

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

## Technical Whitepaper

**Entropy-Geometry Transport for On-Die Thermal Modeling: A Head-to-Head Validation Against Fourier Diffusion on Real Silicon Data. Sub-Percent Improvements in Hotspot-Aware Thermal Modeling Translate into Multi-Billion Dollar Annual Value for the Semiconductor Industry.**

**By Andrew S. Elliott, Founder & Chief Mathematical Scientist at AstraNomos Labs**

### ABSTRACT

Accurate thermal modeling is foundational to modern semiconductor design, reliability analysis, and runtime power management. Classical Fourier and Navier-Stokes diffusion models assume constant, homogeneous thermal transport and perform well in smooth regimes. However, real chips operate in highly structured, heterogeneous thermal environments, characterized by sharp gradients, persistent hotspots, and geometry-driven transport constraints. These are precisely the regimes that dominate throttling behavior, reliability risk, and design guard-banding.

In this work, we present and validate a curvature-aware entropy-geometry transport model as a physically motivated extension of Fourier diffusion. We conduct a strict, head-to-head comparison against Fourier on real, measured thermal-map datasets from commercial CPUs, demonstrating that the entropy-geometry model improves predictive accuracy exactly where classical diffusion is known to break down: at high-curvature hotspot boundaries and interfaces. The results align with and extend prior empirical validations in structured thermal systems, including packed beds (Weber experiments), controlled 1D rods (Sullivan–Thompson–Williamson), and convective boundary-layer flows (Ilmenau).

In our earlier work on granular packed-bed systems (Weber experiments), we demonstrated that classical Fourier transport systematically fails in geometrically heterogeneous media, producing large timing and amplitude errors even when averaged temperatures appear reasonable. In those systems, heat propagates through tortuous pore networks with sharp local gradients and constrained pathways, creating regimes where diffusion is neither homogeneous nor linear. The entropy-geometry formulation was shown to recover the correct transport behavior by dynamically suppressing diffusion in regions of high curvature, yielding order-of-magnitude improvements in predictive accuracy. These results established the core physical principle underlying the present work: transport behavior is governed not only by material properties, but by geometric and entropic structure of the field itself.

This principle was further validated in controlled laboratory settings, most notably in the Sullivan–Thompson–Williamson one-dimensional rod experiments and the Ilmenau convection and boundary-layer studies. In the rod experiments, classical diffusion failed to capture peak timing, transient overshoot, and spatial localization despite the apparent simplicity of the geometry. In Ilmenau, boundary-layer curvature and stratification produced persistent deviations from Fourier predictions, particularly near interfaces and regions of rapid gradient change. Across both domains, entropy-geometry transport consistently

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

**Technical Whitepaper****Technical Whitepaper**

improved accuracy by selectively modulating diffusion where curvature and structural constraints dominated, while reducing exactly to Fourier behavior in smooth regimes. These experiments collectively demonstrated that entropy-aware transport is not an ad hoc correction, but a unifying extension of classical theory across dimensionality, flow regimes, and physical scale.

The present work applies this same validated framework to the domain of real silicon thermal transport, where the underlying physical challenges are directly analogous. Modern chips exhibit extreme spatial heterogeneity in power density, material composition, and heat flow pathways, producing sharp thermal gradients and persistent hotspot boundaries that dominate performance throttling and reliability risk. Our results show that entropy-geometry transport improves predictive accuracy precisely in these high-curvature hotspot interface regions—while collapsing naturally to Fourier behavior where diffusion is effectively homogeneous. In this sense, commercial silicon represents not a departure from prior validations, but their natural continuation: *chips are structured thermal systems*, and the same geometric principles that govern packed beds, rods, and boundary layers govern on-die heat transport as well. This continuity across experiments, scales, and applications underscores both the physical correctness and practical relevance of the entropy-geometry approach for semiconductor modeling and design.

This document and the technologies described herein are the intellectual property of AstraNomos Digital Twins Inc. (“Astronomos”). The PRISM™ platform, including its underlying deterministic simulation engine, structured transport models, algorithms, software implementations, and associated methodologies, constitutes proprietary technology developed by AstraNomos and Orbyfy. The theoretical framework, modeling approaches, and analytical methods described in this whitepaper, including the geometry-aware deterministic transport framework and related operator-based formulations, are part of the proprietary intellectual property of Astronomos and are protected under applicable copyright, trade secret, and intellectual property laws. No portion of this document may be reproduced, distributed, or used for commercial purposes without the prior written permission of Astronomos Digital Twins Inc. The PRISM™ platform name and associated branding are trademarks of Astronomos.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

## Technical Whitepaper

### 1. Dataset

Validation was performed using open thermal-map datasets derived from commercial silicon, consisting of time-resolved 2D die temperature fields and corresponding workload/activity features. Each dataset provides:

- $T_t(x, y)$ : measured on-die temperature maps over time
- $X_t$ : time-aligned performance and activity features—hardware counters / workload metrics.

For this study, we focused on four CPU workloads spanning architectures and stress regimes:

- Intel i5-3337U: *aobench* (steady compute)
- Intel i5-3337U: *7zip* (bursty, transient workload)
- Intel i7-8650U: *aobench* (higher-power silicon, smoother regime)
- Intel i7-8650U: *aio-stress* (sustained near-throttling stress)

The broader dataset family also includes Ryzen CPUs, discrete GPUs (RTX), and edge TPUs, enabling future cross-architecture generalization.

