Article

Extreme Geodesics in Computational Spacetime: Novel Algorithmic Trajectories Beyond Classical Complexity Bounds

Michael Rey [*]

Article

Extreme Geodesics in Computational Spacetime: Novel Algorithmic Trajectories Beyond Classical Complexity Bounds

Michael Rey

Octonion Group, Hong Kong; contact@octoniongroup.com

Abstract

Building upon the foundational framework of computational relativity, this work proposes twelve novel geodesic archetypes that represent extreme trajectories through algorithmic spacetime. These geodesics suggest new approaches to complexity analysis by exploiting previously unexplored dimensions of computational resources—including energy reversibility, quantum coherence maintenance, temporal periodicity, and cross-platform optimization principles. We conjecture that several of these archetypes, particularly Floquet scheduling, Landauer-null trajectories, and hybrid optimization geodesics, may offer significant theoretical insights comparable to Williams' recent breakthrough in space-time trade-offs. Through mathematical formulations and comprehensive visualizations, we establish theoretical frameworks for these extreme geodesics and propose universal principles for algorithm design that move beyond problem-specific optimizations to reveal potential geometric structures in computational spacetime. The theoretical framework is supported by concrete applications spanning matrix multiplication, streaming algorithms, quantum computing, and energy-efficient computation, with each archetype providing testable hypotheses that could advance our understanding of algorithmic efficiency.

Keywords: extreme geodesics; computational spacetime; algorithmic complexity; time-space trade-off; energy reversibility; quantum coherence; Floquet computing; optimization principles; Williams breakthrough

---

1. Introduction

The landscape of computational complexity theory has been fundamentally altered by recent breakthroughs that challenge our classical understanding of resource trade-offs [1]. Most notably, Williams' 2025 result "Simulating Time With Square-Root Space" has established a new paradigm where space and time exhibit deep geometric relationships that can be exploited for algorithmic advantage [2–5]. This breakthrough represents more than an incremental improvement—it reveals that our traditional view of computational resources as separate, independent measures may fundamentally miss important relationships in algorithmic efficiency.

Note on Terminology: Throughout this work, we use geometric terminology such as "geodesics," "null cone," and "spacetime" as computational-geometry analogues to describe algorithmic trajectories through resource space. These are mathematical metaphors for optimization in multi-dimensional resource spaces, not claims about physical spacetime or general relativity.

The framework of computational relativity, introduced by Rey [6], provides a mathematical foundation for understanding these phenomena through the lens of spacetime geodesics. In this geometric theory,

algorithms become trajectories through a five-dimensional manifold (S, T, H, E, C) representing space, time, I/O bandwidth, energy, and quantum coherence respectively. Classical complexity results emerge naturally as specific geodesic types, while the metric structure of this manifold encodes the fundamental trade-offs between computational resources.

However, the exploration of this computational spacetime has only begun to scratch the surface of possible algorithmic trajectories. While previous work has focused on recovering classical results and establishing the basic geometric framework, vast regions of the manifold remain unexplored. These unexplored territories correspond to extreme geodesics—algorithmic trajectories that exploit novel combinations of resources in ways that classical complexity theory has not yet systematically investigated.

This work proposes twelve such extreme geodesic archetypes, each representing a fundamentally different approach to navigating computational spacetime. These are not merely theoretical curiosities but represent concrete algorithmic strategies with the potential to achieve significant improvements in specific domains. The archetypes span a spectrum from energy-reversible computation that approaches thermodynamic limits, to quantum coherence-maintenance strategies, to periodic scheduling that exploits temporal structure as a computational resource.

1.1. Why These Geodesics Matter

The significance of these extreme geodesics extends beyond their individual contributions. They represent a systematic methodology for algorithm discovery based on geometric principles rather than problem-specific insights. By systematically exploring the structure of computational space-time, we can identify potential universal principles that may apply across diverse computational domains. This geometric approach has proven valuable in Williams' breakthrough, which emerged from understanding fundamental geometric relationships in space-time trade-offs.

