Complete Mathematical Formalization of AHP

Proofs of Non-Zeno Guarantees and Stability Analysis

Alexandre Sah*

January 30, 2026

Abstract

This online resource provides complete mathematical formalization of Algorithmic Hysteresis Primacy (AHP), including formal proofs of the Non-Zeno guarantee (minimum hesitation time), stability analysis using Lyapunov methods, and characterization of noise rejection properties. We demonstrate that AHP provides architectural guarantees of minimum hesitation time $\Delta T_{\min} > 0$ that are independent of implementation details or runtime conditions. These mathematical results translate the philosophical concept of "decisional inertia" into provable system properties that enable distributed governance mechanisms across multiple domains. The theorems herein correspond to guarantees cited in the accompanying article and support the technical implementations in supplementary materials.

Keywords: Hysteresis control, Non-Zeno guarantee, Lyapunov stability, Real-time systems, Temporal logic, Distributed governance, Protocol specifications

Contents

*This document provides supplementary mathematical formalization supporting the main article *Algorithmic Hysteresis Primacy (AHP): Temporal Sovereignty in AI Governance*, available at zmem.org and SSRN. The main article presents a complete, self-contained conceptual framework; this material offers technical depth for implementation and formal verification. Additional supplementary materials (protocol specifications, reference implementations, governance frameworks, and validation protocols) are available at zmem.org and SSRN.

1 Introduction: From Concept to Formal Guarantees

The accompanying article *"Algorithmic Hysteresis Primacy (AHP): Temporal Sovereignty in AI Governance"* introduces Algorithmic Hysteresis Primacy (AHP) as an architectural response to the *Speed-Responsibility Paradox.* That text develops a complete, self-contained conceptual and philosophical framework that does not require supplementary materials for comprehension. This document provides the mathematical machinery that renders AHP implementable with provable guarantees for those seeking technical depth beyond the conceptual presentation.

This document establishes three foundational results:

1. **Non-Zeno Guarantee** (Theorem 3.1): State transitions cannot occur arbitrarily fast; there exists a strictly positive minimum time $\Delta T_{\min} > 0$ between decisions.
2. **Stability** (Theorems 4.1 and 4.2): The system is stable in the sense of Lyapunov, with bounded-input bounded-output (BIBO) guarantees.
3. **Noise Rejection** (Theorem 5.1): High-frequency noise is intrinsically filtered by the accumulation mechanism.

Remark 1.1 (Mathematical Foundations for Distributed Governance and Implementation)**.** The formal guarantees established herein—particularly the architectural impossibility of Zeno behavior (Theorem 3.2) and the deterministic bounds on switching frequency (Theorem 3.3)—provide the temporal substrate necessary for protocol specifications (PHA-Hysteresis and ZMEM-Ethics headers), reference implementations (C and Python), Byzantine consensus protocols, and radiation-aware computing validation protocols detailed in supplementary materials.

2 System Dynamics and Definitions

2.1 Notation and State Space

We consider a dynamical system operating in continuous time $t \in \mathbb{R}_{\geq 0}$.

Definition 2.1 (System State)**.** The complete state of an AHP system at time t consists of:

- $I(t) \in [0, I_{\max}] \subset \mathbb{R}_{\geq 0}$: Internal accumulator value (evidence mass)
- $D(t) \in \{0, 1\}$: Binary decision state (0 = Wait, 1 = Act)
- $\varepsilon(t) \in \mathbb{R}$: Error signal or evidence input (external)

The activation function $\Phi : \mathbb{R} \to [0, \Phi_{\max}]$ maps evidence to accumulation rate.