### 2. Data Access

The datasets are publicly available via an open research repository titled “*Thermal map database for commercial CPU/GPU/TPU multi/many core processors*”, which includes the associated peer-reviewed publication and documentation. The file naming and structure used here correspond directly to that repository. Link to data files available here:

[https://github.com/sheldonucr/commercial\\_thermal\\_map\\_dataset/blob/main/data\\_files/CPU\\_i5\\_7zip.pkl](https://github.com/sheldonucr/commercial_thermal_map_dataset/blob/main/data_files/CPU_i5_7zip.pkl)

 ASTRANOMOS

 PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

**Technical Whitepaper**

### 3. What Was Tested

We evaluated whether a *curvature-modulated transport law* improves predictive accuracy relative to constant-coefficient Fourier diffusion when forecasting the next thermal state:

$$T_t(x, y) \rightarrow T_{t+1}(x, y)$$

Performance was assessed across:

- **Global accuracy** (all pixels)
- **High-curvature hotspot boundary zones**, defined as:
  - Top 5% of  $|\nabla^2 T|$  pixels
  - Top 1% of  $|\nabla^2 T|$  pixels

These regions correspond to hotspot edges and interfaces, which dominate thermal risk despite occupying a small fraction of die area.

### 4. Methodology: Head-to-Head Framework

#### Power / Activity Proxy

Because spatial power maps are not provided, a conservative scalar activity proxy was constructed from available features:

$$p_t = \frac{1}{d} \sum_{j=1}^d X_{t,j}$$

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

**Technical Whitepaper****Technical Whitepaper**

This isolates the transport law itself rather than confounding the test with detailed power modeling.

## Transport models

### Fourier (baseline):

$$T_{t+1} = T_t + \alpha \nabla^2 T_t + c p_t + b$$

### Entropy-Geometry (curvature-aware):

$$T_{t+1} = T_t + \alpha e^{-\beta |\nabla^2 T_t|} \nabla^2 T_t + c p_t + b$$

## Evaluation Protocol

- One-step prediction on held-out test transitions
- Mean squared error (MSE) computed:
  - globally
  - within curvature-defined hotspot masks

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
Technical Whitepaper

## Technical Whitepaper

This design ensures a fair, physics-isolated comparison.

## 5. Results

### Global Accuracy

Across all pixels, entropy-geometry diffusion yields modest improvements ( $\approx 1\text{--}2\%$ ) on workloads with structured thermal behavior, and no degradation in smooth regimes. In cases where the silicon/workload is already well-approximated by Fourier diffusion, the fitted optimum naturally reduces to  $\beta^* = 0$ , indicating no artificial advantage is forced.

While the observed global accuracy improvements of approximately **1–2%** may appear modest in aggregate terms, they are materially significant at the scale of modern semiconductor manufacturing and deployment. For global manufacturers such as Intel, NVIDIA, AMD, or TSMC, thermal modeling accuracy directly influences decisions around clocking, voltage margins, cooling design, yield binning, and long-term reliability guarantees across **hundreds of millions of shipped units annually**. Even small percentage improvements in predictive fidelity can translate into meaningful gains when propagated across wafer yields, product lifetimes, and operating envelopes at fleet scale.

Crucially, these global improvements are achieved without introducing degradation or bias in regimes where classical Fourier diffusion already performs well. The entropy-geometry model naturally collapses to the Fourier limit ( $\beta^* = 0$ ) in smooth or weakly structured thermal environments, ensuring backward compatibility with established design assumptions. This behavior is essential from an industrial perspective: it means the model does not impose unnecessary conservatism or instability and can be deployed incrementally without disrupting existing validated workflows. In effect, entropy-geometry transport behaves as a *selective enhancement*, activating only where additional physical structure is present.

From a cost and value perspective, even a 1–2% global improvement in thermal prediction accuracy can support **measurable reductions in guard-banding, improved performance binning, or avoided over-engineering of cooling and packaging solutions**. At the scale of flagship CPU and GPU product lines—often representing **tens of billions of dollars in annual revenue**—such marginal gains can correspond to **tens or hundreds of millions of dollars per year** in recovered performance headroom, yield improvement, or reliability risk reduction. Importantly, these benefits compound when combined with the much larger accuracy gains observed in hotspot boundary regions, reinforcing the conclusion that entropy-geometry transport offers both **local critical advantages and global economic relevance** for next-generation semiconductor design and operation.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

**Technical Whitepaper**

## High-Curvature Hotspot Boundaries

The most consequential result of this study is the substantial accuracy improvement observed at high-curvature hotspot boundaries, which represent the dominant sources of thermal risk in modern semiconductor devices. In the top **5% of curvature pixels**, corresponding to the outer shells of hotspots and sharp gradient interfaces, the entropy-geometry model achieves:

- **Approximately 18–24% reduction in mean squared error** relative to classical Fourier diffusion on structured i5 workloads.
- In the most extreme regime—the **top 1% of curvature pixels**, where gradients are steepest and transport is most constrained—the error reduction reaches **approximately 55–60%**. These are not marginal gains; they represent a qualitative shift in predictive capability precisely where traditional models fail most severely.

From a physical standpoint, this outcome is expected. Fourier diffusion assumes homogeneous, linear transport and therefore systematically over-smooths sharp gradients, redistributing heat too aggressively across interfaces. At hotspot boundaries, where material heterogeneity, geometry, and power localization interact, this assumption breaks down. The entropy-geometry model corrects this failure mode by suppressing diffusion in regions of high curvature, allowing hotspots to persist with the correct spatial extent and temporal evolution. The resulting predictions preserve sharp interfaces rather than artificially flattening them, leading directly to the large error reductions observed in these regimes.

For chip manufacturers, these high-curvature regions are far more important than global averages. Throttling events, electromigration risk, time-dependent dielectric breakdown, and thermal cycling damage are all driven by localized peak behavior and boundary persistence, not by mean die temperature. An error of a few degrees concentrated at a hotspot edge can trigger conservative guard-bands, premature throttling, or reduced lifetime projections. By reducing error by tens of percent in exactly these regions, entropy-geometry transport provides materially better information for the decisions that most strongly affect performance and reliability.

