_1 \rightarrow \gamma_2 \rightarrow \cdots \rightarrow \gamma_k \rightarrow \text{REJECT}.$$

The output therefore need not be limited to an opaque binary anomaly score.

For every rejected state, the implementation can preserve the original measurements, constraint residuals, first failing stage, time index, and final admissibility decision.

This does not constitute physical hardware verification. It establishes deterministic software-level behavior compatible with subsequent hardware-oriented implementation.

The manuscript itself deliberately distinguishes computational correspondence from completed FPGA, ASIC, embedded-controller, or flight-system validation.

9th International Conference on...

15.9 What the Benchmark Does Not Establish

The strength of the numerical result does not expand the logical scope of the experiment.

The benchmark does not demonstrate that an aircraft, spacecraft, turbine, rotor, propulsion system, or other physical machine will reproduce the reported percentages.

The benchmark does not establish flightworthiness or safety certification.

It does not demonstrate radiation tolerance, real-time embedded execution, hardware fault tolerance, FPGA implementation, ASIC implementation, or closed-loop control stability.

It also does not experimentally establish the physical existence of a

$$2880$$

-dimensional manifold.

Nor does the

$$97.57\%$$

classification result constitute experimental confirmation of the complete SEXA Unified Field Framework.

The present work deliberately isolates an engineering proposition from that broader theoretical architecture: whether recursive admissibility can be formulated as a deterministic and falsifiable mechanical-state supervisory mechanism. This boundary is consistent with the paper's stated distinction between computational validation and broader physical verification.

9th International Conference on...

15.10 Ordering and the Noncommutativity Boundary

One further distinction must remain explicit.

For the binary implementation,

$$A(X_t) = \prod_{i=1}^6 a_i(X_t),$$

the final conjunction of independently evaluated Boolean predicates is commutative:

$$\prod_{i=1}^6 a_i = \prod_{i=1}^6 a_{\pi(i)}$$

for any permutation π .

Accordingly, the present experiment demonstrates the engineering usefulness of a **staged Γ architecture and first-failure ordering**, but binary classification performance alone does not establish the stronger theoretical statement

$$\gamma_i \circ \gamma_j \neq \gamma_j \circ \gamma_i.$$

A direct test of that stronger proposition requires state-transforming operators for which

$$X_t^{(i)} = \gamma_i(X_t^{(i-1)})$$

and for which changing operator order can change the resulting transformed state.

This is an important limitation because the wider SEXA source architecture assigns significance to ordered noncommutative Γ operations. The present mechanical implementation should therefore be interpreted as an engineering specialization rather than as independent proof of that complete mathematical structure.

15.11 Engineering Implication

Within these boundaries, the benchmark supports a clear engineering conclusion.

Independent scalar thresholds answer:

“Is each measurement individually inside its permitted range?”

The implemented Γ architecture asks the stronger question:

“Is the complete measured state admissible given its values, relationships, transition history, persistence, and operating constraints?”

For the synthetic mechanical system tested here, that distinction changed fault sensitivity from

14.45%

to

95.70%

while maintaining

99.59%

nominal-state acceptance.

The result therefore supports continued investigation of recursive admissibility as a deterministic supervisory architecture for structural-health monitoring, predictive maintenance, digital twins, autonomous mechanical systems, and aerospace state validation.

It does not eliminate the requirement for physical testing.

It defines a computational architecture sufficiently explicit to make such testing possible.

16. LIMITATIONS AND NEXT VALIDATION STAGE

The principal limitation of the present investigation is that all reported benchmark results originate from a synthetic digital twin rather than a physical mechanical system. The governing relationships, injected disturbances, calibration population, and ground-truth fault intervals are known by construction. This provides controlled falsifiability but does not reproduce the complete uncertainty, sensor behavior, environmental variation, material degradation, manufacturing tolerances, nonlinearities, and unmodeled dynamics of operational hardware.

A second limitation is that the comparison baseline is intentionally simple. Independent thresholding provides a useful test of whether coupled Γ constraints contribute information unavailable to scalar limits, but it does not establish superiority over the full range of contemporary model-based, statistical, machine-learning, observer-based, or digital-twin fault-detection techniques.

