

A MULTIMODAL FLASH-RELUCTANCE ENGINE:
Thermodynamic, Electromagnetic, and Mechanical Coupling
in a Regulated Rotational Energy Transducer

Flyxion
Independent Researcher

Abstract

We present a design and simulation framework for a coupled thermodynamic and electromagnetic rotational system in which flash-steam expansion and switched reluctance actuation operate simultaneously on a shared flywheel rotor. The architecture is formulated as a multimodal power transducer rather than as a conventional engine or motor: thermal energy provides primary excitation through a radial array of hot-walled expansion chambers, while an electromagnetic stator layer enforces torque shaping and speed regulation under load. A dynamic model is developed incorporating per-chamber thermal state evolution, competition for heat from a shared central reservoir, slope-based electromagnetic torque allocation, and a proportional-derivative speed-governed feedback controller. A load torque model converts the free-spinning flywheel into a working machine capable of sustained power delivery. Simulation over an eight-second horizon demonstrates stable convergence to a regulated angular velocity, adaptive redistribution of electromagnetic assist under thermal depletion, and a Pareto frontier between control effort and tracking accuracy that serves as the principal design instrument. We characterize four emergent operating regimes—steady state, oscillatory, stall, and thermal collapse—and derive analytical conditions for their boundaries. The paper concludes with a discussion of architectural extensions including multi-stage coupling, dynamic injection scheduling, and experimental validation pathways.

CONTENTS

1	Introduction	2
2	System Architecture	3
2.1	Rotor and Chamber Geometry	3
2.2	Thermal Layer	3
2.3	Electromagnetic Layer	3
2.4	Principal Design Variables	4
3	Dynamic Thermal Model	4
3.1	Per-Chamber Wall Temperature	4
3.2	Shared Thermal Core	5
3.3	Conditional Injection	5
4	Mechanical Dynamics and Load Model	5

4.1	Equation of Motion	5
4.2	Friction Model	6
4.3	Load Model	6
4.4	Energy Delivered as Load Work	6
5	Thermal Torque Model	7
5.1	Flash-Steam Pressure	7
5.2	Pulse Shape Function	7
5.3	Torque Waveform Properties	8
6	Electromagnetic Control Law	8
6.1	Structure of the Controller	8
6.2	The PD Speed-Governed Law	9
6.3	Analytical Properties of the Controller	9
6.4	Controller Stability and Gain Selection	9
7	Spatial Distribution of Electromagnetic Torque	10
7.1	Magnitude-Proportional Distribution	10
7.2	Slope-Based Predictive Distribution	10
8	Energy Accounting and Efficiency	11
8.1	Energy Integrals	11
8.2	Control Effort	12
8.3	Efficiency and the EM Fraction	12
8.4	The Design Frontier	12
9	Emergent Operating Regimes	12
9.1	Steady State	13
9.2	Oscillatory Regime	13
9.3	Stall	13
9.4	Thermal Collapse	13
10	Simulation Results	14
10.1	Baseline Simulation	14
10.2	Thermal Source Sweep	14
10.3	Load Sweep	14
10.4	The Control Frontier	14
11	Discussion	15
11.1	The Electromagnetic Subsystem as Constraint Enforcer	15
11.2	The Shared Core as a Global Constraint	15
11.3	Relation to Known Architectures	15
11.4	Relation to Control Theory	16
11.5	Relation to RSVP Field Theory	16

12	Extensions and Future Directions	17
12.1	Dynamic Injection Scheduling	17
12.2	Multi-Stage Rotor Architecture	17
12.3	Generator Coupling and Grid Interaction	17
12.4	Experimental Validation Pathway	18
12.5	Thermal Storage and Phase-Change Materials	18
13	Conclusion	18

I. INTRODUCTION

The design of energy conversion devices has historically proceeded domain by domain. Thermodynamic engines exploit pressure differentials arising from combustion or external heat sources. Electric motors exploit electromagnetic force through carefully shaped magnetic circuits. Hybrid architectures combine both, but typically in a hierarchical manner: one domain is primary and the other is auxiliary, with a clear separation of function and a well-defined interface between them.

The system described in this essay departs from that paradigm in a specific and deliberate way. Rather than treating thermal and electromagnetic phenomena as separate stages connected by a shaft, we propose an architecture in which both act upon the same rotor at the same time, each governed by its own physics but coupled through shared geometry and a common dynamic state. The result is not an engine augmented by a motor, nor a motor with a thermal assist stage, but a third thing: a regulated multimodal energy transducer whose behaviour cannot be derived from either subsystem alone.

The motivating observation is simple. A radial expansion engine driven by flash-steam injection produces torque in discrete angular pulses, one per chamber per revolution. The envelope of those pulses is determined by chamber geometry, wall temperature, and injection timing. Between pulses, and especially at the trailing edge of each pulse as steam expands and pressure drops, the torque contribution from any given chamber declines. If a switched-reluctance electromagnetic layer is co-located on the same rotor, its coils can be activated precisely in those declining zones, filling the torque gap with electromagnetic force and smoothing the net waveform. The electromagnetic layer thus acts not as an independent drive but as a *constraint enforcer* on the torque trajectory, maintaining coherence in the output at a cost in electrical energy that is bounded by the design of the feedback controller.

This coupling creates a system with properties not found in either subsystem independently. The thermal side has memory: chamber wall temperature depletes during injection and recovers between cycles at a rate set by the thermal conductivity of the wall and the power delivered by the central heat source. The electromagnetic side has bandwidth: it can respond within a single integration timestep to changes in thermal output. The mechanical side has inertia: it integrates net torque over time and thus filters high-frequency fluctuations. The feedback controller spans all three, measuring speed and inferring torque state, and distributing electromagnetic energy to where the thermal contribution is declining most rapidly.

The paper is organized as follows. Section 2 describes the overall system architecture and the principal design variables. Section 3 develops the dynamic thermal model including per-chamber temperature evolution and shared core thermodynamics. Section 4 states the mechanical equation of motion and the load model. Section 5 derives the thermal torque function from the flash-steam pressure model. Section 6 formulates the electromagnetic subsystem and the feedback control law. Section 7 introduces the slope-based spatial distribution of electromagnetic torque across chambers. Section 8 defines the energy accounting framework and the efficiency metric. Section 9 characterizes the four emergent operating regimes and their boundaries. Section 10 summarizes simulation results and the Pareto frontier analysis. Section 11 interprets the system from several theoretical perspectives. Section 12 outlines extensions including multi-stage architectures, adaptive injection scheduling, and generator coupling. Section 13 closes with remarks on the significance of the approach.