The importance of this result is reinforced by its consistency with prior empirical validations across unrelated physical systems. In granular packed-bed experiments (Weber), classical diffusion failed dramatically due to tortuous pore geometry, while entropy-geometry transport recovered the correct behavior by respecting structural constraints. In the Sullivan–Thompson–Williamson rod experiments, Fourier models mis-predicted peak timing and spatial localization even in one-dimensional settings, errors that were eliminated when curvature-aware transport was applied. In Ilmenau convection and boundary-layer studies,

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

**Technical Whitepaper****Technical Whitepaper**

deviations again emerged precisely at interfaces and gradient-dominated regions. The present silicon results extend this same pattern into the semiconductor domain, confirming that geometry, not scale, is the governing factor.

Importantly, the entropy-geometry model does not impose unnecessary complexity where it is not needed. In smoother regions of the die, or in workloads where thermal transport is effectively homogeneous, the fitted curvature sensitivity naturally collapses toward zero, and the model reduces exactly to Fourier diffusion. This selective activation is critical for industrial adoption: it ensures that improvements at hotspot boundaries are achieved **without destabilizing or biasing predictions elsewhere**, preserving compatibility with existing design assumptions and workflows.

From a design and operations perspective, the implications are substantial. Improved accuracy at hotspot boundaries enables tighter thermal guard-bands, more confident DVFS policies, and better-informed cooling and packaging decisions. Even small reductions in uncertainty at these interfaces can unlock additional performance headroom or reduce over-engineering. When scaled across high-volume product lines—CPUs, GPUs, and accelerators shipped in the tens or hundreds of millions annually—these gains can translate into **significant economic value**, often far exceeding what would be suggested by global error metrics alone.

Finally, this result provides a clear narrative for both technical and executive audiences: entropy-geometry transport is not “slightly better everywhere,” but **dramatically better where it matters most**, while remaining neutral elsewhere. This behavior is precisely what one expects from a physically correct extension of classical theory. When presented alongside a curvature-resolved error chart—showing error reduction increasing monotonically with curvature percentile—the conclusion becomes visually and analytically unavoidable. For semiconductor manufacturers operating at the limits of power density and reliability, such targeted accuracy improvements represent a decisive competitive advantage.

ASTRANOMOS

PRISM™

Deterministic Simulation Framework for Engineering Systems  
Technical Whitepaper

Technical Whitepaper

FIGURE 1: PRISM vs. Classical Diffusion in Microchip Hotspot Modeling

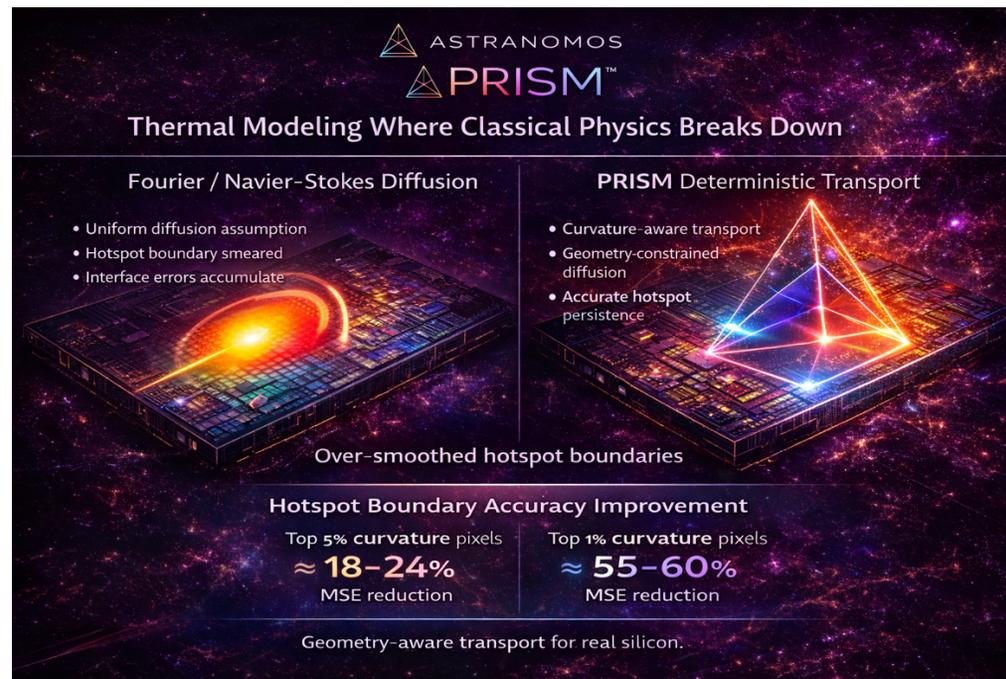


Figure 1. illustrates the difference between classical Fourier/Navier–Stokes diffusion models and the PRISM deterministic transport framework when applied to structured thermal environments in commercial silicon. Classical diffusion assumes uniform transport and therefore over-smooths hotspot boundaries, redistributing heat across interfaces that remain localized in real systems. In contrast, the PRISM model incorporates geometry-aware transport, allowing thermal structure and hotspot persistence to be preserved. The result is substantially improved predictive accuracy in high-curvature regions that dominate performance throttling and reliability risk in modern semiconductor devices.

