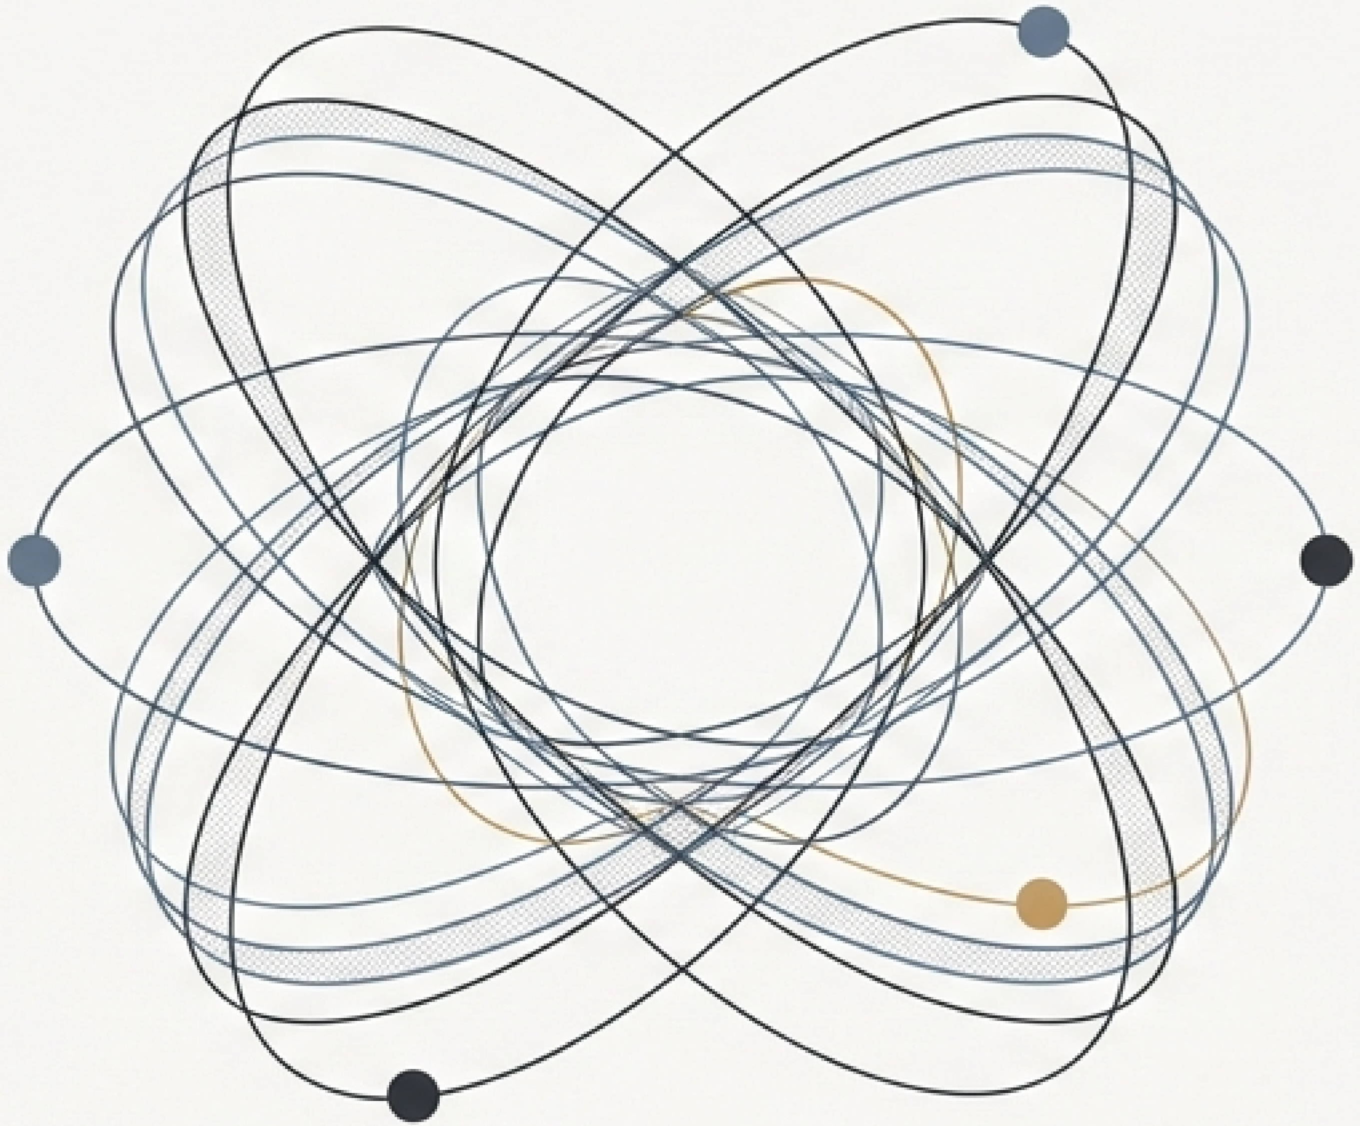

Vetting the Corrective Algorithm

Technical Diligence, Viability Analysis, and Validation Plan
for Proposed Hamiltonian Middleware

High potential upside coupled with significant prior-art risk requires a gated technical evaluation

The Pitch

A middleware algorithm claiming **~300x** accuracy gains and **<2%** compute overhead for N-body and Hamiltonian simulations.

The Reality

The underlying math (symplectic correctors) is credible and well-documented. However, IP defensibility is low due to 1990s prior art, and commercial delivery as a hosted API introduces major security friction.

The Recommendation

Proceed to a strict, multi-stage technical gate. Do not escalate to partnership or investment discussions without empirical, reproducible artifact validation against modern baselines.

Deconstructing the core claims from the introductory screening call

Performance

~300x <2%

