Title: SSM: A Physics-Inspired Adaptive Runtime for Autonomous Optimization Across Heterogeneous Workloads

Principal Investigator: Robert A. James

Email: rajames440@gmail.com

Phone: (216) 239-6600

Affiliation: TBD

Requested Funding: $300,000

1. Executive Summary

The Steady State Machine (SSM) is a novel adaptive execution architecture for computational systems. It autonomously regulates workloads through coordinated feedback loops, statistical inference, and supervisory mode selection. Designed for real-time, embedded, and mission-critical environments, the SSM ensures deterministic performance through self-regulation without manual tuning. It enables resilient and self-optimizing execution in unpredictable or adversarial conditions.

2. Problem Statement

Modern runtimes fail under workload heterogeneity due to static configurations and brittle heuristic tuning. These limitations yield instability, poor scalability, and elevated variance—unacceptable in defense applications where computational predictability is essential. Current solutions either require manual tuning or rely on simplistic adaptation with limited robustness.

3. Technical Objectives

- Dynamically classify workload behavior (e.g., stable, volatile)
- Coordinate feedback loops to regulate internal runtime parameters
- Ensure shape-invariant performance across waveform families
- Converge to steady-state using conservation laws and thermodynamic models
- Achieve zero-variance behavior validated through extensive empirical trials

4. Technical Approach

The core innovation is the Runtime State Vector (RSV), tracking real-time metrics like heat, entropy, and decay rates. These drive seven feedback loops (L1–L7) and a supervisory Jacquard Mode Selector (L8). The system enforces James Law ($K \equiv 1.0$) and uses bounded hysteretic switching to ensure stable convergence. Adaptive modes are selected based on entropy slope, heat distribution, and statistical inference. The design supports deterministic operation across platforms.

5. Key Innovations

- Multi-loop feedback architecture (L1–L7)
- Supervisory mode selector (L8) with bounded transitions
- K-statistic conservation law (James Law)
- Thermodynamic self-organization and universal frequencies
- Shape-invariant performance across waveform classes

Deterministic Stability and Variance Characterization

Rather than asserting absolute variance elimination, the Steady State Machine (SSM) defines *deterministic stability* in terms of **coefficient of variation (CV%) convergence** over bounded operational horizons. Variance is measured relative to specific runtime observables, including instantaneous heart rate (HR), ΔHR (first derivative), convergence time, and execution jitter. Stability is achieved when the CV% of these observables asymptotically approaches a narrow bound under sustained execution, despite workload heterogeneity or regime transitions.

In practice, SSM enforces a conservation constraint (James Law, $K \equiv 1.0$) that drives feedback-mediated convergence toward a steady-state attractor. This results in *shape-invariant performance*: while absolute values may differ across workload classes, the normalized dispersion (CV%) collapses toward a common bound. This property enables predictable behavior across stable, temporal, transition-heavy, and volatile workloads without manual tuning.

Experimental Validation to Date

A working SSM prototype has been implemented in ANSI C99 and FORTH, comprising a runtime state vector (RSV), seven coordinated feedback loops (L1–L7), and a supervisory Jacquard Mode Selector (L8). Extensive empirical validation has been conducted using synthetic workloads spanning diverse execution profiles, including stable, omni, temporal, transition, and volatile regimes.

Experiments were executed across hundreds of independent runs with randomized seeds and controlled replicates. High-resolution heartbeat instrumentation was used to capture HR, ΔHR, entropy proxies, and mode transitions at runtime. Phase-space density analysis reveals the emergence of bounded attractor regions across all workload classes, with convergence toward a common equilibrium point. Notably, while transient dispersion differs by workload, the steady-state CV% consistently collapses within a narrow band, demonstrating robust self-regulation and deterministic convergence.

These results indicate that SSM does not merely optimize for mean performance, but actively suppresses operational variance through feedback-driven self-organization. Ongoing experiments extend this validation across larger run matrices and longer horizons to further characterize convergence rates, stability margins, and recovery behavior under induced perturbations.

6. Work Plan / Milestones

Phase I (Complete): Development of a working VM prototype in ANSI C99 and FORTH, with core runtime state vector and feedback loop mechanisms implemented and verified.

Phase II (4–6 mo): Expand synthetic workload validation. Conduct performance benchmarking across multiple waveform families to test shape-invariance, feedback stability, and K-statistic convergence.

Phase III

Planned Formal Verification Strategy
While the Steady State Machine (SSM) is fundamentally an empirical, feedback-driven runtime, selected components will be subjected to **targeted formal verification** using Isabelle/HOL. Rather than attempting full-system verification, formal methods will be applied selectively to **critical invariants and safety properties** that govern stability and bounded behavior.