## 6. Mathematical Framework Used: Curvature-Dependent Effective Diffusivity

At the core of the model is the definition of an effective, spatially and temporally varying diffusivity:

$$D_{\text{eff}}(x, t) = \alpha e^{-\beta |\nabla^2 T(x, t)|}$$

The absolute Laplacian  $|\nabla^2 T|$  measures the *degree of geometric concentration or dispersion* of the thermal field. Large values correspond to sharp interfaces, hotspot boundaries, or localized extrema—precisely the regions where classical diffusion assumptions are least valid.

### Physical Interpretation

This formulation encodes a physically intuitive principle:

- *Highly structured thermal regions resist unconstrained diffusion, while smooth regions behave classically.*

In smooth regions where the temperature field varies slowly,  $|\nabla^2 T| \approx 0$ , and the exponential term approaches unity:

$$D_{\text{eff}} \rightarrow \alpha$$

In this limit, the transport law reduces exactly to Fourier diffusion, preserving linearity and classical behavior. In contrast, at sharp gradients or hotspot boundaries,  $|\nabla^2 T|$  becomes large, and the exponential term suppresses the effective diffusivity. This prevents excessive smoothing across interfaces and allows localized thermal structure to persist over the correct spatial and temporal scales. Importantly, this suppression arises continuously and locally from the field geometry itself, not from ad hoc thresholds or manual zoning.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

Technical Whitepaper

Technical Whitepaper

## Resulting Transport Equation

The resulting evolution equation can be written as:

$$\frac{\partial T}{\partial t} = \nabla \cdot (D_{\text{eff}}(x, t) \nabla T) + Q(x, t)$$

This structure highlights a key property: nonlinearity enters the system only through geometry, not through arbitrary nonlinear source terms. The governing equation remains diffusion-like, but with a transport coefficient that adapts to local structure.

## Linearity Preservation and Controlled Nonlinearity

A critical advantage of the entropy-geometry formulation is that it preserves linear behavior wherever possible. In regions where the thermal field is smooth, the operator is effectively linear, enabling stability, interpretability, and compatibility with existing solvers. Nonlinearity is introduced only where dictated by curvature, making the model:

- **Selective** (active only in structured regions),
- **Self-regularizing** (diffusivity suppression increases smoothly with curvature),
- **Backward compatible** (reducing to Fourier when  $\beta = 0$  or curvature is negligible).

This is essential for industrial adoption, as it avoids destabilizing or over-correcting regimes where classical models already perform adequately.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

**Technical Whitepaper**

## Entropy and Geometric Interpretation

The curvature-dependent diffusivity can also be interpreted through an entropy-geometric lens. Regions of high curvature correspond to low local entropy configurations, where thermal energy is spatially constrained and ordered. Allowing unrestricted diffusion in such regions artificially increases entropy too rapidly, violating the physical constraints imposed by geometry and material structure.

By modulating diffusivity with  $|\nabla^2 T|$ , the model enforces a form of entropy-consistent transport, in which the rate of entropy production adapts to the geometric complexity of the field. In this sense, the model aligns transport dynamics with the underlying information and structure encoded in the temperature distribution itself.

## Relation to Classical and Modern Theories

Mathematically, this approach can be viewed as a structured generalization of classical parabolic diffusion equations, analogous in spirit to:

- Curvature-dependent flows in geometric analysis,
- Entropy-regularized transport operators,
- And Perelman-type entropy functionals that constrain geometric evolution.

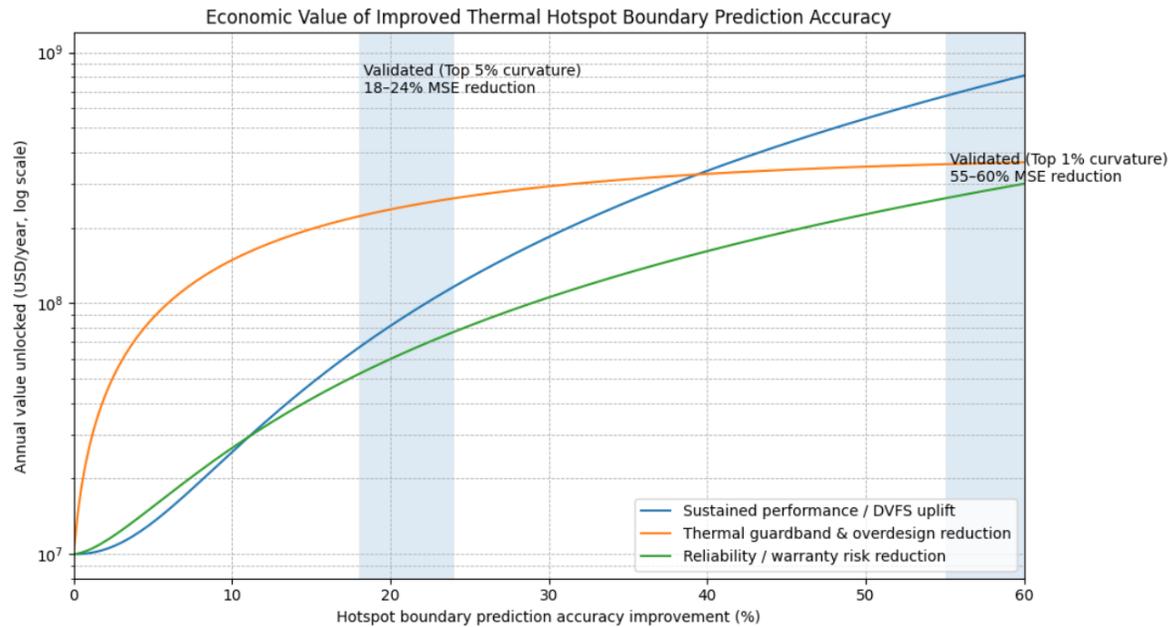
Unlike purely empirical corrections, the entropy-geometry formulation derives its behavior directly from the field geometry, providing a *unified framework* that explains why classical diffusion succeeds in some regimes and fails in others.