2.2 Bounded Accumulator Dynamics

The evolution of the accumulator is governed by:

$$I(t) = \text{clamp}\left(I(t_0) + \int_{t_0}^{t} \Phi(\varepsilon(\tau))\, d\tau,\ 0,\ I_{\max}\right) \tag{1}$$

where the clamping function ensures saturation:

$$\text{clamp}(x, a, b) = \min(\max(x, a), b) \tag{2}$$

The critical architectural constraint is the **Governance Cap**:

$$\forall \varepsilon \in \mathbb{R}, \quad 0 \leq \Phi(\varepsilon) \leq \Phi_{\max} < \infty \tag{3}$$

Remark 2.1 (Governance Cap as Ethical Parameter)**.** The bound $\Phi_{\max}$ is *not* merely a physical constraint (e.g., sensor bandwidth) but a *deliberate design choice* encoding ethical and stability requirements. This parameter operationalizes the "right to hesitation" discussed in the accompanying article and serves as a **governance mechanism**—analogous to a speed limit—that constrains system velocity regardless of underlying computational capability. This architectural choice enables the "temporal non-compressibility" required for meaningful human oversight, protocol-level hesitation enforcement, and radiation-aware computing where transient errors must be temporally bounded.

2.3 Hysteretic Decision Rule

The decision output follows a Schmitt trigger (hysteretic comparator):

$$D(t) = \begin{cases} 1 & \text{if } I(t) \geq \Gamma_{\max} \quad \text{(Act)} \\ 0 & \text{if } I(t) \leq \Gamma_{\min} \quad \text{(Wait)} \\ D(t^-) & \text{otherwise} \quad \text{(Hold)} \end{cases} \tag{4}$$

where $0 < \Gamma_{\min} < \Gamma_{\max} < I_{\max}$ define the **hysteresis band**, and $D(t^-)$ denotes the left-hand limit (previous state).

Definition 2.2 (Hysteresis Regions)**.** Equation (4) partitions the state space:

- **Act Region**: $\mathcal{A} = [\Gamma_{\max}, I_{\max}]$ (commitment enforced)
- **Wait Region**: $\mathcal{W} = [0, \Gamma_{\min}]$ (quiescence enforced)
- **Hysteresis Band**: $\mathcal{H} = (\Gamma_{\min}, \Gamma_{\max})$ (memory preserved)

2.4 Assumptions and Modeling Scope

Assumption 2.1 (Continuous-Time Approximation)**.** The continuous-time model in (1) provides a valid approximation for digital implementations when $\Delta T_{\min} \gg T_s$, where T_s is the system clock period. For typical parameter values ($\Delta T_{\min} \geq$ 50ms) and modern computing ($T_s \sim$ 1ns), this assumption holds with high fidelity. Discrete-time implementations maintain these guarantees with appropriate scaling (see Theorem 6.1).

Assumption 2.2 (Constant Governance Cap)**.** The governance cap $\Phi_{\max}$ is treated as constant during operation. While adaptive variants are possible, the core guarantees require $\Phi_{\max}$ to be fixed or bounded during any decision interval.

Assumption 2.3 (Accumulator Integrity)**.** The accumulator implementation is free from catastrophic failures (e.g., overflow, corruption). This is ensured through standard engineering practices (bounds checking, watchdog timers) detailed in reference implementations and validated through fault injection protocols.

3 The Non-Zeno Guarantee

3.1 Main Theorem

Theorem 3.1 (Architectural Non-Zeno Property)**.** For the system defined by (1) and (4), any state transition $D : i \to j$ (where $i \neq j$) requires minimum time:

$$\Delta T_{\min} = \frac{\Gamma_{\max} - \Gamma_{\min}}{\Phi_{\max}} > 0 \tag{5}$$

Proof. Consider the transition $D : 0 \to 1$ (Wait $\to$ Act). By (4), this requires $I(t_2) \geq \Gamma_{\max}$ at some exit time t_2. Since the system was in state $D = 0$, there exists an entry time $t_1 < t_2$ when $I(t_1) \leq \Gamma_{\min}$.