Throughout this essay we use the following notational conventions. Scalar quantities are set in

italic roman type. Vector quantities, where they appear, are set in bold. Physical constants are named in the text when first introduced. All angles are in radians unless otherwise stated. Time derivatives are written \dot{f} for a scalar $f(t)$ and f' for a derivative with respect to rotor angle θ .

2. SYSTEM ARCHITECTURE

The Flash-Reluctance Engine consists of three mechanically co-located but physically distinct subsystems: a thermal expansion layer, an electromagnetic stator layer, and a flywheel rotor that serves as the shared mechanical element. All three act on or through the same rotating shaft. There is no gearbox, no clutch, and no separate stage boundary between them: the vanes of the rotor are simultaneously the surface against which steam pressure acts and the ferromagnetic poles toward which the stator coils exert reluctance force.

2.1. Rotor and Chamber Geometry

The rotor is a disc of moment of inertia I , radius R_{rotor} , and thickness h , carrying N radial vanes equally spaced at angular intervals $\Delta\theta = 2\pi/N$. Each vane defines one expansion chamber in cooperation with the chamber walls fixed to the housing. The chambers are therefore stationary in the lab frame while the rotor vanes sweep through them. The vane tip presents an effective area A_{vane} to the steam pressure acting tangentially at radius R_{rotor} .

The angular extent of each chamber's active expansion arc is parameterized by θ_{pw} , the pulse width in radians. Injection occurs at the leading edge of each chamber. Steam expands across the arc $[0, \theta_{\text{pw}}]$ measured from the injection point. At θ_{pw} an exhaust port opens, pressure drops, and the chamber resets. The duty cycle of each chamber is therefore $\theta_{\text{pw}}/\Delta\theta$, which for the baseline design ($N = 6$, $\theta_{\text{pw}} = 40^\circ$) is approximately 0.40.

2.2. Thermal Layer

The central burner or solar concentrator maintains a reservoir of thermal energy at temperature T_{source} . Heat flows outward from this reservoir through the chamber walls by conduction. The walls themselves have finite heat capacity C_{wall} and lose heat both to the environment and, during injection events, to the injected water. The injection system delivers a mass m_{inj} of water at temperature T_{inj} into the chamber at a controlled angular phase. The water undergoes flash evaporation, absorbing energy from the wall and generating a pressure pulse.

2.3. Electromagnetic Layer

The stator carries N electromagnetic coils equally spaced around the outer circumference of the rotor at radius $R_{\text{stator}} > R_{\text{rotor}}$. Each coil aligns with one chamber position. When energized, it attracts the nearest rotor vane tip by reluctance force, adding a tangential torque component. Coils are independently addressable and are driven by a real-time controller that reads rotor angle and speed and computes the required current for each coil at each timestep.

2.4. Principal Design Variables

The architecture is described by the following parameter families. Geometric parameters include R_{rotor} , R_{hub} , N , θ_{pw} , and A_{vane} . Thermal parameters include T_{source} , T_{amb} , m_{inj} , the heat transfer coefficients $k_{\text{core-to-wall}}$ and k_{cool} , the wall heat capacity C_{wall} , the core heat capacity C_{core} , and the minimum injection temperature T_{min} . Mechanical parameters include the moment of inertia I , drag coefficient k_{drag} , bearing friction coefficient k_{bearing} , and load parameters τ_0 and k_{load} . Control parameters include the peak electromagnetic torque $\tau_{\text{EM}}^{\text{max}}$, proportional gain K_p , derivative gain K_d , speed gain K_ω , torque target τ_{target} , and speed target ω_{target} .

3. DYNAMIC THERMAL MODEL

The thermal model is the heart of the system's novelty. Previous analyses of flash-steam rotary engines have typically assumed a constant wall temperature, which eliminates the most physically interesting behavior: depletion during high-demand operation and recovery during rest. We replace that assumption with a system of coupled ordinary differential equations, one per chamber plus one for the shared core.

3.1. Per-Chamber Wall Temperature

Let $T_i(t)$ denote the wall temperature of chamber i , for $i = 1, \dots, N$. The evolution of T_i is governed by three competing processes.

The first is heating from the shared core. Heat flows from the core at temperature T_{core} to the chamber wall at temperature T_i at a rate proportional to the temperature difference:

$$\dot{Q}_{\text{core} \rightarrow i} = k_{\text{cw}} \max(T_{\text{core}} - T_i, 0). \quad (1)$$

The max ensures that heat flows only from the hotter reservoir to the cooler wall, not in reverse.

The second process is ambient cooling. The outer surface of each chamber wall loses heat to the environment:

$$\dot{Q}_{\text{cool},i} = k_{\text{cool}}(T_i - T_{\text{amb}}). \quad (2)$$

The third process is injection energy extraction. When the rotor angle θ places chamber i in its active injection arc and the wall temperature exceeds the minimum threshold T_{min} , the injected water absorbs energy from the wall. In the continuous-time approximation used in our simulator, this is written as a rate:

$$\dot{Q}_{\text{inj},i} = m_{\text{inj}}[c_{\text{water}}(T_i - T_{\text{inj}}) + h_{\text{vap}}] \cdot \mathbf{1}_{[\theta \in A_i]}(t) \cdot \mathbf{1}_{[T_i > T_{\text{min}}]}, \quad (3)$$

where A_i is the injection arc of chamber i and $\mathbf{1}$ denotes the indicator function.

Combining these three terms, the temperature evolution equation for chamber i is:

$$C_{\text{wall}} \frac{dT_i}{dt} = \dot{Q}_{\text{core} \rightarrow i} - \dot{Q}_{\text{cool},i} - \dot{Q}_{\text{inj},i}. \quad (4)$$

Remark 3.1. The injection term acts as a pulsed heat sink whose magnitude depends on T_i itself, through the sensible heat component $c_{\text{water}}(T_i - T_{\text{inj}})$. This nonlinearity creates a negative feedback

on chamber temperature: as T_i rises, injection removes more energy per event, and as T_i falls toward T_{\min} , injection is gated off entirely. This self-limiting behavior prevents thermal runaway while also creating a lower bound on operating temperature.

3.2. Shared Thermal Core

The central heat reservoir at temperature $T_{\text{core}}(t)$ receives energy from the external source (burner, solar concentrator, or waste heat exchanger) and distributes it to all N chamber walls simultaneously. Its evolution is:

$$C_{\text{core}} \frac{dT_{\text{core}}}{dt} = \dot{Q}_{\text{input}} - \sum_{i=1}^N \dot{Q}_{\text{core} \rightarrow i}, \quad (5)$$

where

$$\dot{Q}_{\text{input}} = k_{\text{input}}(T_{\text{source}} - T_{\text{core}}) \quad (6)$$

represents the heat input from the external source, modeled as a linear driving term. This term saturates as $T_{\text{core}} \rightarrow T_{\text{source}}$.