Claims ~300x accuracy improvement with less than 2% compute overhead. Mentions a “14^2” magnitude.

Impact

Solves **“energy drift”** in long-horizon predictions **(100-200 years out)**. Translates software efficiency directly into GPU energy savings.

Delivery & Hardware

Packaged as an API / “middle layer algorithm”. Stated to be strictly **hardware agnostic**.

Target Domains

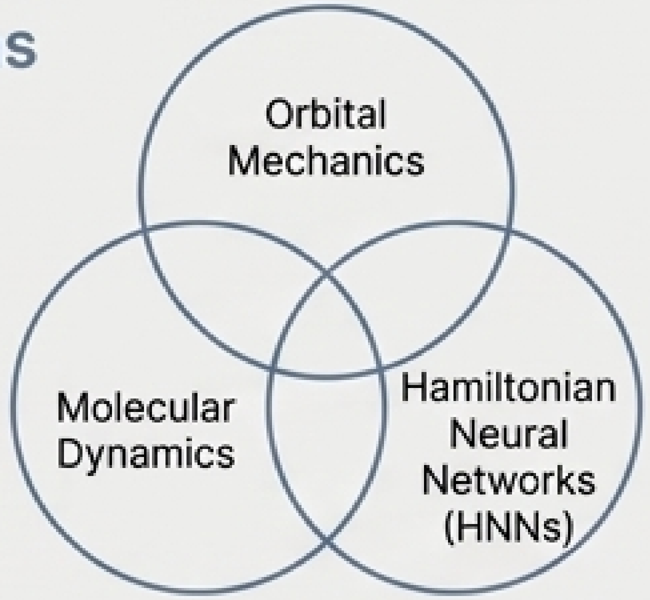

The vocabulary maps directly to established structure-preserving numerical methods

Correctional Algorithm	**Symplectic Correctors / Processed Integrators** **Context:** Introduced by Jack Wisdom et al. in the 1990s to reduce spurious oscillations in N-body problems.
<2% Overhead	**Output-Point Application** **Context:** Correctors applied only at output points bypass per-step compute costs, perfectly explaining the minimal overhead claim.
Energy Drift	**Shadow Hamiltonians** **Context:** Symplectic methods preserve geometric structure, keeping energy behavior well-behaved over massive timescales.