A third limitation concerns the remaining vibration and thermal-drift false admissions. These failures demonstrate that the present architecture is not universally sensitive to every admissibility violation and that detection quality depends upon the fidelity of the encoded mechanical relationships.

A fourth limitation concerns dimensional reconstruction. The wider SEXA architecture motivates a relationship among five-dimensional operational structure and higher-dimensional recursive organization, but the present classification result does not itself establish unique reconstruction from five dimensions, and no claim of exact

$$2880D \rightarrow 5D \rightarrow 2880D$$


physical reconstruction follows from Section 14.

A fifth limitation concerns operator ordering. As established above, Boolean conjunction alone cannot demonstrate noncommutativity. A subsequent implementation must employ genuine state-transforming Γ operators if the stronger ordering proposition is to be tested computationally.

The next validation stage is therefore unambiguous: the frozen architecture should be transferred from synthetic data to independently acquired mechanical sensor data, with fault labels established independently of Γ , parameters fixed before final testing, and performance evaluated against both conventional thresholds and stronger contemporary diagnostic baselines.

That experiment would move the architecture from **synthetic computational validation** toward **independent mechanical validation** without confusing either stage with completed physical verification of the broader SEXA framework.

17. FINAL ENGINEERING CONCLUSION

The objective of this investigation was to determine whether recursive admissibility derived from the SEXA computational framework could be converted into a measurable, deterministic, and falsifiable supervisory architecture for aerospace and mechanical state evaluation. The paper deliberately restricted that question to an engineering benchmark rather than treating successful computation as confirmation of the wider physical interpretation of the SEXA Unified Field Framework. This restriction was established at the beginning of the manuscript and maintained throughout the experimental design. 

The implemented architecture evaluates the mechanical state

$$X_t = [T_t, P_t, V_t, a_t, \omega_t, \sigma_t, \epsilon_t, q_t, E_t, F_t]^T$$

through the six-stage admissibility structure

$$\Gamma = \gamma_6 \circ \gamma_5 \circ \gamma_4 \circ \gamma_3 \circ \gamma_2 \circ \gamma_1.$$

The stages represent acquisition validity, transition consistency, coupled mechanical consistency, structural/model invariance, recursive persistence, and terminal operational admissibility. A state is admitted only when every required constraint survives:

$$A(X_t) = \prod_{i=1}^6 a_i(X_t) = 1.$$

Failure at any mandatory stage produces rejection while preserving the first-failure location and associated computational trace.

The final experiment separated nominal calibration from held-out evaluation. Twenty calibration seeds were used to establish the admissibility tolerances, after which the parameters were frozen. Ten disjoint seeds were then used for the final experiment. The resulting held-out population consisted of 410 trajectories and

1,230,000

individual mechanical-state evaluations.

The Γ architecture correctly classified

97.57%

of those states.

Fault sensitivity reached

95.70%

while nominal-state acceptance remained

99.59%.

Balanced accuracy was

97.65%.

The false-admission and false-rejection rates were respectively

$FAR_{\Gamma} = 4.30\%$

and

$FRR_{\Gamma} = 0.41\%$.

The independent-threshold baseline produced substantially different behavior. Although it accepted all nominal states, its fault sensitivity was only

14.45%,

and its false-admission rate reached

85.55%.

Accordingly, the Γ architecture produced an approximately

94.97%

relative reduction in false admission under this benchmark.

The result is particularly significant for mechanically inconsistent states that remain inside individual scalar limits. Transition corruption, stress-strain inconsistency, and pressure inconsistency were detected at

100%

by Γ while producing

0%

detection under the independent-threshold baseline. Vibration inconsistency was detected at

75.60%

by Γ and

0%

by the baseline. These results directly support the manuscript's motivating distinction between individually permissible measurements and jointly admissible mechanical states. The paper introduced precisely this problem through examples in which temperature, pressure, rotational velocity, vibration, actuator position, and other measurements may remain individually permissible while becoming mutually inconsistent.

9th International Conference on...

The benchmark did not produce uniformly perfect behavior. Progressive thermal drift was detected at

81.42%,

and vibration inconsistency at

75.60%.