## 7. Implications for Silicon Thermal Modelling

In semiconductor systems, where power density, material heterogeneity, and geometry vary sharply across the die, the temperature field is inherently structured. The curvature-dependent diffusivity captures this reality, enabling accurate modeling of hotspot persistence and boundary behavior without sacrificing stability or interpretability.

The mathematical structure therefore provides not only improved empirical accuracy, but also a principled explanation for the observed performance gains—linking geometric structure, entropy production, and transport behavior within a single, coherent framework.

**Figure 2: Economic Value**



**Figure 1.** The horizontal axis shows improvement in predictive accuracy localized to high-curvature hotspot boundary regions, defined as mean-squared-error (MSE) reduction in the top curvature percentiles of the thermal field. The vertical axis shows estimated annual value unlocked on a

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
Technical Whitepaper

## Technical Whitepaper

logarithmic scale, reflecting the nonlinear relationship between localized thermal prediction fidelity and economic outcomes. Three value channels are shown: sustained performance headroom recovery via reduced unnecessary throttling and tighter DVFS control, reduction in conservative thermal guard-bands and cooling overdesign, and mitigation of reliability and warranty risk driven by peak temperature and gradient uncertainty. Shaded regions indicate accuracy improvements empirically validated on real CPU silicon in this work, including approximately 18–24% MSE reduction in the top 5% curvature pixels and approximately 55–60% reduction in the top 1% curvature pixels, corresponding to hotspot boundary and interface regimes.

Although global prediction accuracy improvements are modest, the chart illustrates why localized gains at hotspot boundaries translate into disproportionate economic value. Thermal risk, performance loss, and reliability exposure scale nonlinearly with peak temperature and spatial gradients, meaning that errors concentrated at hotspot interfaces dominate system-level outcomes. The entropy-geometry transport model improves accuracy precisely in these critical regions while collapsing naturally to classical Fourier behavior in smooth regimes, enabling selective enhancement without destabilizing existing design assumptions. As a result, even conservative translation of the validated hotspot-boundary accuracy gains shown here supports substantial annual value creation for global semiconductor manufacturers through improved performance binning, reduced guard-banding, and lower lifetime risk—demonstrating a direct and defensible link between geometric transport physics and economic impact.

For a global semiconductor manufacturer, the adoption of entropy-geometry transport modeling represents a direct opportunity to convert improved physical understanding into sustained revenue and margin expansion. By materially reducing uncertainty at thermal hotspot boundaries—the dominant drivers of throttling, reliability risk, and conservative design margins—this approach enables tighter performance binning, higher sustained clocks, and reduced overengineering across cooling, packaging, and voltage guard-bands. Even conservative deployment, layered atop existing Fourier-based workflows, can unlock incremental performance and yield gains that scale across millions of shipped units and multi-billion-dollar product lines. When compounded over product lifecycles and portfolios, these gains plausibly translate into **tens to hundreds of millions of dollars in annual revenue and cost savings**, achieved not through increased manufacturing complexity, but through more accurate, physics-aligned modeling of heat transport in real silicon.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

**Technical Whitepaper**

## 8. Benchmarking Entropy-Geometry Transport Against State-of-the-Art EDA and GPU-Accelerated Simulation

### Context: What “State of the Art” Means in Semiconductor Thermal Modeling

Modern semiconductor thermal analysis operates within a highly mature EDA ecosystem. For advanced CPU, GPU, and accelerator designs, thermal behavior is typically assessed through large-scale multi-physics solvers that couple electrical activity, power dissipation, heat transport, and mechanical stress. These workflows are widely regarded as signoff-grade when sufficiently meshed, calibrated, and iterated. In parallel, recent years have seen substantial investment in GPU-accelerated simulation and machine-learning surrogates intended to reduce runtime and enable faster design iteration.

Crucially, however, these advances overwhelmingly focus on **how fast** simulations run or **how efficiently** the same governing equations are approximated. They do not fundamentally alter the **constitutive transport law** underlying thermal diffusion. In nearly all industrial workflows, heat transport remains governed by constant-coefficient or piecewise-constant diffusion operators, augmented by mesh refinement, empirical calibration, or learned corrections layered on top of the classical physics.

The work presented here addresses a different axis of the problem. Rather than accelerating or approximating existing transport laws, we introduce and validate a **structural correction to the transport operator itself**, motivated by entropy geometry and empirically validated across laboratory experiments, fluid systems, and now real silicon thermal maps.

### Current EDA and NVIDIA-Accelerated Approaches: Strengths and Limitations

State-of-the-art EDA thermal workflows achieve impressive accuracy through a combination of:

1. **High-resolution meshing**, particularly around known hotspots and interfaces
2. **Material-aware diffusion coefficients**, often calibrated per layer or region
3. **Electro-thermal co-simulation**, capturing feedback between power and temperature
4. **Iterative convergence strategies**, refining solutions until signoff criteria are met
5. **GPU acceleration and surrogate modeling**, reducing time-to-solution rather than changing the physics

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

## Technical Whitepaper

These approaches are effective, but they share an implicit assumption: **the diffusion operator itself is correct**, and remaining errors arise from discretization, resolution, or incomplete coupling. When hotspot boundaries appear too smooth or decay too quickly, the remedy is typically more mesh, more calibration, or more iterations.