A 300x improvement is highly contingent on the strength of the baseline

Against Generic Methods

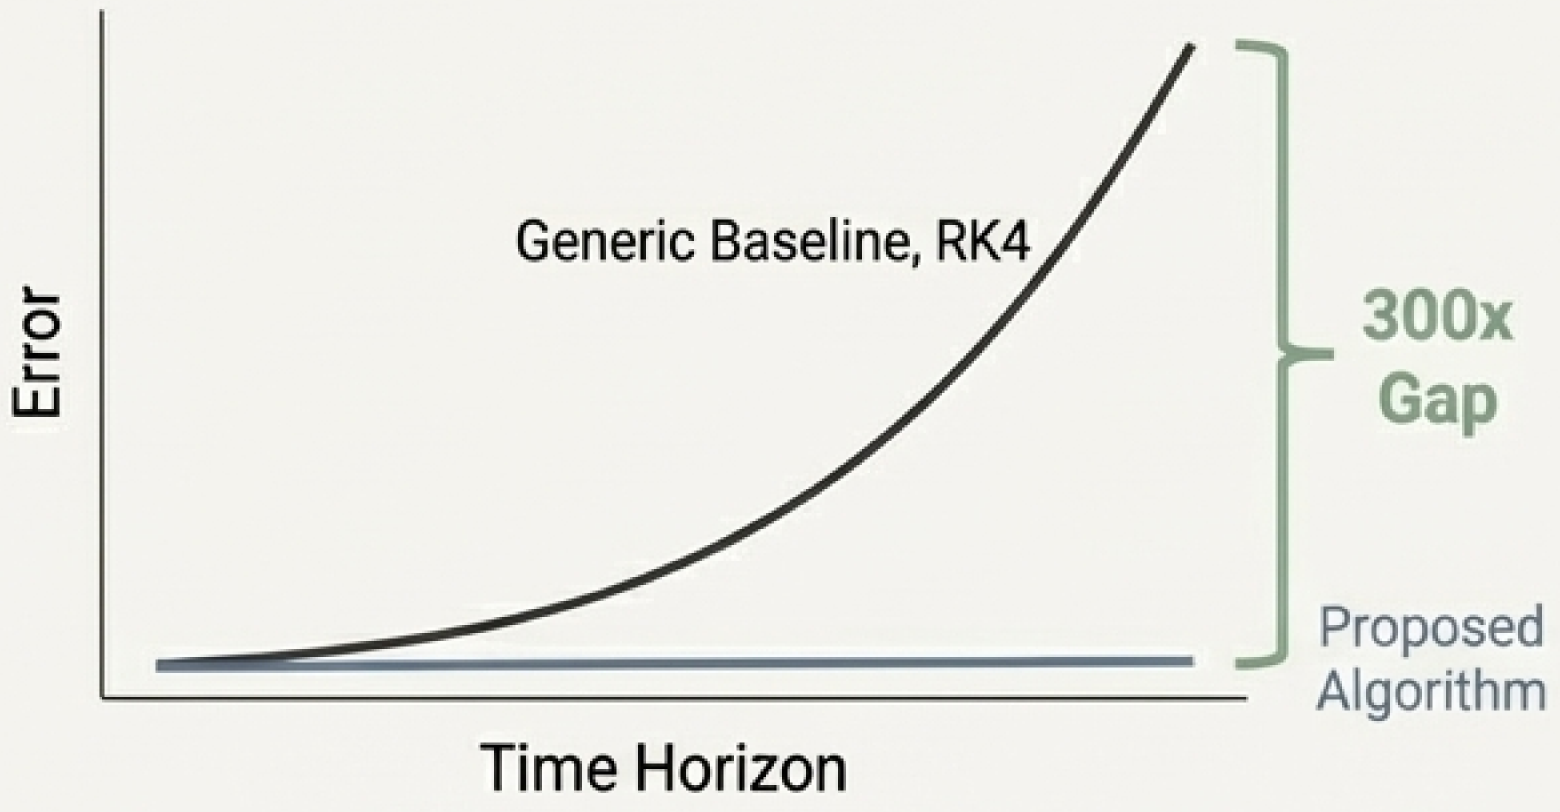

A 300x leap is highly credible if the baseline is a naive integrator (e.g., standard RK4) that fails to preserve Hamiltonian structure over long horizons.