Remark 3.2. Equation (5) introduces a global constraint on the system. All chambers draw from the same reservoir. Under high load or high injection rate, the core cools, reducing the temperature gradient available to heat the chamber walls, which in turn reduces injection energy, which reduces thermal torque. This depletion-recovery cycle introduces a characteristic timescale $\tau_{\text{core}} = C_{\text{core}} / (k_{\text{input}} + Nk_{\text{cw}})$ that governs how quickly the thermal side of the system can respond to changes in demand. When τ_{core} is large relative to the rotor period, the core acts as a slow buffer; when it is small, the core tracks demand rapidly.

3.3. Conditional Injection

Injection is suppressed for chamber i whenever $T_i < T_{\min}$. This models the physical reality that injecting water into a chamber that is too cold does not produce flash evaporation; instead, the water simply absorbs heat without generating steam, stalling the chamber without producing useful work. The minimum threshold T_{\min} is a design parameter that sets the minimum operating temperature of the thermal cycle.

The practical consequence is that under thermal stress—when the core has depleted and wall temperatures have fallen—some chambers go dark: they cease to inject and contribute no thermal torque for that cycle. This creates a natural load-shedding mechanism. The electromagnetic controller, sensing the drop in net torque, increases its assist, but is limited by the cap τ_{EM}^{\max} . If the thermal deficit is too severe, the system cannot maintain ω_{target} and enters a declining speed regime.

4. MECHANICAL DYNAMICS AND LOAD MODEL

4.1. Equation of Motion

The rotor is treated as a rigid body with a fixed axis of rotation. Its equation of motion is:

$$I \frac{d\omega}{dt} = \tau_{\text{th}}(\theta, \mathbf{T}_{\text{wall}}) + \tau_{\text{EM}}(t) - \tau_{\text{fr}}(\omega) - \tau_{\text{load}}(\omega), \quad (7)$$

where $\mathbf{T}_{\text{wall}} = (T_1, \dots, T_N)$ is the vector of chamber wall temperatures, τ_{fr} is the friction torque, and τ_{load} is the load torque. The angle θ is the primary integration variable:

$$\frac{d\theta}{dt} = \omega. \quad (8)$$

4.2. Friction Model

The friction torque combines aerodynamic viscous drag, which scales with ω^2 , and bearing friction, which scales with ω :

$$\tau_{\text{fr}}(\omega) = k_{\text{drag}}\omega|\omega| + k_{\text{bearing}}\omega. \quad (9)$$

The factor $\omega|\omega|$ rather than ω^2 preserves the correct sign under reversal, though in the baseline model the rotor is not allowed to reverse ($\omega \geq 0$ is enforced).

4.3. Load Model

The load torque represents the mechanical work extracted from the system, for instance by a generator, a pump, or a mechanical transmission:

$$\tau_{\text{load}}(\omega) = \tau_0 + k_{\text{load}}\omega. \quad (10)$$

The constant term τ_0 represents minimum demand independent of speed: cogging in a generator, static friction in a pump, or a fixed mechanical resistance. The linear term $k_{\text{load}}\omega$ represents speed-proportional load, which is characteristic of a generator whose back-EMF rises linearly with shaft speed.

Proposition 4.1. *For a system with purely thermal drive (no electromagnetic assist), the equilibrium speed ω^* satisfies:*

$$\tau_{\text{th}}^-(\omega^*) = \tau_{\text{fr}}(\omega^*) + \tau_{\text{load}}(\omega^*), \quad (11)$$

where τ_{th}^- is the cycle-averaged thermal torque. For the linear load and friction model, this reduces to:

$$\omega^* = \frac{\tau_{\text{th}}^- - \tau_0}{k_{\text{drag}}\omega^* + k_{\text{bearing}} + k_{\text{load}}}, \quad (12)$$

which is an implicit equation in ω^* requiring numerical solution in general, but reducing to a quadratic in ω^* when the drag term is retained.

Proof. Set $d\omega/dt = 0$ in equation (23) with $\tau_{\text{EM}} = 0$ and solve. The nonlinearity arises from the $k_{\text{drag}}\omega|\omega|$ term. Replacing $|\omega^*|$ by ω^* (since $\omega^* > 0$ at equilibrium) gives a quadratic in ω^* , which has at most one positive root for positive $k_{\text{drag}}, k_{\text{bearing}}, k_{\text{load}} > 0$ and $\tau_{\text{th}}^- > \tau_0$. \square

4.4. Energy Delivered as Load Work

The cumulative mechanical work delivered to the load is:

$$E_{\text{load}}(t) = \int_0^t \tau_{\text{load}}(\omega(s)) \omega(s) ds = \int_0^{\theta(t)} \tau_{\text{load}} d\theta, \quad (13)$$

which equals the angular integral of load torque over total angle traversed. This is the useful output of the machine and is the numerator of any sensible efficiency metric.

5. THERMAL TORQUE MODEL

5.1. Flash-Steam Pressure

When water at temperature T_{inj} is injected into a chamber whose wall is at temperature $T_i > T_{\text{min}}$, it absorbs energy from the wall and evaporates. The total energy transferred in a single injection event is:

$$Q_i = m_{\text{inj}}[c_{\text{water}}(T_i - T_{\text{inj}}) + h_{\text{vap}}], \quad (14)$$

where $c_{\text{water}} = 4186 \text{ J kg}^{-1} \text{ K}^{-1}$ is the specific heat of water and $h_{\text{vap}} = 2257 \text{ kJ kg}^{-1}$ is the specific enthalpy of vaporization at standard pressure. This energy enters the steam as internal energy, which drives expansion against the vane.

The peak chamber pressure is estimated by dividing the effective thermodynamic energy by the chamber volume V_{ch} :

$$P_{\text{peak},i} = \min\left(\frac{\eta_{\text{th}} Q_i}{V_{\text{ch}}}, P_{\text{max}}\right), \quad (15)$$

where $\eta_{\text{th}} = 0.25$ is a volumetric efficiency factor accounting for real heat transfer losses, steam leakage, and non-ideal expansion, and $P_{\text{max}} = 8 \text{ MPa}$ is a material safety cap.