Our contributions are organized around three central themes. First, we establish mathematical frameworks for extreme geodesic analysis, providing rigorous formulations for each archetype that demonstrate their theoretical foundations. Second, we analyze the potential significance of these geodesics—their capacity to provide new insights into complexity theory. Third, we provide concrete algorithmic instantiations and testable hypotheses that enable experimental validation of the theoretical framework.

The twelve extreme geodesic archetypes we propose are:

Energy-Reversible Geodesics: Landauer-null trajectories that approach thermodynamic efficiency limits through reversible computation, potentially establishing energy as a fundamental complexity measure alongside space and time.

Coherence-Maintenance Geodesics: Quantum trajectories that maintain coherence above critical thresholds through geometric control, potentially enabling quantum advantages in previously classical domains.

Temporal Geodesics: Floquet schedules that exploit periodic control to potentially achieve better cycle-averaged metrics than static schedules, introducing temporal structure as an algorithmic resource.

Hybrid Geodesics: Optimization principles that govern transitions between computational platforms (CPU, GPU, quantum, photonic), proposing universal principles for heterogeneous computing.

Information-Theoretic Geodesics: Advice-augmented states that decouple computational complexity from verification complexity through succinct certificates, potentially enabling ecosystem-scale optimization.

Streaming Geodesics: Ultra-streaming approaches that explore whether limited-pass streaming can yield nontrivial improvements within known lower bounds, investigating the boundaries of streaming computation while respecting established $\Omega(n)$ space lower bounds for exact connectivity and related problems.

Stochastic Geodesics: Entropy-injection trajectories that use structured randomness to potentially improve convergence, establishing thermodynamic principles for algorithmic optimization.

Adaptive Geodesics: Curvature-following trajectories that adjust granularity based on local geometric properties, enabling real-time optimization of computational paths.

Redundancy Geodesics: Latency-minimizing trajectories that trade space and energy for improved tail performance, optimizing for modern distributed computing requirements.

I/O Geodesics: Bursty trajectories that may achieve global optimality through locally suboptimal I/O patterns, challenging smooth optimization assumptions.

Compression Geodesics: Holographic trajectories that project computation onto lower-dimensional manifolds while preserving essential structure.

Verification Geodesics: Proof-carrying trajectories that amortize computational cost across multiple consumers through cryptographic certificates [7].

Each of these archetypes represents a potential direction for algorithm design research. The Floquet schedules, for instance, suggest that temporal periodicity could be as fundamental to algorithmic efficiency as the spatial locality that drives cache optimization. The optimization principles propose systematic approaches for hybrid computing that could transform how we design algorithms for heterogeneous platforms. The Landauer-null geodesics indicate that energy efficiency might not require sacrificing computational performance, potentially advancing green computing research.

The theoretical framework we develop provides not only mathematical descriptions of these geodesics but also concrete criteria for identifying when they apply and how to implement them. We establish significance metrics that quantify each archetype's potential to advance complexity theory, drawing parallels to Williams' breakthrough to contextualize their importance. Through comprehensive visualizations of computational spacetime, we make the geometric intuitions accessible while maintaining mathematical rigor.

The remainder of this paper develops these ideas systematically. Section 2 establishes the mathematical framework for extreme geodesic analysis, extending the basic computational relativity formalism to handle the novel resource combinations these archetypes exploit. Section 3 provides detailed formulations for each of the twelve extreme geodesic archetypes, including their mathematical descriptions, implementation strategies, and theoretical properties. Section 4 analyzes the potential significance of these geodesics, establishing criteria for identifying which archetypes have the greatest potential for advancing complexity theory. Section 5 demonstrates concrete algorithmic applications across diverse computational domains, from matrix multiplication to quantum computing to distributed systems. Section 6 discusses the broader implications for complexity theory and algorithm design, while Section 7 outlines future research directions and open problems.