Against Modern Best-in-Class

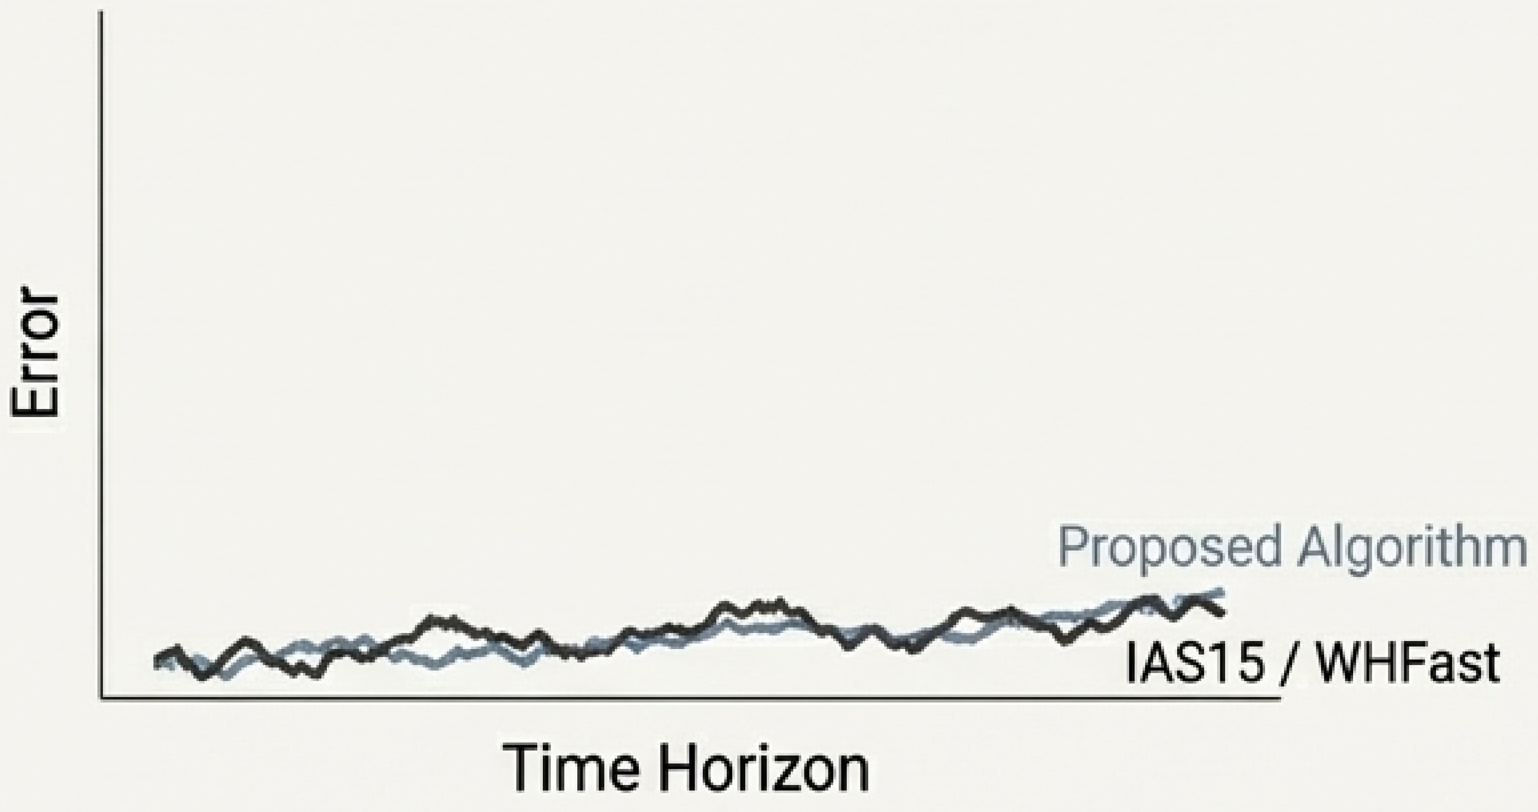

A 300x leap is highly **unlikely** if compared to optimized methods like **IAS15** (15th-order adaptive) or **WHFast** (which already utilizes symplectic correctors). In high-accuracy regimes, energy error already behaves like a random walk (Brouwer's law).

Energy savings are only viable if the numerical integrator is the true system bottleneck