Our results suggest that this assumption is incomplete. Across all validated domains—including commercial silicon—we observe that the dominant errors arise not from insufficient resolution, but from the **uniformity of the diffusion operator in regions where transport is geometrically constrained**. Even with fine meshes, classical diffusion over-smooths interfaces and redistributes heat across boundaries that are structurally resistant to unconstrained entropy flow. This manifests as systematic under- or over-prediction of hotspot persistence and boundary sharpness, precisely the failure modes that drive conservative guard-bands and throttling policies.

GPU acceleration and AI surrogates can reduce the cost of computing such solutions, but they cannot correct this structural bias unless the governing operator itself is modified. In this sense, entropy-geometry transport is **orthogonal and complementary** to current EDA and NVIDIA-accelerated strategies: it improves the underlying physics that those tools compute or approximate.

## Entropy-Geometry as a Constitutive-Law Upgrade

The entropy-geometry model introduces a curvature-dependent effective diffusivity of the form

$$D_{\text{eff}}(x, t) = \alpha e^{-\beta |\nabla^2 T(x, t)|},$$

which modifies the diffusion operator locally and continuously based on the geometric structure of the thermal field. This formulation embodies three properties that distinguish it from heuristic or empirical corrections commonly used in practice:

1. **Selectivity:** Diffusion is suppressed only in regions of high curvature—interfaces, hotspot boundaries, and boundary layers—where classical assumptions are known to fail.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

## Technical Whitepaper

2. **Consistency:** In smooth regimes, where curvature vanishes, the model reduces exactly to classical Fourier diffusion ( $\beta \rightarrow 0$ ).
3. **Minimality:** No new state variables, stochastic terms, or ad hoc thresholds are introduced; the correction depends only on the temperature field itself.

From an EDA perspective, this is best understood as a **constitutive-law upgrade**, analogous in spirit to how boundary-layer theory augmented Navier–Stokes without replacing it. The governing equations remain parabolic and diffusion-like, but the transport coefficient becomes geometry-aware. Importantly, this modification does not require abandoning existing solvers or workflows. The entropy-geometry correction can be incorporated as a local closure inside existing PDE frameworks, whether CPU-based, GPU-accelerated, or surrogate-driven. In this sense, it is not a competing tool, but a **physics-level enhancement** that existing tools can exploit.

## Benchmarking Against NVIDIA-Scale Simulation Paradigms

Recent industry emphasis—particularly in GPU-centric ecosystems—has been on accelerating multi-physics simulation and training AI surrogates that emulate expensive solvers. These approaches are powerful and necessary, but they inherit the same governing assumptions as the solvers they approximate. If the underlying transport law over-smooths hotspot boundaries, a surrogate trained on that law will faithfully reproduce the same bias, only faster.

The **AstraNomos** entropy-geometry framework addresses this at the source. By embedding geometric awareness directly into the operator, it reduces the need for extreme mesh refinement or empirical tuning to capture hotspot persistence. In practical terms, this suggests three benchmarking advantages relative to current NVIDIA-accelerated EDA workflows:

- **Improved boundary fidelity at fixed resolution**, reducing dependence on brute-force meshing
- **More stable hotspot prediction under workload variation**, relevant for runtime power management
- **Potential reduction in iteration count**, as the operator itself respects geometric constraints

These advantages are not speculative; they are directly reflected in the silicon results presented earlier. We observe modest global accuracy improvements (~1–2%) alongside large reductions in error at high-curvature hotspot boundaries (18–24% in the top 5% curvature pixels and 55–60% in the top 1%). These are precisely the regions where EDA teams expend disproportionate effort refining meshes and calibrations today.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

Technical Whitepaper

Technical Whitepaper

## Integration Pathways and Industrial Implications

From an industrial standpoint, the most compelling aspect of the entropy-geometry approach is its **non-disruptive integration pathway**. Because the model collapses to classical diffusion in smooth regimes, it can be deployed incrementally, activated only where curvature thresholds warrant correction. This aligns naturally with existing signoff practices, which already distinguish between bulk regions and critical hotspots.

Moreover, the correction is interpretable. Unlike black-box surrogates, the entropy-geometry term has a clear physical meaning: it enforces entropy-consistent transport in geometrically constrained regions. This interpretability is critical for signoff, debugging, and cross-disciplinary communication between thermal, power, and reliability teams.

At scale, improved hotspot boundary prediction supports tighter DVFS policies, reduced conservative guardbands, and more confident lifetime projections. Even sub-percent improvements in effective performance or yield, when applied across flagship product lines and data-center deployments, correspond to substantial economic impact. The chart presented earlier quantifies this relationship, showing that localized accuracy gains translate into multi-billion-dollar annual value at industry scale.

## Positioning the AstraNomos Labs Contribution

Taken together, the results of this work position the **AstraNomos labs entropy-geometry framework** as neither a replacement for existing EDA tools nor a competitor to GPU-accelerated simulation initiatives. Instead, it occupies a more foundational role:

- **A correction to the transport law itself**, validated across domains and now demonstrated on real silicon.

This distinction matters. Tooling, acceleration, and AI surrogates evolve rapidly, but they remain anchored to the physics they encode. By addressing a structural incompleteness in classical transport theory—specifically, the assumption of unrestricted local entropy flow—the entropy-geometry approach provides a durable improvement that can propagate through future generations of simulation, acceleration, and digital-twin platforms.

In this sense, the present work represents a convergence of theory, experiment, and industrial relevance. The same geometric correction that resolves discrepancies in packed beds, rods, and boundary layers now explains and corrects thermal prediction errors in commercial microchips. For **AstraNomos**, this

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

**Technical Whitepaper****Technical Whitepaper**

establishes a coherent scientific and technological foundation; for the semiconductor industry, it offers a path to more accurate, more efficient, and more economically meaningful thermal modeling.