The ultimate goal of this work is not merely to catalog novel algorithmic techniques but to establish a systematic methodology for algorithm discovery based on geometric principles. By understanding the structure of computational spacetime and systematically exploring its extreme regions, we can identify potential universal principles that transcend specific problem domains. This geometric approach has already proven its value through Williams' breakthrough, and we believe the extreme geodesics introduced here represent promising directions for continued research in this geometric approach to computational complexity theory.

2. Mathematical Framework for Extreme Geodesic Analysis

The analysis of extreme geodesics requires extending the basic computational relativity formalism to handle novel resource combinations and geometric structures that classical complexity theory has not systematically explored. This section establishes the mathematical foundations necessary for rigorous treatment of the twelve archetypes we propose.

2.1. Extended Spacetime Manifold

We model computational processes as trajectories through a five-dimensional manifold $\mathcal{M} = (S, T, H, E, C)$ where each coordinate represents a fundamental computational resource:

$$S(\tau) : \text{Space complexity (memory usage)} \tag{1}$$
$$T(\tau) : \text{Time complexity (computational steps)} \tag{2}$$
$$H(\tau) : \text{I/O bandwidth (data movement rate)} \tag{3}$$
$$E(\tau) : \text{Energy consumption rate} \tag{4}$$
$$C(\tau) : \text{Quantum coherence level} \tag{5}$$

The parameter τ represents computational time, distinct from the complexity measure $T(\tau)$ which counts algorithmic steps. This distinction becomes important for geodesics that exploit temporal structure.

The metric tensor $g_{\mu\nu}(\mathbf{x})$ encodes the local cost structure of resource trade-offs:

$$g_{\mu\nu} = \begin{pmatrix} w_S(\mathbf{x}) & \gamma_{ST} & \gamma_{SH} & \gamma_{SE} & \gamma_{SC} \\ \gamma_{ST} & w_T(\mathbf{x}) & \gamma_{TH} & \gamma_{TE} & \gamma_{TC} \\ \gamma_{SH} & \gamma_{TH} & w_H(\mathbf{x}) & \gamma_{HE} & \gamma_{HC} \\ \gamma_{SE} & \gamma_{TE} & \gamma_{HE} & w_E(\mathbf{x}) & \gamma_{EC} \\ \gamma_{SC} & \gamma_{TC} & \gamma_{HC} & \gamma_{EC} & w_C(\mathbf{x}) \end{pmatrix}$$

where $w_\mu(\mathbf{x})$ represents the local cost of resource μ and $\gamma_{\mu\nu}$ captures coupling between resources. For extreme geodesics, these coupling terms become essential as they encode the novel resource interactions that may enable improved performance.

2.2. Generalized Action Principle

The action functional for computational trajectories takes the form:

$$\mathcal{A}[\mathbf{x}] = \int_{\tau_1}^{\tau_2} L(\mathbf{x}, \dot{\mathbf{x}}, \ddot{\mathbf{x}}, \tau) \, d\tau$$

where the Lagrangian L may depend on higher-order derivatives to capture the temporal structure exploited by extreme geodesics. The standard geodesic Lagrangian:

$$L_{\text{standard}} = \frac{1}{2} g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau}$$

can be extended for extreme geodesics to include:

Temporal Structure Terms:

$$L_{\text{temporal}} = \sum_{n=1}^{N} \alpha_n \cos\left(\frac{2\pi n \tau}{\Theta}\right) g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau}$$

for Floquet geodesics with period Θ.

Curvature Coupling Terms:

$$L_{\text{curvature}} = \kappa R(\mathbf{x}) + \lambda R_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau}$$

for adaptive geodesics that respond to local spacetime curvature.