Algorithmic Efficiency → Reduced Time-to-Solution → Lower GPU Power Draw

The Macro Context: Strategic relevance is high. DOE and IEA project massive data center power demand growth through 2030.

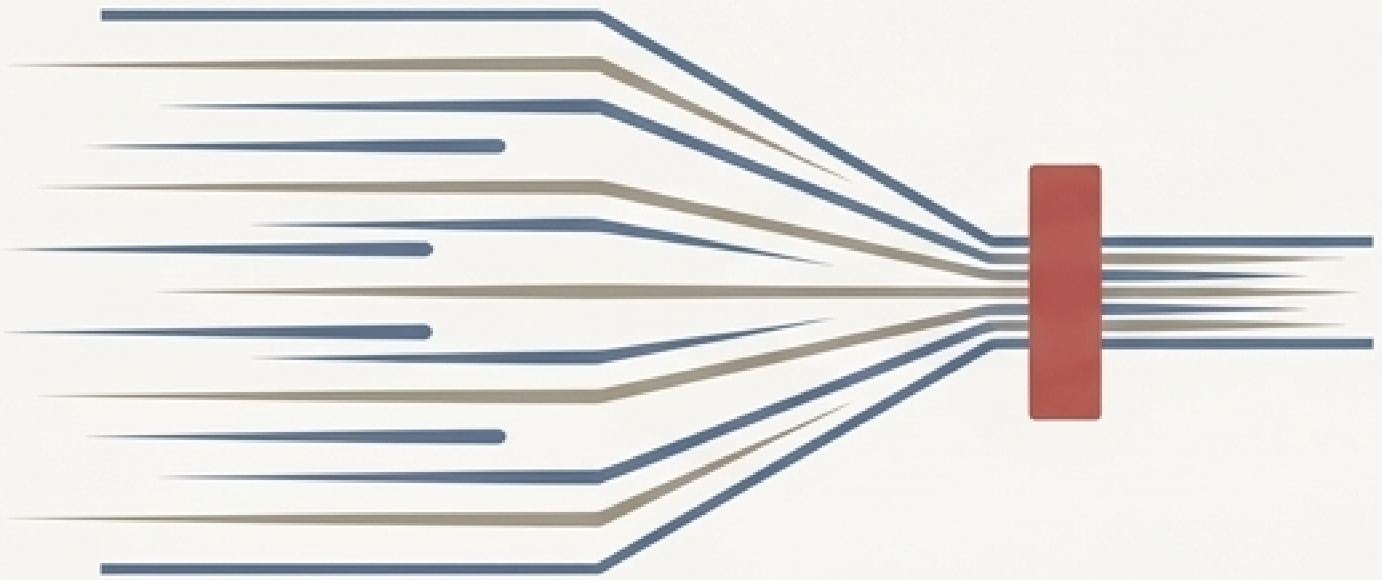

Warning: **Kernel gains do not equal workflow gains.** If the operational pipeline is bottlenecked by I/O, communication, or memory constraints, the total energy savings will be negligible.