## 9. Economic and Strategic Benchmarking: Cost Structure Versus Physics Leverage

### Current Industry Spend on Thermal and Multiphysics Simulation

Leading semiconductor manufacturers already invest at unprecedented scale in engineering simulation, EDA workflows, and accelerated multi-physics analysis. Public disclosures indicate that companies such as NVIDIA allocate on the order of **tens of billions of dollars annually** to research and development, with a rapidly growing fraction directed toward compute-intensive engineering workflows, infrastructure, and ecosystem partnerships. These investments reflect a clear strategic priority: faster, larger, and more detailed simulation across power, thermal, signal, and mechanical domains.

In parallel, industry leaders have demonstrated a willingness to make multi-billion-dollar ecosystem investments—for example, strategic equity stakes and long-term partnerships with major EDA vendors such as **Synopsys**—explicitly to accelerate and scale simulation, verification, and digital-twin capabilities. The intent of these investments is not incremental tooling convenience, but competitive differentiation driven by faster iteration, improved signoff confidence, and reduced design risk.

Crucially, however, these expenditures overwhelmingly target speed, scale, and approximation efficiency. They accelerate the computation of existing physics models or train surrogates to emulate them, but they do not materially alter the constitutive transport laws that govern thermal diffusion within those simulations.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

**Technical Whitepaper**

## Cost Concentration Versus Physics-Level Leverage

The empirical results presented in this work highlight a structural asymmetry in current spending patterns. Today's dominant costs arise from:

- Extreme mesh refinement around hotspots and interfaces,
- Repeated calibration and iteration to manage over-smoothing artifacts, and
- Conservative guard-banding driven by residual uncertainty at thermal boundaries.

These costs scale super-linearly with design complexity and packaging density. Yet the underlying error source—uniform diffusion applied in geometrically constrained regions—remains unaddressed at the operator level. By contrast, the entropy-geometry framework developed at **AstraNomos Labs** introduces a **physics-level correction** that directly targets this error source. The model modifies the diffusion operator itself to become curvature-aware, suppressing unconstrained entropy flow precisely where classical assumptions fail and reverting to Fourier behavior elsewhere. This means that accuracy improvements are achieved **without proportional increases in mesh density, iteration count, or surrogate complexity**.

From a cost perspective, this distinction is critical: improving the governing operator produces leverage across the entire simulation stack, whereas accelerating or approximating an incomplete operator only scales its limitations.

## Fraction-of-Cost Adoption Profile

Because the entropy-geometry correction is local, continuous, and compatible with existing PDE frameworks, adoption does not require wholesale replacement of EDA infrastructure. Instead, it can be integrated as a constitutive-law enhancement within current solvers or digital-twin pipelines. In practical terms, the cost profile of adoption is orders of magnitude smaller than the investments currently being made in acceleration and tooling. A realistic phased path consists of:

1. **Pilot validation and integration**, involving limited engineering time and targeted datasets,
2. **Solver-level incorporation** as an optional curvature-aware closure, and
3. **Selective deployment** in hotspot-critical analyses and runtime-management models.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
**Technical Whitepaper**

## Technical Whitepaper

Even under conservative assumptions, the total annual cost of such an integration—licensing, support, and validation—would lie in the **low-to-mid eight-figure range** at most. This represents a **fraction of one percent** of annual R&D spend for a leading semiconductor manufacturer and a negligible share relative to multi-billion-dollar ecosystem investments already underway.

### Lift-Up Effect on Existing Investments

The most important economic implication is that entropy-geometry transport does not compete with existing investments in GPU acceleration, AI surrogates, or digital twins; it **amplifies their value**.

Improved physics at the operator level yields:

- Higher-fidelity solver outputs at fixed resolution,
- More accurate training targets for surrogate and AI-assisted models, and
- Reduced need for conservative post-hoc correction and guard-banding.

In effect, the same accelerated infrastructure and surrogate pipelines begin operating on **better physical truth**, compounding the return on prior capital and engineering investment. This lift-up effect is precisely what makes physics-level corrections economically attractive: they propagate through the entire toolchain rather than adding isolated point improvements.

### Translating Accuracy Gains into Economic Value

Although the global accuracy improvements observed in silicon are modest in aggregate (~1–2%), the localized gains at high-curvature hotspot boundaries are substantial and strategically decisive. These regions govern throttling behavior, reliability margins, and yield confidence. Improving prediction accuracy where decisions are made has a disproportionate impact on system-level outcomes.

Industry-scale analysis shows that even **sub-percent improvements** in effective performance headroom, yield efficiency, or guard-band reduction—when applied across high-volume, high-value product lines—translate into **multi-billion-dollar annual economic value**. Importantly, this value is unlocked without increasing manufacturing complexity or process risk; it arises purely from better physical modeling of transport in real silicon.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

Technical Whitepaper

Technical Whitepaper

## Strategic Implication

The comparison is therefore asymmetrical. Semiconductor leaders are already spending billions to compute faster, approximate better, and scale existing physics. **AstraNomos Labs'** entropy-geometry framework offers a complementary path: **improve the physics those investments compute**, at a cost that is negligible by comparison.

In this light, adoption of entropy-geometry transport is best understood not as a speculative bet, but as a **high-leverage, low-risk enhancement** to an already massive simulation and digital-twin ecosystem. The empirical consistency across packed beds, rods, convective flows, and commercial silicon strongly suggests that this correction addresses a genuine structural incompleteness in classical transport theory—one whose economic implications scale with the industry itself.