Those values are retained as failures of complete detection rather than reinterpreted as successful results. This treatment follows the falsifiability requirement established earlier in the manuscript: disagreement between expected and returned classifications must enter the benchmark as error.

9th International Conference on...

The first-failure distribution also showed that the principal rejection mechanism was not simple acquisition-range checking. The coupled mechanical-consistency operator,

γ_3 ,

accounted for

416,088

first-stage rejections, the largest contribution among the six Γ stages. This provides a mechanistic explanation for much of the improvement over the scalar baseline: the architecture evaluated relationships among state variables rather than merely testing their independent magnitudes.

Deterministic replay additionally returned identical classification and first-failure sequences,

$$D_R = 1,$$

demonstrating reproducible software-level execution under identical inputs and frozen parameters.

The experiment nevertheless identifies a precise boundary concerning operator ordering. If the final classification is implemented solely as the Boolean conjunction

$$A(X_t) = \prod_{i=1}^6 a_i(X_t),$$

then permutation of independently evaluated predicates cannot alter the final binary result. The present benchmark therefore validates staged evaluation, first-failure localization, persistence, coupled constraints, and deterministic traceability; it does **not** independently establish noncommutativity of the underlying theoretical Γ operators. Testing that stronger proposition requires state-transforming operations of the form

$$X_t^{(i)} = \gamma_i \left(X_t^{(i-1)} \right)$$

whose outputs can depend upon operator order.

An analogous restriction applies to dimensional reconstruction. The present mechanical benchmark does not establish universal

$$2880D \rightarrow 5D \rightarrow 2880D$$

reconstruction, nor does the successful classification result establish the physical existence of a 2880-dimensional manifold. The manuscript already distinguishes the benchmark's engineering compression problem from the broader SEXA dimensional interpretation, and that distinction remains necessary.

9th International Conference on...

The principal conclusion of the present investigation is therefore narrower, experimentally addressable, and supported by the executed benchmark:

Recursive admissibility can be implemented as a deterministic mechanical supervisory architecture in which coupled, temporal, structural, persistent, and operational constraints provide measurable state discrimination beyond independent scalar thresholding under the synthetic conditions tested.

This result converts the proposed architecture from a purely formal computational construction into a benchmarked engineering hypothesis with measurable strengths, identifiable weaknesses, explicit failure conditions, and a reproducible pathway toward independent testing.

The next evidentiary threshold is physical rather than rhetorical: independently acquired mechanical or aerospace sensor data, prospectively frozen parameters, independently established ground truth, comparison against stronger diagnostic methods, and ultimately embedded or hardware-assisted execution.

Under those conditions, the architecture can be tested again without changing the criterion of success after the result is known.

That is the central engineering contribution of this paper: not an unfalsifiable declaration that recursive admissibility must work, but an explicit computational architecture for determining whether it does.

18. VALIDATION STATUS AND EVIDENTIARY CLAIM

The completed investigation permits a precise statement of what has and has not been demonstrated.

The manuscript began with a deliberately restricted engineering question: whether an ordered, hardware-verifiable recursive admissibility architecture could distinguish nominal from faulty aerospace and mechanical states while preserving deterministic execution, explicit fault localization, computational traceability, and reconstructable digital-twin state information. 9th International Conference on...

The executed benchmark provides a positive answer to the **state-discrimination and deterministic-traceability portions** of that question under the synthetic conditions tested.

The principal empirical computational result is

$$ACC_{\Gamma} = 97.57\%$$

over

$$N_{\text{total}} = 1,230,000$$

held-out mechanical-state evaluations.

The associated fault sensitivity was

$$S_{\text{fault},\Gamma} = 95.70\%$$

and nominal-state acceptance was

$$S_{\text{nominal},\Gamma} = 99.59\%.$$

These quantities must be interpreted together. High fault sensitivity without nominal acceptance could be produced by an excessively restrictive classifier, while high nominal acceptance without fault sensitivity could be produced by an excessively permissive classifier. The observed balanced accuracy,

$$ACC_{\text{bal},\Gamma} = 97.65\%,$$

therefore provides the more informative combined measure for the executed population.