The proposed API delivery model fundamentally clashes with target customer reality

Tension Diagram

Hosted API Endpoint

- Black box
- Remote compute
- Recurring revenue model

Secure Operational Environment

- High data gravity
- Extreme security
- Zero latency tolerance

> Aerospace and defense workloads (e.g., NASA CARA standards for operational conjunction assessment tracking ~$10^4$4 objects) mandate **in-process libraries (C/C++/Fortran) or isolated in-VPC deployments.** A remote API is commercially unviable.

High prior-art risk dictates that defensibility will stem from execution, not patents

- ✓ **Crowded Concept Space**: Symplectic correctors and processed-splitting methods have a rich, published literature trail dating back to the 1990s.
- ✓ **The ‘No Free Lunch’ Principle:** Algorithmic improvements must exploit specific structural features. Gains in one class of Hamiltonian problems imply reduced advantages in others.

> Assume the moat is execution and distribution—novel corrector construction, tight C++
> Assume the moat is execution and distribution—novel corrector construction, tight C++
> implementation, or automated splitting discovery—rather than broad patent protection.

Pre-requisites to prevent black-box optimism and enable empirical evaluation

1. Defined Problem Statement

Explicit boundaries. What is in scope (separable Hamiltonians, perturbations) and what is strictly out of scope.

2. Precise Accuracy Metrics

Definition of success for demo domains (RMS state error, maximum position error, conserved quantity drift).

3. Stated Baselines

Required minimum comparisons against modern reference integrators (e.g., IAS15, WHFast).

4. Reproducible Harness

Delivery of a containerized environment, fixed seeds, configuration files, and independent run instructions.

Executing a two-tier technical evaluation to isolate kernel truth and workflow reality