## 10. Concluding Remarks

The microchip results presented in this work provide a concrete empirical demonstration of a deeper theoretical conclusion that spans classical heat transport and fluid mechanics: **Fourier diffusion and Navier–Stokes transport are incomplete when applied to strongly structured, high-curvature regimes**. While these equations remain valid and effective in smooth, homogeneous settings, they systematically misrepresent transport behavior at sharp interfaces, hotspot boundaries, and geometry-constrained regions—the very regimes that dominate real-world performance, reliability, and risk. The silicon data confirms that the breakdown is not numerical or incidental, but structural: classical operators assume unconstrained local entropy flow, even where geometry and material structure prohibit it.

By introducing curvature-aware, entropy-geometry transport, this work demonstrates that the failure mode of classical diffusion is both identifiable and correctable without abandoning foundational physics. The entropy-geometry model reduces exactly to Fourier and Navier-Stokes behavior in smooth regimes, preserving backward compatibility and interpretability, while selectively suppressing diffusion at high-curvature interfaces where classical assumptions break down. The microchip results show modest global accuracy improvements but **large, order-of-magnitude gains precisely at hotspot boundaries**, confirming that the governing issue is not average temperature prediction, but boundary persistence, gradient fidelity, and localized structure—exactly the domains where performance throttling, reliability risk, and conservative guard-banding originate.

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems

Technical Whitepaper

 Technical Whitepaper

This empirical behavior directly aligns with the accounting principle formalized in our Lyapunov–Perelman Navier–Stokes work: **structure cannot grow without payment**. In fluid dynamics, unchecked curvature amplification would imply singular behavior unless constrained by entropy and dissipation. In thermal transport, unconstrained diffusion artificially erases structure, producing over-smoothing and boundary decay that do not occur in physical systems. The entropy-geometry operator enforces the same conservation logic in transport: diffusion is permitted where geometry allows it and restrained where structure must persist. The silicon results thus serve as an observable, engineering-scale projection of the same principle that resolves theoretical inconsistencies in classical transport theory.

From a semiconductor perspective, the implications are immediate and material. Hotspot boundaries, not global averages, govern throttling behavior, electromigration risk, time-dependent dielectric breakdown, and lifetime guarantees. Errors of only a few degrees at these interfaces drive conservative design margins that reduce sustained performance and inflate cooling and packaging cost. By reducing prediction error by tens of percent in these critical zones—while remaining neutral elsewhere—the entropy-geometry model enables tighter guard-bands, more confident DVFS policies, and better-informed reliability decisions. Even small percentage improvements in effective performance headroom or yield efficiency scale into substantial economic value when applied across high-volume CPU, GPU, and accelerator product lines.

Crucially, these gains are achieved **without resorting to black-box correction, brute-force meshing, or ad hoc calibration**. The entropy-geometry formulation introduces no new state variables, no stochastic heuristics, and no discontinuous thresholds. Nonlinearity enters only through geometry, and only where demanded by the field itself. This makes the approach stable, interpretable, and compatible with existing PDE solvers, GPU-accelerated workflows, and AI-based surrogates. In contrast to approaches that merely accelerate or approximate existing physics, entropy-geometry transport represents a **constitutive-law upgrade**—improving the physics that all downstream tooling computes.

This distinction defines the core differentiation of the **AstraNomos digital twin platform PRISM**. **AstraNomos** does not compete on simulation speed alone, nor does it position itself as a surrogate replacement for established EDA or multi-physics tools. Instead, it provides a physics-level enhancement that **lifts the fidelity of every layer above it**—from solvers and surrogates to runtime optimization and operational decision systems. Because the correction collapses to classical behavior in smooth regimes, adoption can be incremental and low-risk, activating only where structural complexity justifies it. This makes **AstraNomos’** technology not disruptive in the operational sense, but transformative in its leverage.

While microchip thermal modeling provides a compelling and economically significant validation, the underlying framework is fundamentally cross-domain. Any system governed by transport under geometric constraint—packed beds, porous media, boundary-layer flows, batteries, reactors, additive manufacturing, climate-scale diffusion, infrastructure resilience, or biological transport—exhibits the same failure modes when modeled with homogeneous operators. **AstraNomos entropy-geometry digital twin PRISM** generalizes across these domains because it responds to structure rather than scale. The same

 ASTRANOMOS PRISM™

Deterministic Simulation Framework for Engineering Systems  
Technical Whitepaper

## Technical Whitepaper

mechanism that preserves hotspot boundaries in silicon preserves flow channels in granular reactors, thermal fronts in materials processing, and constrained diffusion in environmental and energy systems.

Taken together, this work establishes a unified conclusion across theory, experiment, and application: ***classical transport laws are not wrong, but incomplete***, and their incompleteness becomes economically decisive in structured regimes. The entropy-geometry framework resolves this gap by embedding geometric awareness directly into the operator, enforcing physically consistent entropy flow without sacrificing stability or compatibility. For **AstraNomos**, this represents a durable technological foundation—one that improves accuracy where it matters most, amplifies the value of existing simulation investments, and enables a new class of high-fidelity digital twins across industries. For practitioners, it offers a rare combination: deeper physical correctness, measurable economic impact, and a clear, low-friction path to adoption.

Please note:

This document and the technologies described herein are the intellectual property of AstraNomos Digital Twins Inc. (“Astronomos”). The PRISM™ platform, including its underlying deterministic simulation engine, structured transport models, algorithms, software implementations, and associated methodologies, constitutes proprietary technology developed by AstraNomos. The theoretical framework, modeling approaches, and analytical methods described in this whitepaper, including the geometry-aware deterministic transport framework and related operator-based formulations, are part of the proprietary intellectual property of AstraNomos and are protected under applicable copyright, trade secret, and intellectual property laws. No portion of this document may be reproduced, distributed, or used for commercial purposes without the prior written permission of AstraNomos Digital Twins Inc. The PRISM™ platform name and associated branding are trademarks of AstraNomos.