The accumulator must traverse the entire hysteresis band:

$$\Delta I = I(t_2) - I(t_1) \geq \Gamma_{\max} - \Gamma_{\min} \tag{6}$$

From (1) and (3), the rate of change is bounded:

$$\frac{dI}{dt} = \Phi(\varepsilon(t)) \leq \Phi_{\max} \tag{7}$$

Integrating over $[t_1, t_2]$:

$$\Delta I = \int_{t_1}^{t_2} \frac{dI}{dt}\, d\tau \leq \Phi_{\max} \cdot (t_2 - t_1) \tag{8}$$

Rearranging and combining with the traversal requirement:

$$t_2 - t_1 \geq \frac{\Gamma_{\max} - \Gamma_{\min}}{\Phi_{\max}} = \Delta T_{\min} \tag{9}$$

Since $\Gamma_{\max} > \Gamma_{\min}$ and $\Phi_{\max} > 0$ by construction, $\Delta T_{\min} > 0$ strictly. The reverse transition $D : 1 \to 0$ follows by symmetry (descending traversal). □

Corollary 3.2 (Zeno Impossibility)**.** Infinite state transitions in finite time (Zeno behavior) is architecturally impossible under AHP. This guarantee is critical for preventing pathological synchronization in distributed systems, bounding radiation-induced transient error propagation, and ensuring protocol-level temporal guarantees.

Corollary 3.3 (Maximum Switching Frequency)**.** The maximum toggle frequency is bounded:

$$f_{\max} = \frac{1}{2\Delta T_{\min}} = \frac{\Phi_{\max}}{2(\Gamma_{\max} - \Gamma_{\min})} \tag{10}$$

This provides deterministic performance bounds essential for hard real-time system certification, radiation-aware computing where SEU rates must be bounded, and protocol timing constraints.

3.2 Numerical Illustration

Example 3.1 (High-Frequency Trading Parameters)**.** Consider a financial trading system with:

- $\Phi_{\max} = 1000$ units of evidence per second
- $\Gamma_{\max} - \Gamma_{\min} = 50$ units (hysteresis band)

The minimum hesitation time is:

$$\Delta T_{\min} = \frac{50}{1000} = 0.05\,\text{s} = 50\,\text{ms}$$

This aligns with circuit-breaker regulations (50-200 ms) cited in the accompanying article. The maximum switching frequency is:

$$f_{\max} = \frac{1}{2 \times 0.05} = 10\,\text{Hz}$$

These parameters map directly to the PHA-Hysteresis protocol specification (`delay=50` parameter).

Example 3.2 (Radiation-Aware Computing Parameters)**.** For space systems applications, typical parameters might be:

- $\Phi_{\max} = 200$ units/s (conservative for radiation environments)
- $\Gamma_{\max} - \Gamma_{\min} = 40$ units (wider hysteresis for SEU resilience)

Yielding:

$$\Delta T_{\min} = \frac{40}{200} = 0.2\,\text{s} = 200\,\text{ms}$$

This provides sufficient temporal buffer for radiation-induced transient errors to settle without triggering spurious state transitions, as validated in falsification protocols.

3.3 Physical Interpretation

Remark 3.1 (Digital Inertia)**.** Theorem 3.1 establishes a computational analogue to Newton's First Law. Just as physical mass resists instantaneous velocity changes, the hysteresis band $\Delta\Gamma = \Gamma_{\max} - \Gamma_{\min}$ acts as *inertial mass* in the decision process, while $\Phi_{\max}$ acts as the maximum applicable "force" (evidence rate). This "digital inertia" is what enables:

- *Reversible decisions* during the accumulation window—a necessary condition for the Zero-Mean Execution Memory (ZMEM) concept developed in the accompanying article
- *Protocol-level hesitation* in PHA-Hysteresis headers
- *Operationalization in distributed governance* scenarios
- *Radiation resilience* by providing temporal buffering against transient errors

4 Stability Analysis

4.1 Lyapunov Stability

We demonstrate that AHP systems converge to stable equilibria under constant inputs.

Definition 4.1 (Lyapunov Candidate)**.** For constant input ε_0 with corresponding equilibrium I_{eq}, define:

$$V(I) = \frac{1}{2}(I - I_{\text{eq}})^2 \tag{11}$$

Theorem 4.1 (Asymptotic Stability in $\mathcal{H}$)**.** The AHP system is Lyapunov stable with respect to the hysteresis band $\mathcal{H}$. For any bounded input, the accumulator converges to either: (i) saturation at $I_{\max}$ (Act), (ii) quiescence at 0 (Wait), or (iii) a limit cycle within $\mathcal{H}$.