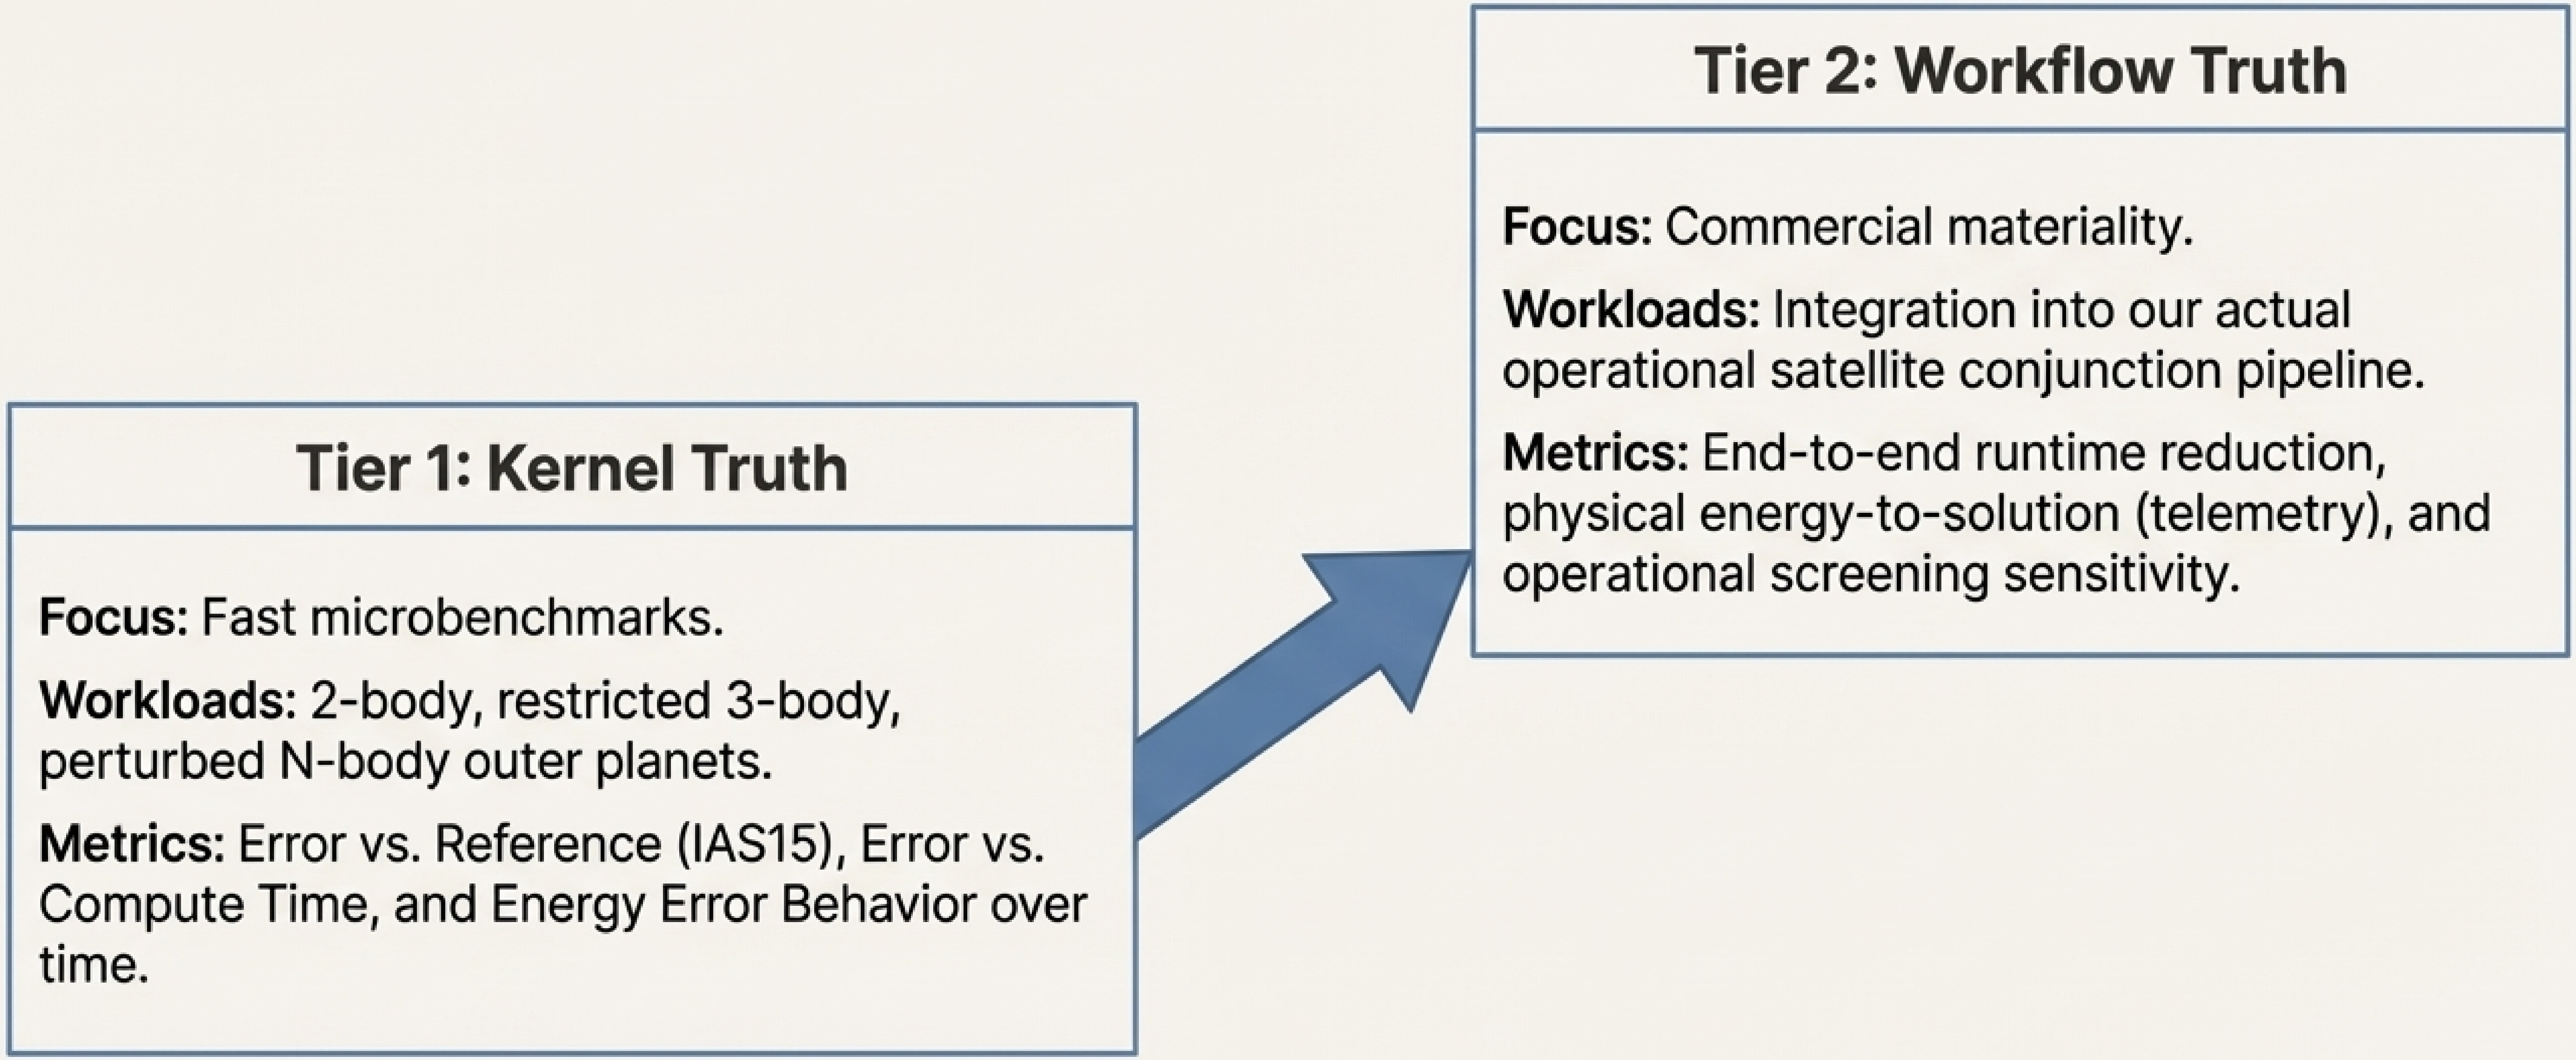

Tier 1: Kernel Truth

Focus: Fast microbenchmarks.

Workloads: 2-body, restricted 3-body, perturbed N-body outer planets.

Metrics: Error vs. Reference (IAS15), Error vs. Compute Time, and Energy Error Behavior over time.

Tier 2: Workflow Truth

Focus: Commercial materiality.

Workloads: Integration into our actual operational satellite conjunction pipeline.

Metrics: End-to-end runtime reduction, physical energy-to-solution (telemetry), and operational screening sensitivity.

Objective go / no-go thresholds for post-validation escalation

Proceed to Partnership Escalation

- ✓ Demonstrates material step-size increases while preserving operational accuracy.
- ✓ Unlocks net-new product capabilities (e.g., less frequent repropagations, longer-horizon stability).
- ✓ Maintains superior performance when tested against strong, modern baselines.

Halt Investigation

- ✗ Performance gains evaporate when tested against anything stronger than naive baselines.
- ✗ Algorithm fails or loses stability under real-world perturbations and non-idealities.
- ✗ A straightforward reimplementation of known correctors with no discernible defensible moat.