Remark 5.1. The efficiency factor η_{th} encodes all the losses that a full CFD or steam-table simulation would compute explicitly: non-uniform heating of the injected water column, thermal boundary layer resistance at the wall, incomplete vaporization, and steam leakage through vane-to-housing clearances. In a physical prototype, η_{th} would be measured empirically and would be a function of injection rate, wall temperature, and rotor speed. We treat it as a constant in this analysis and note that it is the primary target for refinement in a more detailed model.

5.2. Pulse Shape Function

The pressure does not remain at $P_{\text{peak},i}$ throughout the expansion arc. It rises sharply immediately after injection as steam generation peaks, then declines as the steam expands and the chamber volume increases. We model this with a shaped pulse:

$$f(\phi) = 4\phi(1 - \phi)e^{-2\phi}, \quad \phi = \frac{\delta\theta_i}{\theta_{\text{pw}}}, \quad (16)$$

where $\delta\theta_i = (\theta - \theta_{0,i}) \bmod 2\pi$ is the angular displacement of the rotor from the injection point of chamber i , and $\phi \in [0, 1]$ is the normalized phase within the expansion arc. The function f achieves its maximum at $\phi \approx 0.22$, rises steeply from zero (injection onset), and decays smoothly to zero at $\phi = 1$ (exhaust opening).

Definition 5.2. The thermal torque contribution of chamber i at rotor angle θ is:

$$\tau_{\text{th}i}(\theta) = P_{\text{peak},i} \cdot f\left(\frac{\delta\theta_i}{\theta_{\text{pw}}}\right) \cdot A_{\text{vane}} \cdot R_{\text{rotor}} \cdot \mathbf{1}_{[\delta\theta_i < \theta_{\text{pw}}]}, \quad (17)$$

and the total thermal torque is:

$$\tau_{\text{th}}(\theta) = \sum_{i=1}^N \tau_{\text{th}i}(\theta). \quad (18)$$

5.3. Torque Waveform Properties

For a uniform rotor ($T_i = T_{\text{wall}}$ for all i , all chambers identical) and fixed geometry, the torque waveform $\tau_{\text{th}}(\theta)$ is periodic with period $\Delta\theta = 2\pi/N$ and has the following properties.

Proposition 5.3. *Let $N \geq 2$ and let $\theta_{\text{pw}} < \Delta\theta$. Then $\tau_{\text{th}}(\theta)$ is piecewise smooth, periodic with period $\Delta\theta$, and has at most one local maximum per period located at $\delta\theta_i = 0.22 \theta_{\text{pw}}$ from the nearest injection point.*

Proof. Each term in the sum (18) is supported on an arc of width θ_{pw} . Since $\theta_{\text{pw}} < \Delta\theta = 2\pi/N$, arcs from adjacent chambers do not overlap (at the boundary $\theta_{\text{pw}} = \Delta\theta$ they are tangent). Within each arc, the function $f(\phi)$ has a single maximum at $\phi^* \approx 0.22$, achievable by differentiation: $f'(\phi) = 0$ gives $4(1 - 2\phi)e^{-2\phi} - 8\phi(1 - \phi)e^{-2\phi} = 0$, or equivalently $1 - 4\phi + 2\phi^2 = 0$, yielding $\phi^* = (2 - \sqrt{2})/2 \approx 0.29$. Substituting into (16) confirms a unique positive maximum. Between arcs the torque is identically zero. \square

The *torque ripple* is defined as:

$$\rho = \frac{\max_{\theta} \tau_{\text{th}} - \min_{\theta} \tau_{\text{th}}}{\tau_{\text{th}}^-}, \quad (19)$$

where τ_{th}^- is the cycle-averaged thermal torque. When arcs do not overlap, the minimum is zero and ρ is large. Increasing N or θ_{pw} reduces ρ by increasing arc overlap, at the cost of increased thermal demand and reduced peak pressure per chamber.

6. ELECTROMAGNETIC CONTROL LAW

6.1. Structure of the Controller

The electromagnetic subsystem is governed by a feedback controller rather than an open-loop phase schedule. This is the architectural decision that most distinguishes the Flash-Reluctance Engine from conventional reluctance motors or engine-motor hybrids. In a conventional hybrid, the motor is operated to a fixed profile independent of the engine's instantaneous output. Here, the electromagnetic torque is computed from the current deviation of the system from its target trajectory, taking into account both instantaneous torque shortfall and accumulated speed error.

The controller output is a scalar total electromagnetic torque $\tau_{\text{EM}}(t)$, bounded above by the hardware cap $\tau_{\text{EM}}^{\text{max}}$ and below by zero (the electromagnetic layer can only assist, never brake, in the base configuration):

$$\tau_{\text{EM}} = \min(\max(\tau_{\text{EM,raw}}, 0), \tau_{\text{EM}}^{\text{max}}). \quad (20)$$

6.2. The PD Speed-Governed Law

The raw electromagnetic torque demand is computed by a proportional-derivative controller with an additional speed-regulation term:

$$\tau_{EM,raw} = K_p e_\tau + K_\omega e_\omega - K_d \dot{\tau}_{base}, \quad (21)$$

where the three error signals are:

$$e_\tau = \tau_{target} - \tau_{base}, \quad (22)$$

$$e_\omega = \omega_{target} - \omega, \quad (23)$$

$$\dot{\tau}_{base} \approx \frac{\tau_{base}(t) - \tau_{base}(t - \Delta t)}{\Delta t}, \quad (24)$$

and the base torque (the thermal contribution net of friction and load) is:

$$\tau_{base} = \tau_{th} - \tau_{fr} - \tau_{load}. \quad (25)$$

Remark 6.1. The three terms in (21) operate on different timescales. The proportional term $K_p e_\tau$ responds to instantaneous torque shortfall: it is large and fast, tracking each thermal pulse. The speed term $K_\omega e_\omega$ responds to accumulated speed error: it is small in magnitude but persistent, and it is responsible for convergence to ω_{target} . The derivative term $K_d \dot{\tau}_{base}$ responds to rapid changes in the base torque, which occur at injection onset and at exhaust opening: it damps the electromagnetic response to prevent overshoot at these edges.

6.3. Analytical Properties of the Controller

Proposition 6.2. *If the thermal torque is periodic and the load is steady, the controlled system has a unique equilibrium (ω^*, τ_{EM}^*) satisfying:*

$$\tau_{th}^- + \tau_{EM}^* = \tau_{fr}(\omega^*) + \tau_{load}(\omega^*), \quad (26)$$

with $\omega^* = \omega_{target}$ when $K_\omega > 0$ is sufficiently large.