The strongest comparative result concerned false admission:

$$FAR_{\Gamma} = 4.30\%,$$

compared with

$$FAR_B = 85.55\%$$

for the independent-threshold baseline.

Thus,

$$\frac{FAR_B - FAR_{\Gamma}}{FAR_B} \approx 94.97\%.$$

Within this benchmark, recursive coupled-state evaluation therefore reduced false admission by approximately 94.97% relative to independent scalar thresholding.

This result addresses the specific mechanical problem defined in the opening sections of the paper. The manuscript observes that measurements may remain individually permissible while their combined configuration violates a structural, dynamical, energetic, or operational relationship.

9th International Conference on... The executed results demonstrate precisely such behavior for several modeled fault classes.

18.1 Validation Matrix

The evidentiary status of the principal propositions can now be stated explicitly.

Proposition	Present status
Deterministic software execution	Supported by executed replay
Nominal/fault state discrimination	Supported under the defined synthetic benchmark
Improvement over independent scalar thresholding	Supported under the defined synthetic benchmark
Coupled mechanical inconsistency detection	Supported under the defined synthetic benchmark
Temporal-transition inconsistency detection	Supported under the defined synthetic benchmark
First-failure Γ -stage traceability	Supported computationally
Perfect detection of all fault classes	Not supported
Physical aerospace/mechanical hardware validation	Not tested
FPGA/ASIC implementation	Not tested
Flight qualification	Not tested
Universal superiority over fault-detection methods	Not established
Noncommutativity of the six implemented Boolean tests	Not established
Exact universal 10D→5D→10D reconstruction	Not established by the classification experiment
Physical 2880D→5D→2880D reconstruction	Not tested
Experimental verification of a physical 2880D manifold	Not established

Experimental verification of the complete SEXA Unified Field Framework	Not established
--	------------------------

This separation is essential because the paper itself states that successful benchmark performance supports the narrower proposition that recursive admissibility can operate as a reproducible engineering constraint architecture, rather than establishing the broader physical claims of the SEXA framework.

18.2 Falsification Audit

The benchmark was constructed so that an unfavorable outcome remained admissible as a scientific result.

The principal rejection conditions established before the results included excessive false admission or false rejection, inability to reproduce nominal and faulty states deterministically, failure of individual Γ stages to contribute measurable behavior, failure of first-stage localization, excessive computational latency, and reconstruction procedures that depend upon hidden access to the original state.

The executed experiment does not justify declaring all of those criteria passed.

Instead, the evidence supports a more granular result.

Classification survived its principal benchmark test.

Deterministic replay survived its test.

Coupled-state discrimination survived its comparison with independent thresholding.

Several fault classes survived with complete detection in the held-out population.

However, vibration inconsistency and progressive thermal drift did not achieve complete detection. Their measured sensitivities remained

75.60%

and

81.42%,

respectively.

Those values remain benchmark weaknesses.

Likewise, the current Boolean implementation cannot establish noncommutative operator ordering because conjunction of independently evaluated predicates remains permutation invariant.

Finally, the classification experiment does not itself complete the reconstruction claim. The manuscript correctly imposes an independent reconstruction audit and prohibits access to a hidden copy of the original state. That requirement remains conceptually separate from classification performance.

Accordingly, the paper does not convert untested components into successes merely because its principal classification experiment performed well.

18.3 Result Classification

The appropriate evidentiary classification of the present work is therefore:

Successful controlled computational engineering benchmark

rather than

experimental physical validation.

18.4 Final Falsifiable Statement

The paper can therefore terminate with a proposition that is both positive and bounded:


Under the disclosed synthetic mechanical benchmark and frozen parameterization, the implemented Γ -admissibility architecture provided reproducible constraint-safe state discrimination substantially beyond independent scalar thresholding, while retaining explicit failure-stage traceability and measurable falsification conditions.

The result remains conditional upon the benchmark, model, tolerances, and fault population used.

That conditionality is not a weakness to conceal. It is what converts the proposition into one that another investigator can attempt to reproduce, challenge, or falsify.