Proof. $V(I)$ is positive definite and radially unbounded. Within $\mathcal{H}$, the time derivative along trajectories is:

$$\dot{V} = (I - I_{\text{eq}})\Phi(\varepsilon_0) \tag{12}$$

The hysteretic rule (4) ensures that once I exits $\mathcal{H}$, it remains in the corresponding region ($\mathcal{A}$ or $\mathcal{W}$) until evidence reverses sufficiently to cross the opposite threshold. This prevents chattering (oscillation around a single threshold) characteristic of simple relay controllers.

Remark 4.1 (Scope of Analysis and Implementation)**.** This stability result assumes constant or slowly varying inputs $\varepsilon(t)$. For rapidly changing inputs, the system exhibits bounded oscillations but remains BIBO stable per Theorem 4.2. In radiation environments, this stability guarantee ensures bounded response to transient faults. Reference implementations preserve this stability property in discrete-time systems.

□

4.2 Bounded-Input Bounded-Output (BIBO) Stability

Theorem 4.2 (BIBO Guarantee)**.** For any input signal with $|\varepsilon(t)| \leq M < \infty$, the accumulator satisfies:

$$0 \leq I(t) \leq I_{\max}, \quad \forall t \geq 0 \tag{13}$$

Proof. Immediate from the clamping operation in (1). □

Remark 4.2 (Implications for Fault Tolerance and Governance)**.** The BIBO guarantee ensures that even under worst-case evidence inputs (including radiation-induced transients), the accumulator state remains bounded within $[0, I_{\max}]$. This property is crucial for ensuring system safety under fault conditions, providing deterministic bounds for regulatory compliance, enabling predictable behavior in distributed consensus, and supporting reliable protocol implementations.

5 Noise Rejection Properties

5.1 High-Frequency Attenuation

The bounded accumulation rate $\Phi_{\max}$ acts as an intrinsic low-pass filter.

Proposition 5.1 (Noise Rejection)**.** Consider zero-mean noise $n(t)$ with characteristic frequency $f_n \gg 1/\Delta T_{\min}$. The probability of noise-induced state transition vanishes as $f_n \to \infty$:

$$\lim_{f_n \to \infty} \mathbb{P}[\text{transition} \mid n(t)] = 0 \tag{14}$$

Proof Sketch. The accumulator integrates noise over the minimum window $\Delta T_{\min}$:

$$\Delta I_{\text{noise}} = \int_{t}^{t+\Delta T_{\min}} \Phi(n(\tau)) d\tau \tag{15}$$

For zero-mean high-frequency noise, positive and negative contributions cancel (by the Law of Large Numbers). The traversal threshold $\Delta\Gamma$ requires sustained bias, which random noise cannot provide over $\Delta T_{\min}$.

Remark 5.1 (Formal Treatment and Practical Implications)**.** A more rigorous proof can be constructed using spectral analysis: the transfer function from noise to accumulator state has a low-pass characteristic with cutoff frequency $f_c = 1/(2\pi\Delta T_{\min})$. For $f_n \gg f_c$, the attenuation factor is approximately $(f_c/f_n)^2$. This property is essential for rejecting radiation-induced transient noise in space systems, filtering high-frequency market noise in financial applications, providing robustness against sensor artifacts in medical AI systems, and implementing reliable network protocols.

□

5.2 Comparison with Relay Controllers

Remark 5.2 (Governance Advantages and Protocol Integration)**.** The comparison in Table 1 highlights AHP's architectural advantages for governance applications. The guaranteed $\Delta T_{\min} > 0$ enables meaningful human oversight intervals, regulatory compliance with temporal requirements, distributed coordination without pathological synchronization, protocol-level hesitation enforcement, and validation through falsifiable protocols. These advantages are realized through reference implementations provided in supplementary materials.