In particular, Isabelle/HOL will be used to model and validate:

- Conservation constraints associated with James Law ($K \equiv 1.0$)
- Boundedness guarantees on feedback loop interactions (L1–L7)
- Hysteretic transition conditions within the supervisory Jacquard Mode Selector (L8)
- Absence of unbounded oscillatory or divergent state trajectories under defined assumptions

These proofs will focus on *local correctness* and *invariant preservation* rather than end-to-end functional completeness. The goal is to formally demonstrate that, given admissible inputs and bounded perturbations, the SSM control structure cannot violate defined stability constraints.

Formal results will be cross-referenced with empirical phase-space observations to ensure alignment between theoretical models and measured runtime behavior. This hybrid approach leverages formal methods to harden guarantees around safety-critical properties while preserving the adaptive, physics-inspired nature of the system.

Phase IV (7–9 mo): Integrate SSM engine into embedded RTOS environments (e.g., FreeRTOS) and JIT virtual machines (e.g., LuaJIT). Validate feedback loop control under resource-constrained conditions.

Phase V (10–12 mo): Evaluate under mission-mimicking workloads, including burst-like, adversarial, and transition-heavy execution scenarios. Record metrics on convergence time, jitter, and computational stability.

Optional Hardware Substrate Evaluation

While the Steady State Machine (SSM) is designed to be platform-agnostic, an FPGA-based implementation is planned as an **optional experimental substrate** to evaluate feedback dynamics under maximally deterministic execution conditions. FPGA deployment is not required for core development, but serves as a controlled environment to isolate runtime feedback behavior from operating system scheduling artifacts, cache interference, and asynchronous software noise.

On FPGA, selected SSM components—such as heartbeat generation, feedback loop interactions, and supervisory mode transitions—may be instantiated as hardware or hybrid hardware/software modules. This enables cycle-accurate measurement of convergence behavior, boundedness, and recovery under induced perturbations. Results obtained on FPGA will be compared directly against software implementations to validate that observed steady-state attractors and CV% convergence properties are intrinsic to the SSM architecture rather than artifacts of a specific execution platform.

This approach positions FPGA as a **verification and characterization tool**, not a deployment dependency, reinforcing the portability and robustness of the SSM design.

7. National Security Relevance

The SSM architecture empowers next-gen embedded platforms, AI agents, and edge devices to operate under unstable or adversarial conditions without performance collapse. Applicable to autonomous vehicles, tactical AI, aerospace platforms, and embedded defense systems, SSM enables runtime adaptability, resource-efficient optimization, and algorithmic resilience critical to mission assurance.

8. PI Background

Robert A. James is the inventor of the SSM. With expertise in runtime system design and physics-based computation, his work bridges adaptive execution theory and control systems. He is currently unaffiliated with an institution but open to partnerships.

Appendix A: Architecture Diagram

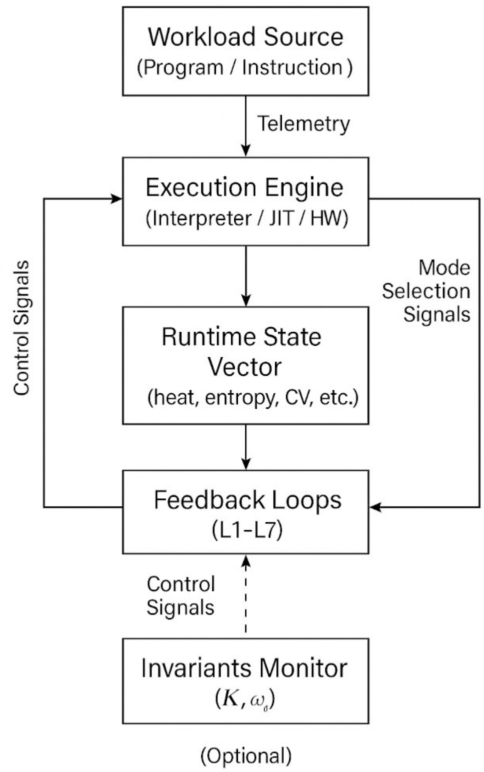

Figure A – Solid State Machine Architecture (SSM)

This figure illustrates the high-level architecture of the SSM as described in the original patent. It includes the runtime state vector, feedback loops L1–L7, the Jacquard mode selector (L8), and the execution loop. This system coordinates adaptive workload regulation using feedback-driven convergence and supervisory control.