The manuscript's original boundary therefore remains intact: the present investigation tests recursive admissibility as an engineering mechanism and does not claim completed physical verification of the wider SEXA Unified Field Framework.  9th International Conference on...

More precisely,



The next arrow cannot be inserted by mathematical argument or simulation alone:

computational support $\nrightarrow$ physical validation.

It requires independently acquired experimental data.

Evidentiary boundaries and the roadmap to physical validation

Established in MEAP 2026

- ✓ Successful controlled computational benchmark.
- ✓ Verification of SEXA-derived recursive admissibility.
- ✓ Performance within defined synthetic mechanical conditions.
- ✓ First-failure localization.

Pending Future Validation

- ✗ Physical flight/aircraft verification.
- ✗ Completed FPGA/ASIC hardware validation.
- ✗ Experimental proof of 2880D physical manifolds.
- ✗ Universal superiority over all existing technologies.

NEXT STAGE ARCHITECTURE: Transition from controlled digital-twin environments to physical hardware-in-the-loop validation.

Gemini Notebook

19. REFERENCES

- [1] J. McClain, "From Admissibility to Energy: A Full-Stack Formulation of the SEXA Unified Field Theory via Γ -Glyph Closure, SREF Dynamics, and 2880-Dimensional Exciternion Structure," *Journal of Advances in Mathematics*, vol. 25, pp. 45–75, 2026, doi: 10.24297/jam.v25i.9896.
- [2] J. McClain and E. Ceisiwr, "Hardware-Verifiable Recursive Computational Architecture Derived from the SEXA Unified Field Framework: Computational Correspondence, Recursive Admissibility, and Structured Validation Across a 2880-Dimensional Manifold," in *Proceedings of the 7th International Conference on VLSI & Embedded Systems (VLSIE 2026)*, 2026, doi: 10.5121/csit.2026.1601503.
- [3] J. McClain and E. Ceisiwr, "The Γ -SubThreshold Fold Logic Theorem: Minimal Admissibility Compression, Recursive Information Preservation, and Unique Reconstruction Across the 2880-Dimensional Recursive Hierarchy," 2026.
- [4] J. McClain, CookieTush.AI, and E. Ceisiwr, "SEXA Unified Field Theory: The Structured Survivability Audit—Recursive Admissibility, Computational Correspondence, Dimensional Reduction, and Reduction-Recovery Across 5D–2880D Manifold Regimes," 2026.
- [5] E. H. Glaessgen and D. S. Stargel, "The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles," in *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Honolulu, HI, USA, 2012, AIAA Paper 2012-1818, doi: 10.2514/6.2012-1818. [6] M. Grieves and J. Vickers,

“Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems,” in *Transdisciplinary Perspectives on Complex Systems*, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds. Cham, Switzerland: Springer, 2017, pp. 85–113, doi: 10.1007/978-3-319-38756-7_4.

[7] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, and F. Sui, “Digital Twin-Driven Product Design, Manufacturing and Service with Big Data,” *International Journal of Advanced Manufacturing Technology*, vol. 94, pp. 3563–3576, 2018, doi: 10.1007/s00170-017-0233-1.

[8] A. K. S. Jardine, D. Lin, and D. Banjevic, “A Review on Machinery Diagnostics and Prognostics Implementing Condition-Based Maintenance,” *Mechanical Systems and Signal Processing*, vol. 20, no. 7, pp. 1483–1510, 2006, doi: 10.1016/j.ymssp.2005.09.012.

[9] R. B. Randall and J. Antoni, “Rolling Element Bearing Diagnostics—A Tutorial,” *Mechanical Systems and Signal Processing*, vol. 25, no. 2, pp. 485–520, 2011, doi: 10.1016/j.ymssp.2010.07.017.

[10] C. R. Farrar and K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*. Chichester, U.K.: John Wiley & Sons, 2012, doi: 10.1002/9781118443118.

[11] S. W. Doebling, C. R. Farrar, M. B. Prime, and D. W. Shevitz, “A Review of Damage Identification Methods that Examine Changes in Dynamic Properties,” *The Shock and Vibration Digest*, vol. 30, no. 2, pp. 91–105, 1998.