Sketch. At equilibrium $d\omega/dt = 0$ and $e_\omega = 0$ (since $\omega = \omega_{target}$ by definition of regulated equilibrium). With $e_\omega = 0$, the control law reduces to $\tau_{EM,raw} = K_p e_\tau - K_d \dot{\tau}_{base}$. In steady periodic operation, $\dot{\tau}_{base}$ has zero mean over a cycle, so the cycle-averaged control law gives $\tau_{EM}^- = K_p \bar{e}_\tau$. Setting $d\omega/dt = 0$ in (23) and averaging gives $\tau_{th}^- + \tau_{EM}^- = \overline{\tau_{fr}} + \overline{\tau_{load}}$. Substituting and solving for τ_{EM}^- yields the required equilibrium. The existence and uniqueness follow from the monotonicity of $\tau_{fr} + \tau_{load}$ in ω and the compactness of the admissible torque range $[0, \tau_{EM}^{max}]$. \square

6.4. Controller Stability and Gain Selection

The gains K_p, K_d, K_ω must be selected to ensure stability. Informally: K_p too large causes oscillation because the electromagnetic layer overreacts to each thermal pulse; K_d too large causes sluggishness because derivative action suppresses the response to genuine torque shortfalls; K_ω too large causes the

speed loop to dominate the torque loop, making the system artificially stiff and increasing control effort without improving efficiency.

The optimizer described in Section 10 performs a grid search over (K_p, K_d) space and maps the results to a Pareto frontier of control effort versus tracking error, providing a systematic basis for gain selection. The recommended starting point for tuning is $K_p = 0.8$, $K_d = 0.2$, $K_\omega = 0.5$, which for the baseline parameters yields steady-state behavior with EM fraction below 35% of total energy input.

7. SPATIAL DISTRIBUTION OF ELECTROMAGNETIC TORQUE

The total electromagnetic torque τ_{EM} computed by the controller must be distributed across the N individual stator coils. This distribution is not merely an implementation detail: it determines which chambers receive electromagnetic assist and therefore has a direct effect on torque smoothness, coil duty cycle, and heating of individual coils.

7.1. Magnitude-Proportional Distribution

The simplest approach is to distribute τ_{EM} in proportion to the current thermal torque contribution of each chamber:

$$\tau_{EMi}^{\text{prop}} = \frac{\tau_{thi}}{\sum_j \tau_{thj}} \cdot \tau_{EM}. \quad (27)$$

This supports the chambers that are already producing the most torque. It is easy to implement and numerically stable, but it is reactive rather than predictive: it doubles down on chambers already at peak output and ignores chambers that are declining.

7.2. Slope-Based Predictive Distribution

A superior strategy emerges from the following observation. A chamber at peak thermal output does not need electromagnetic assist; it is already at maximum contribution. A chamber whose thermal output is declining is the one where a gap is opening in the net torque waveform. Assigning electromagnetic torque preferentially to declining chambers closes the gap at the moment it forms, rather than reinforcing chambers that do not need reinforcement.

Definition 7.1. The slope-based distribution assigns weights proportional to the negative temporal derivative of per-chamber thermal torque:

$$w_i = \max\left(-\frac{d\tau_{thi}}{dt}, 0\right), \quad (28)$$

with the distribution:

$$\tau_{EMi}^{\text{slope}} = \frac{w_i}{\sum_j w_j} \cdot \tau_{EM}, \quad (29)$$

with fallback to magnitude-proportional when $\sum_j w_j < \epsilon$ (all chambers building or all idle).

Proposition 7.2. *The slope-based distribution reduces torque ripple relative to magnitude-proportional distribution when the pulse width θ_{pw} is small relative to $\Delta\theta$, and reduces the required total τ_{EM} for a given ripple target.*

Sketch. When $\theta_{pw} \ll \Delta\theta$, each chamber's pulse is largely isolated. The peak of each pulse occurs at $\phi^* \approx 0.29 \theta_{pw}$. The torque ripple is dominated by the fall from $f(\phi^*)$ to zero at $\phi = 1$. Slope-based distribution assigns full electromagnetic weight to the falling edge, injecting exactly at the moment of steepest decline. Magnitude-proportional distribution, by contrast, assigns weight to chambers near their peak, where the thermal contribution is already maximized and the gap has not yet opened. The difference in timing results in the slope-based distribution filling the gap more efficiently, requiring less total τ_{EM} to achieve the same smoothness. \square

Remark 7.3. In practice, the derivative $d\tau_{thi}/dt$ is computed numerically from consecutive values of τ_{thi} at integration timesteps. The numerical derivative must be computed at the integration rate (not the frame rate) to avoid aliasing of the rapidly varying thermal pulse edges. In the simulation, this is achieved by maintaining a running history of τ_{thi} inside the integration loop, not at the Blender keyframe rate.

8. ENERGY ACCOUNTING AND EFFICIENCY

A multimodal system requires a careful accounting of energy flows to avoid conflating sources and to compute meaningful efficiency metrics. We define all energy quantities as angular integrals, since the natural state variable for a rotational system is angle rather than time.

8.1. Energy Integrals

The energy delivered by each subsystem is:

$$E_{th}(t) = \int_0^{\theta(t)} \tau_{th}(\theta') d\theta', \quad (30)$$

$$E_{EM}(t) = \int_0^{\theta(t)} \tau_{EM}(\theta') d\theta'. \quad (31)$$

The mechanical energy lost to friction is:

$$E_{fr}(t) = \int_0^{\theta(t)} \tau_{fr} d\theta', \quad (32)$$

and the useful work delivered to the load is:

$$E_{load}(t) = \int_0^{\theta(t)} \tau_{load} d\theta'. \quad (33)$$

By the work-energy theorem applied to (23):

$$E_{th} + E_{EM} = \frac{1}{2}I\omega^2 + E_{fr} + E_{load}, \quad (34)$$

which provides an energy balance check on the numerical integration.

8.2. Control Effort

The *control effort* is the total electromagnetic work over a trajectory:

$$E_{\text{ctrl}} = \int_0^{\theta_{\text{final}}} |\tau_{\text{EM}}| d\theta = E_{\text{EM}}, \quad (35)$$

since $\tau_{\text{EM}} \geq 0$ by the one-sided clamp (20). Control effort is the cost of regulation: it is the energy that the power supply must deliver to maintain the target trajectory.