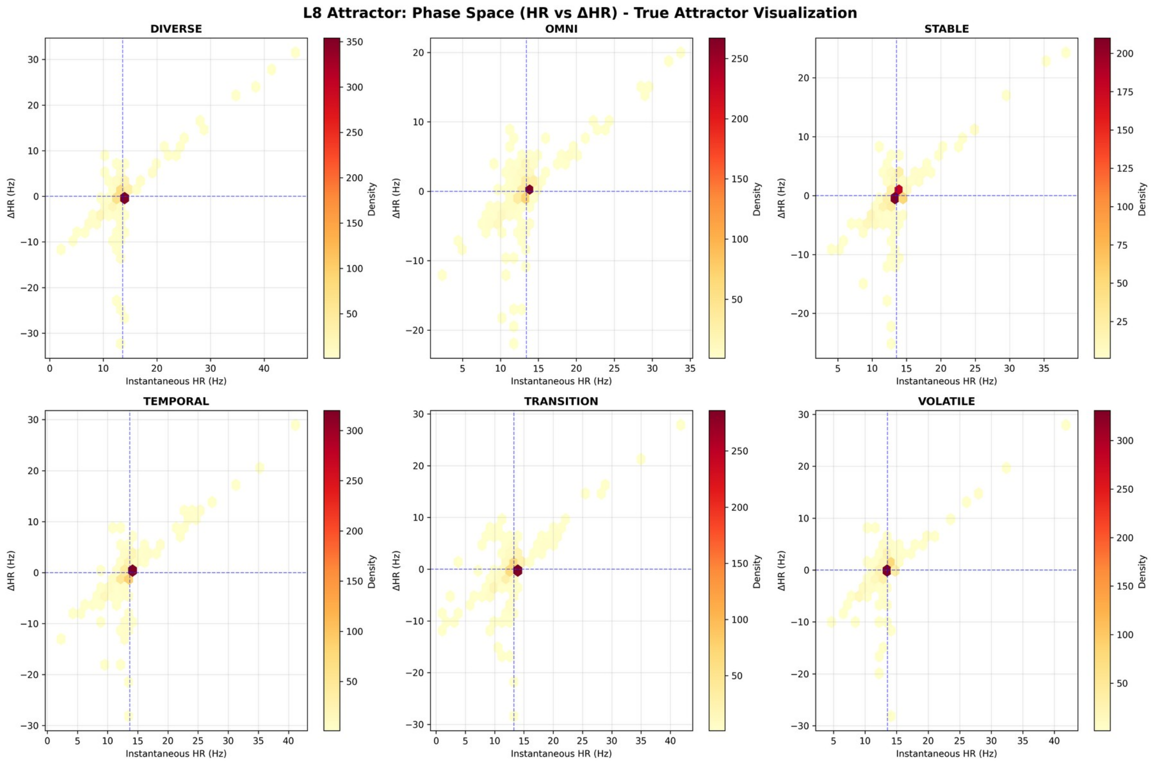

Figure B – Phase Space portrait

Figure B – Phase Space Portrait of the Jacquard Mode Selector (L8)

This diagram presents a **phase space visualization** of the supervisory controller L8's operation within the **Steady State Machine (SSM)** architecture. Each subplot illustrates the relationship between:

- **Instantaneous Heart Rate (HR) —** a proxy for workload activity frequency (x-axis)
- **ΔHR —** the first derivative of HR (rate of change, y-axis)
- **Density coloring —** higher densities of operating states are shown in dark red, indicating attractor regions

The six panels correspond to different classified workload behaviors:

- **DIVERSE**: High dispersion, indicating adaptation across multiple attractor states.
- **OMNI**: A general-purpose regime showing a broad operational envelope with central attractor convergence.
- **STABLE**: Tightly clustered around the origin, reflecting minimal change in workload frequency and strong convergence behavior.

- **TEMPORAL**: Slightly elongated along the horizontal axis, indicating gradual and predictable shifts in workload.
- **TRANSITION**: Broader spread in ΔHR, consistent with regime-shifting or boundary-phase workloads.
- **VOLATILE**: Shows vertical dispersion with large fluctuations in ΔHR, characteristic of bursty or chaotic workloads.

Each plot shows a **central attractor point near (15 Hz, 0 ΔHR)**, reinforcing the **steady-state equilibrium target** of the SSM architecture. The clustering patterns and anisotropies across modes validate the **shape-aware feedback dynamics** of L8 and demonstrate its ability to maintain bounded, adaptive convergence across heterogeneous workloads.