Table 1: Comparison: AHP vs. Standard Relay Control

Property	**Relay (Bang-Bang)**	**AHP**
Minimum switching time	$\Delta T = 0$	$\Delta T_{\min} > 0$ (guaranteed)
Chattering (Zeno)	Possible	Architecturally impossible
Noise sensitivity	High (threshold crossing)	Low (integration filter)
Implementation	Simple (comparator)	Moderate (accumulator)
Minimum enforced inter-decision time	None	Guaranteed ($\Delta T_{\min} > 0$)
Protocol support	Limited	Native (PHA-Hysteresis headers)

6 Extensions and Practical Considerations

6.1 Discrete-Time Implementation

While the analysis uses continuous-time models, practical implementations are discrete. Let T_s be the sampling period.

Proposition 6.1 (Discrete-Time Guarantee)**.** For a discrete-time implementation with sampling period T_s, the minimum hesitation time is bounded by:

$$\Delta T_{\min}^{\text{disc}} \geq \max\left(\Delta T_{\min}, T_s\right)$$

where $\Delta T_{\min}$ is the continuous-time bound from (5).

Proof. The accumulator update becomes:

$$I[k+1] = \text{clamp}\left(I[k] + \Phi(\varepsilon[k])T_s, 0, I_{\max}\right)$$

The minimum time to traverse $\Delta\Gamma$ is at least $\lceil \Delta\Gamma / (\Phi_{\max} T_s) \rceil \cdot T_s \geq \Delta T_{\min}$. □

Remark 6.1 (Implementation Guidance and Reference Code)**.** This guarantee ensures that digital implementations maintain the core temporal properties. Implementation details and reference code (C and Python) are provided in supplementary materials, with validation protocols. The discrete-time implementation preserves all key properties for protocol integration.

6.2 Sensitivity Analysis

Proposition 6.2 (Parameter Sensitivity)**.** The minimum hesitation time $\Delta T_{\min}$ has first-order sensitivity:

$$\frac{\partial \Delta T_{\min}}{\partial \Phi_{\max}} = -\frac{\Delta\Gamma}{\Phi_{\max}^2}, \quad \frac{\partial \Delta T_{\min}}{\partial \Delta\Gamma} = \frac{1}{\Phi_{\max}}$$

Relative sensitivity to $\Phi_{\max}$ is -1 (inverse proportional), while sensitivity to $\Delta\Gamma$ is $+1$ (linear).

Proof. Direct differentiation of (5). □

Remark 6.2 (Design Implications and Parameter Calibration)**.** This sensitivity analysis informs parameter selection for different application domains:

- For high-stakes decisions requiring longer deliberation: increase $\Delta\Gamma$
- For domains requiring responsiveness with minimal delay: increase $\Phi_{\max}$
- The trade-off is captured mathematically in the accompanying article's calibration methodology
- Protocol parameters map directly to these mathematical variables (`delay` corresponds to $\Delta T_{\min}$)
- Validation protocols test sensitivity bounds

7 Computational Complexity

Proposition 7.1 (Complexity Bounds)**.** The AHP update algorithm has:

- **Time**: $O(1)$ — constant 4 arithmetic operations and 2 comparisons
- **Space**: $O(1)$ — 4 state variables $(I, D, \Gamma_{\min/\max}, \Phi_{\max})$
- **WCET**: Deterministic, data-independent execution path

Proof. Analysis of the reference implementation. □

Remark 7.1 (Practical Deployability Across Domains)**.** The $O(1)$ complexity and deterministic WCET make AHP suitable for hard real-time systems, embedded systems with limited resources, high-frequency applications where computational overhead must be minimized, network protocol implementations requiring predictable performance, and distributed consensus algorithms with strict timing requirements.