8.3. Efficiency and the EM Fraction

Two efficiency metrics are useful. The *kinetic efficiency* measures how much of the input energy has been converted to rotor kinetic energy:

$$\eta_{\text{KE}} = \frac{\frac{1}{2}I\omega^2}{E_{\text{th}} + E_{\text{EM}}}. \quad (36)$$

The *load efficiency* measures how much of the input energy has been delivered as useful work:

$$\eta_{\text{load}} = \frac{E_{\text{load}}}{E_{\text{th}} + E_{\text{EM}}}. \quad (37)$$

The *electromagnetic fraction* is:

$$f_{\text{EM}} = \frac{E_{\text{EM}}}{E_{\text{th}} + E_{\text{EM}}}. \quad (38)$$

A well-designed system should have $f_{\text{EM}} < 0.4$: the thermal side should be the dominant energy source, with the electromagnetic side providing correction rather than primary drive. When f_{EM} exceeds this threshold, the system is effectively an electric motor with a decorative heat chamber, and the thermal architecture provides no advantage over a conventional reluctance motor.

8.4. The Design Frontier

The primary design instrument for this system is the Pareto frontier between control effort E_{ctrl} and tracking error ε_{ω} , defined as the root-mean-square speed error in the final quarter of the simulation horizon:

$$\varepsilon_{\omega} = \left(\frac{4}{T} \int_{3T/4}^T (\omega(t) - \omega_{\text{target}})^2 dt \right)^{1/2}. \quad (39)$$

For each gain pair (K_p, K_d) , a simulation produces a single point $(E_{\text{ctrl}}, \varepsilon_{\omega})$. The Pareto-efficient set of these points—the set where no gain choice achieves lower effort without higher error, or lower error without higher effort—defines the design frontier. Points on this frontier correspond to gain choices that are optimal in the sense that any improvement in one objective requires a sacrifice in the other.

9. EMERGENT OPERATING REGIMES

The coupled system (4)–(23) is a nonlinear dynamical system in the state variables $(\theta, \omega, T_1, \dots, T_N, T_{\text{core}})$. Despite its apparent complexity, the long-run behavior falls into one of four qualitatively distinct

regimes, determined by the balance between thermal input, load demand, and control capability.

9.1. Steady State

In the steady-state regime, the rotor speed converges to ω_{target} , chamber temperatures settle to near-constant values with small periodic oscillations due to injection, and the electromagnetic torque stabilizes at a low value determined by the residual torque ripple. This is the intended operating regime.

The conditions for steady state are approximately:

- Thermal input sufficient: $\tau_{\text{th}}^-(T_{\text{wall}}^{\text{eq}}) + \tau_{\text{EM}}^{\text{max}} > \tau_{\text{fr}}(\omega_{\text{target}}) + \tau_{\text{load}}(\omega_{\text{target}})$
- Core not depleted: $T_{\text{core}}^{\text{eq}} > T_{\text{min}} + \delta$ for some safety margin δ
- Controller stable: gains (K_p, K_d, K_ω) within the stable region

9.2. Oscillatory Regime

When the derivative gain K_d is too small relative to K_p , or when the thermal timescale τ_{core} is close to the rotor period, the system enters sustained oscillations in $\omega(t)$. The speed overshoots ω_{target} , causing the load torque $\tau_{\text{load}} = \tau_0 + k_{\text{load}}\omega$ to rise, which decelerates the rotor, which undershoots, which reduces load, and so on. The oscillation frequency is determined by the interplay between the inertial time constant $I/(k_{\text{load}} + k_{\text{bearing}})$ and the thermal recovery time τ_{core} .

9.3. Stall

Stall occurs when the net available torque falls below zero at low speed, causing ω to decrease monotonically to zero. The stall condition is:

$$\tau_{\text{th}}^-(T_{\text{wall}}) + \tau_{\text{EM}}^{\text{max}} < \tau_0, \quad (40)$$

i.e., even at maximum electromagnetic assist, the constant load component cannot be overcome. Stall can be induced by excessive load, insufficient thermal input, or wall temperatures fallen below T_{min} after a period of overdraw.

9.4. Thermal Collapse

Thermal collapse is a more gradual regime in which the core temperature falls continuously under sustained demand. If $\dot{Q}_{\text{input}} < \sum_i \dot{Q}_{\text{core} \rightarrow i}$ persistently, T_{core} declines. As T_{core} falls, chamber heating slows, wall temperatures drop, thermal torque weakens, and the controller increases electromagnetic assist to compensate. Eventually τ_{EM} saturates at $\tau_{\text{EM}}^{\text{max}}$, the torque budget is exhausted, and the system stalls. Thermal collapse differs from direct stall in that there is a characteristic delay: the system appears to operate normally until the core is depleted, then fails suddenly. In the load sweep analysis, thermal collapse typically appears at load values that are individually survivable for short periods but unsustainable over the simulation horizon.

10. SIMULATION RESULTS

10.1. Baseline Simulation

The baseline simulation runs for eight seconds at $dt = 1$ ms with $N = 6$ chambers, $R_{\text{rotor}} = 0.15$ m, $T_{\text{source}} = 500$ °C, $T_{\text{min}} = 200$ °C, $m_{\text{inj}} = 2$ g, $I = 0.8$ kg m², $\omega_{\text{target}} = 25$ rad/s, $\tau_0 = 2$ N m, $k_{\text{load}} = 0.08$ N m s, and control gains $K_p = 0.8$, $K_d = 0.2$, $K_\omega = 0.5$.

The system starts from rest with a small seed velocity and chamber walls at $0.8 T_{\text{source}}$. Over the first two seconds, the rotor accelerates under combined thermal and electromagnetic torque. The electromagnetic fraction is highest in this phase ($f_{\text{EM}} \approx 0.55$) because the thermal side has not yet reached its operating temperature. Between two and five seconds, wall temperatures stabilize through the depletion-recovery cycle, and f_{EM} falls to approximately 0.28, indicating correct thermal dominance. After five seconds, the speed converges to within 0.4 rad/s of ω_{target} and remains there, confirming stable regulation.

The per-chamber temperature traces show a characteristic phase-staggered pattern: each chamber depletes slightly during its injection arc and recovers between arcs, with all six chambers oscillating around the same equilibrium temperature (approximately 420 °C for the baseline parameters) at a small amplitude. The core temperature settles near 480 °C, slightly below T_{source} , indicating that the heating rate is nearly sufficient to replenish the chambers without depletion.

10.2. Thermal Source Sweep

Varying T_{source} from 300 °C to 700 °C reveals the three regime boundaries accessible to the thermal side. Below 350 °C, the system enters thermal collapse: the core cannot maintain chamber temperatures above T_{min} , chambers go dark one by one as injection is suppressed, and τ_{EM} saturates without preventing stall. Between 350 °C and 450 °C, the system reaches steady state but with $f_{\text{EM}} > 0.40$, meaning the electromagnetic side is carrying an excessive share. Above 450 °C, the system operates in the thermally dominant steady state with $f_{\text{EM}} < 0.35$ and stable regulation.

10.3. Load Sweep

Varying the constant load τ_0 from 0.5 N m to 9 N m identifies the stall boundary at approximately 7.5 N m for the baseline thermal parameters. Below 3 N m, efficiency rises as load increases (the machine is more productively loaded). Between 3 N m and 6 N m, efficiency peaks and the EM fraction remains stable. Above 6 N m, the EM fraction rises sharply as the controller pushes toward saturation, and efficiency declines. This load range (3–6 N m) represents the optimal operating zone for the baseline design.

10.4. The Control Frontier

The Pareto frontier analysis over a 9×8 grid of (K_p, K_d) values reveals that control effort E_{ctrl} and tracking error ε_ω are inversely related in a curved frontier. The knee of this curve, where additional control effort yields diminishing improvement in tracking, occurs at approximately $K_p = 0.8$, $K_d = 0.15$. Lower K_d moves along the frontier toward lower effort and higher error (looser control); higher K_p moves toward lower error and higher effort (tighter control). The Pareto-optimal gain choice

depends on the application: a generator requiring tight frequency regulation demands low ε_ω and accepts higher E_{ctrl} ; a flywheel energy store tolerates higher ε_ω in exchange for lower E_{ctrl} and thus better round-trip efficiency.

II. DISCUSSION

II.1. *The Electromagnetic Subsystem as Constraint Enforcer*

The central theoretical reinterpretation proposed here is that the electromagnetic subsystem should not be understood as a *motor* in the conventional sense. A motor converts electrical energy into mechanical energy as its primary function. In the Flash-Reluctance Engine, the electromagnetic subsystem converts electrical energy into coherence in the torque waveform. It is spending energy to enforce a target trajectory in torque-space rather than to produce motion independently.

This distinction matters for how one reasons about the system. The relevant performance question for a motor is: how much torque can it deliver? The relevant question for the electromagnetic layer of a Flash-Reluctance Engine is: how much energy is required to maintain a given torque profile? The answer depends not only on the gains (K_p, K_d, K_ω) but on the shape and reliability of the thermal torque waveform. A high-quality thermal side with smooth, predictable pulses requires very little electromagnetic correction. A degraded thermal side—cold chambers, inconsistent injection, worn nozzles—demands heavy electromagnetic compensation, which appears as a rising f_{EM} and is thus a diagnostic signal as well as a performance metric.

II.2. *The Shared Core as a Global Constraint*

The shared thermal core introduces a feature absent from systems where each chamber is independently heated: global resource competition. All six chambers draw from the same reservoir, and if the aggregate demand exceeds the input rate \dot{Q}_{input} , the core cools. This creates an implicit coupling between chambers that is not present in the individual temperature equations (4) but emerges from the shared dynamics (5).

This coupling has a self-organizing character. Chambers that are currently injecting cool the core, which slows the heating of other chambers, which reduces their injection energy, which reduces their cooling effect on the core. The system reaches a quasi-steady state in which the aggregate injection rate is automatically regulated by the core temperature, without any explicit inter-chamber controller. This emergent regulation is a form of resource-constraint feedback that does not require any additional control logic.

II.3. *Relation to Known Architectures*

The Flash-Reluctance Engine bears family resemblance to several existing technologies without being identical to any of them. It resembles a Stirling engine in the sense that it exploits a temperature gradient between a hot source and a cold sink, but unlike a Stirling engine, the working fluid (water-steam) is consumed rather than recycled. It resembles a flash-steam engine in its use of rapid phase change, but unlike conventional flash-steam engines, the expansion chambers are arranged radially and the steam acts directly on rotor vanes rather than on pistons. It resembles a switched-reluctance motor in its stator topology, but unlike a standalone reluctance motor, the electromagnetic layer operates under feedback

from a thermal primary drive rather than as the sole source of torque. And it resembles a flywheel energy storage system in its use of rotor inertia as a buffer, but unlike a passive flywheel, the rotor is actively driven and loaded.

The key novelty is the simultaneous co-location of all three functions—thermal expansion, electromagnetic control, and inertial storage—on a single rotor with no interface separating them. This co-location eliminates the transmission losses inherent in hierarchical hybrid architectures and allows the electromagnetic layer to respond to thermal fluctuations on a timescale limited only by the coil’s electrical dynamics and the controller’s update rate.

11.4. Relation to Control Theory

From the perspective of control theory, the system is a nonlinear SISO (single-input, single-output) plant—with ω as the output and τ_{EM} as the control input—driven by a stochastic-like exogenous disturbance $\tau_{th}(\theta, \mathbf{T}_{wall})$. The plant nonlinearity arises from the quadratic drag term in τ_{fr} , the nonlinear load-speed relationship, and the temperature-dependent thermal torque. The controller (21) is a PD law augmented with a speed integral term (through the $K_{\omega}e_{\omega}$ component), which in the language of control theory makes it a PID controller in the speed loop with an inner proportional-derivative loop in the torque space.

The Pareto frontier analysis is equivalent to a multi-objective optimization in the gain space $(K_p, K_d) \in \mathbb{R}_+^2$ with objectives E_{ctrl} and ε_{ω} . The Pareto-efficient set forms a curve in $(E_{ctrl}, \varepsilon_{\omega})$ space that is convex (in the sense that any point on the frontier is a convex combination of no two other frontier points with strictly lower values of both objectives). The knee of this curve is the natural gain selection criterion for applications where both objectives are equally weighted.

11.5. Relation to RSVP Field Theory

The multimodal coupling in the Flash-Reluctance Engine has a structural analogy with the Relativistic Scalar-Vector Plenum (RSVP) framework developed in the author’s parallel theoretical work. In RSVP, a scalar density field Φ , a vector flow field \mathbf{v} , and an entropy field S interact through coupled evolution equations in which energy is continuously redistributed across domains. The Flash-Reluctance Engine instantiates an analogous triad: the thermal field (scalar density of heat, concentrated in the chamber walls and core), the electromagnetic field (vector potential driving coil currents), and the mechanical field (angular momentum stored in the flywheel). The feedback controller is the analog of the RSVP coupling kernel: it mediates energy transfer between domains in response to local gradients (torque error and speed error), maintaining a target trajectory that is the mechanical analog of the plenum’s dynamical attractor.