8 Conclusion: From Mathematical Guarantees to Governance Infrastructure

This supplement has established that AHP provides:

1. **Temporal Guarantees**: Minimum hesitation $\Delta T_{\min}$ independent of runtime conditions (Theorem 3.1)
2. **Stability**: Lyapunov-stable dynamics with BIBO guarantees (Theorems 4.1 and 4.2)
3. **Robustness**: Intrinsic noise filtering without external debouncing (Theorem 5.1)
4. **Feasibility**: $O(1)$ complexity suitable for embedded systems (Theorem 7.1)
5. **Practicality**: Discrete-time implementations preserve guarantees (Theorem 6.1)

Remark 8.1 (From Mathematical Foundations to Implemented Systems)**.** These mathematical properties translate the philosophical "ethics of hesitation" into verifiable architectural constraints. Crucially, the **architectural impossibility of Zeno behavior** (Theorem 3.2) and the **bounded switching frequency** (Theorem 3.3) provide the deterministic temporal substrate required for:

- *Protocol specifications*—enabling network-level hesitation enforcement
- *Reference implementations*—providing executable code for validation
- *Distributed consensus protocols*—supporting Byzantine fault-tolerant mechanisms for multi-jurisdictional AI governance
- *Radiation-aware computing*—enabling reliable operation of commercial-off-the-shelf hardware in space environments
- *Regulatory compliance*—providing auditable evidence of temporal governance

Without the Non-Zeno guarantee, coordination across institutional boundaries would remain vulnerable to the pathological acceleration that AHP architecturally precludes. These mathematical foundations enable the practical implementations, protocol specifications, governance frameworks, and validation protocols detailed in supplementary materials.

Remark 8.2 (Integrated Research Contribution)**.** The formal guarantees established in this online resource directly support and enable the complete AHP research ecosystem:

- **Theorem 3.1** provides the mathematical foundation for the "Speed-Responsibility Paradox" (accompanying article)
- **Governance cap** $\Phi_{\max}$ operationalizes the "right to hesitation" (accompanying article)
- **Calibration methodology** builds upon the sensitivity analysis in Proposition 7.2
- **Noise rejection properties** enable the case studies in the accompanying article
- **All supplementary materials** derive from these mathematical foundations, creating a coherent, falsifiable research contribution

This mathematical formalization transforms philosophical critique into falsifiable technical infrastructure spanning protocols, implementations, governance frameworks, and validation methodologies.

Related Work

While the AHP framework introduces novel architectural primitives for AI governance, it builds upon established control theory concepts:

- **Hysteresis Control**: Classical treatments of relay systems with hysteresis [3, 4]
- **Schmitt Triggers**: Analysis of threshold-based systems with memory [5]
- **Non-Zeno Systems**: Formal methods for ensuring finite transitions in hybrid systems [6]
- **Temporal Logic in AI**: Formal methods for reasoning about time in autonomous systems [7]

AHP's contribution lies in *re-purposing* these mechanisms for *governance* rather than merely control, embedding ethical constraints directly into system architecture. This approach aligns with Value Sensitive Design principles [8] while providing mathematical rigor often lacking in ethical AI frameworks.

References

[1] Sah, A. (2026). *Algorithmic Hysteresis Primacy (AHP): Temporal Sovereignty in AI Governance*. Working paper available at zmem.org and SSRN.

[2] Sah, A. (2026). *Supplementary Materials: Protocol Specifications, Implementation Examples, Governance Frameworks, and Validation Protocols*. Working paper available at zmem.org and SSRN.

[3] Khalil, H. K. (2002). *Nonlinear Systems* (3rd ed.). Prentice Hall. (Chapter on relay and hysteresis systems)

[4] Åström, K. J., & Murray, R. M. (2008). *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press.

[5] Mead, C. (1989). *Analog VLSI and Neural Systems*. Addison-Wesley. (Chapter on Schmitt triggers and hysteretic circuits)

[6] Alur, R., & Dill, D. L. (1994). A theory of timed automata. *Theoretical Computer Science*, 126(2), 183–235.

[7] Pnueli, A. (1977). The temporal logic of programs. *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, 46–57.

[8] Friedman, B., Kahn Jr, P. H., & Borning, A. (2002). Value sensitive design: Theory and methods. *University of Washington Technical Report*, 2–12.