This analogy is not merely metaphorical. Both systems are governed by the principle that energy redistribution is driven by constraint violation: in RSVP, constraints arise from the field equations; in the Flash-Reluctance Engine, constraints arise from the target torque profile and speed setpoint. The architecture of the controller—proportional response to instantaneous error, derivative damping of rapid fluctuations, integral accumulation of persistent offset through the speed term—mirrors the structure of variational methods in field theory, where minimization of an action functional drives the system toward a ground state.

12. EXTENSIONS AND FUTURE DIRECTIONS

12.1. *Dynamic Injection Scheduling*

In the baseline model, injection occurs unconditionally whenever the rotor angle is within a chamber's active arc and the wall temperature exceeds T_{\min} . A natural extension is to make the injection decision a control variable: at each chamber activation, the controller determines not only whether to inject but how much mass $m_{\text{inj},i}$ to inject, based on current wall temperature, rotor speed, and torque demand. This transforms the system from a fixed-cycle engine into a programmable thermodynamic machine capable of adjusting its thermodynamic work per cycle in response to load.

Variable injection mass modifies the flash pressure (15) and injection energy loss (14) proportionally. The controller gains a new degree of freedom at the cost of a more complex actuator (a variable-rate solenoid valve rather than a fixed-orifice nozzle). The tradeoff between thermal efficiency and control bandwidth becomes explicit: large injection mass per cycle produces high torque but cools the wall quickly; small injection mass per cycle preserves wall temperature but produces low torque per cycle and requires higher RPM to maintain average power.

12.2. *Multi-Stage Rotor Architecture*

A two-stage rotor places a second set of chambers and stator coils on a concentric inner rotor sharing the same shaft. The inner rotor operates at a different temperature range than the outer rotor, allowing a temperature cascade: the outer rotor extracts work from the high-temperature source, its exhaust steam is used to heat the inner rotor's chambers, and the inner rotor extracts additional work from the intermediate-temperature steam. This is a rotary implementation of the cascaded heat engine concept, analogous to combined-cycle gas turbines but in a compact radial geometry.

The shared shaft means the two rotors must be in torque balance: the controller must distribute electromagnetic assist across both stages simultaneously, managing the thermal budget of each independently. The shared core model extends naturally to this architecture by introducing separate core temperatures for the two stages, coupled through a heat exchanger that transfers waste heat from the outer stage to the inner stage's heat source.

12.3. *Generator Coupling and Grid Interaction*

The load model (10) approximates a generator as a linear speed-proportional resistance. A more realistic generator model would include back-EMF dynamics, field winding response, and the interaction with an external electrical grid or load bus. Coupling the mechanical simulation to an electrical circuit model would allow the study of transient responses to load steps (sudden increase or decrease in electrical demand), which are the most demanding scenarios for speed regulation.

The electromagnetic stator layer could also be operated in generator mode during phases of excess mechanical energy, converting some of the flywheel's kinetic energy back into electrical energy during deceleration. This bi-directional operation would make the electromagnetic layer a true power converter rather than a unidirectional assist, at the cost of a more complex power electronics interface.

12.4. Experimental Validation Pathway

The simulation framework generates falsifiable predictions that can be tested on a physical prototype. The most tractable first experiment is a single-chamber bench test: a fixed chamber with a controlled heater, a single injection solenoid, and a small linear actuator representing the vane, allowing direct measurement of injection force as a function of wall temperature and water mass. This would validate the flash pressure model (15) and the efficiency factor η_{th} .

The second experiment is a multi-chamber static test: a full rotor with all chambers but no rotation (locked shaft), measuring torque as a function of angle for a single revolution under controlled injection. This would validate the pulse shape function (16) and the per-chamber temperature dynamics (4).

The third and most informative experiment is a dynamic test under load, using a calibrated brake dynamometer to impose a known $\tau_{load}(\omega)$ and comparing measured speed trajectories, electromagnetic current draws, and chamber temperatures against simulation predictions. Discrepancies between simulation and experiment would identify the dominant modeling errors and guide refinement of η_{th} , the heat transfer coefficients, and the controller gains.

12.5. Thermal Storage and Phase-Change Materials

Replacing the simple wall heat capacity model with a phase-change material (PCM) in the chamber wall would introduce a qualitatively new thermal behavior. A PCM stores large amounts of energy at nearly constant temperature during the phase transition, then releases it at the same temperature during solidification. If the PCM transition temperature is chosen near the optimal wall temperature for flash evaporation, the chamber wall would act as a thermal buffer that resists temperature excursions in both directions. This would reduce the amplitude of the depletion-recovery oscillations, stabilize the thermal torque profile, and potentially allow operation at lower source temperatures by concentrating heat at the transition temperature rather than distributing it across a broad range.

13. CONCLUSION

The Flash-Reluctance Engine is not easily classified. It is not an engine in the Carnot sense: it does not operate on a closed thermodynamic cycle. It is not a motor in the electromagnetic sense: it cannot produce motion from electrical energy alone. It is not a hybrid in the conventional sense: there is no separation of function between a primary thermal stage and an auxiliary electromagnetic stage. It is a new class of machine in which thermal, electromagnetic, and mechanical phenomena are co-equal, co-located, and dynamically coupled through a feedback architecture that treats torque coherence as a design objective rather than an incidental outcome.

The principal contributions of this paper are as follows. We have developed a dynamic thermal model incorporating per-chamber temperature evolution and shared core thermodynamics, replacing the constant-wall-temperature assumption of earlier analyses with a physically realistic depletion-recovery cycle. We have formulated a PD speed-governed controller that simultaneously shapes the torque waveform and regulates rotor speed under load, converting the free-spinning flywheel model into a working machine. We have introduced the slope-based spatial distribution of electromagnetic torque, a predictive allocation strategy that assigns electromagnetic assist to chambers where thermal torque is declining most rapidly, reducing torque ripple with lower control effort than reactive strategies.

We have defined the control effort versus tracking error Pareto frontier as the primary design instrument, replacing ad hoc gain tuning with a systematic optimization framework. And we have characterized four emergent operating regimes—steady state, oscillatory, stall, and thermal collapse—and identified their boundaries in the parameter space.

The deepest insight is perhaps the simplest one to state. In a machine where multiple energy domains act on the same rotor, the design question is not how to maximize each domain's contribution independently, but how to coordinate them so that their combined trajectory is better than any single domain could achieve. The Flash-Reluctance Engine answers that question with a feedback architecture that continuously measures the gap between actual and desired behavior and assigns electromagnetic energy precisely to close it, while allowing the thermal side to carry as much of the load as the physical constraints permit. The result is a machine that is, in a meaningful sense, aware of its own performance and responsive to its own limitations—not by design of any individual component, but as an emergent property of the coupling between them.