Constraint Terms:

$$L_{\text{constraint}} = \sum_i \lambda_i \phi_i(\mathbf{x}, \dot{\mathbf{x}})$$

where ϕ_i represent physical constraints such as energy conservation, coherence bounds, or information-theoretic limits.

2.3. Extreme Geodesic Classification

We classify extreme geodesics based on their geometric properties and the novel resource combinations they exploit:

Type I: Resource Substitution Geodesics These geodesics may achieve improved performance by substituting abundant resources for scarce ones in ways that classical analysis has not systematically explored. Williams' breakthrough exemplifies this type, showing that space can substitute for time with surprising efficiency. Our Landauer-null and coherence-maintenance geodesics extend this principle to energy and quantum resources respectively.

Type II: Temporal Structure Geodesics These geodesics exploit periodic or quasi-periodic control to potentially achieve better time-averaged performance than static strategies. The key insight is that temporal structure itself may become a computational resource. Floquet schedules and time-crystal geodesics exemplify this type.

Type III: Hybrid Platform Geodesics These geodesics optimize transitions between different computational platforms (classical, quantum, photonic, neuromorphic) using optimization principles analogous to refraction laws. The potential lies in establishing systematic principles for heterogeneous computing that transcend platform-specific optimizations.

Type IV: Information-Theoretic Geodesics These geodesics exploit the gap between computational complexity and verification complexity, using cryptographic techniques to decouple producer costs from consumer costs. The potential comes from enabling ecosystem-scale optimization where computational work can be amortized across many consumers.

3. Twelve Extreme Geodesic Archetypes

This section presents detailed mathematical formulations and analysis for twelve extreme geodesic archetypes that represent potentially new approaches to navigating computational spacetime. Each archetype exploits previously unexplored combinations of computational resources to potentially achieve improved performance that classical complexity analysis has not systematically investigated.

3.1. Type I: Landauer-Null Geodesics (Energy-Reversible Trajectories)

The Landauer-null geodesic represents perhaps the most physically grounded of our extreme archetypes, directly building upon the fundamental thermodynamic limits of computation established by Landauer's principle [8].

Conjecture 1 (Landauer-Null Optimality). *For algorithms that can be made reversible, there exist computational trajectories that approach the Landauer limit for energy dissipation while maintaining polynomial time complexity.*

Theoretical Foundation: Landauer's principle establishes that irreversible bit operations require minimum energy dissipation $\Delta E \geq k_B T \ln 2$ per erased bit. However, this bound applies only to irreversible operations. Reversible computation, in principle, can approach zero energy dissipation at the cost of increased space and time complexity.

Mathematical Formulation: The Landauer-null geodesic minimizes energy dissipation through the Lagrangian:

$$L_{\text{Landauer}} = \alpha \left(\frac{dS}{d\tau}\right)^2 + \beta \left(\frac{dH}{d\tau}\right)^2 + \gamma \left(\frac{dT}{d\tau}\right)^2 + \varepsilon \left(\frac{dE}{d\tau}\right)^2$$

with $\varepsilon \ll \alpha, \beta, \gamma$ to heavily penalize energy changes. The constraint from Landauer's principle becomes:

$$\frac{dE}{d\tau} \geq k_B T_{\text{phys}} \ln 2 \cdot \frac{dN_{\text{erases}}}{d\tau}$$

where $N_{\text{erases}}(\tau)$ counts irreversible bit operations up to time τ.

Implementation Strategy: Landauer-null geodesics can be implemented through Bennett-style reversible computation [9] with systematic checkpoint recycling. The space-time trade-off becomes:

$$S \cdot T \geq \text{const} \cdot N_{\text{ops}}^2$$

due to checkpoint storage requirements, but energy consumption approaches the fundamental thermodynamic limit.

Why This Matters: This archetype has potential significance because it could elevate energy to a fundamental complexity measure alongside space and time. If practical reversible algorithms can be developed for broad problem classes, it would provide a systematic framework for energy-aware algorithm design, which is increasingly important as energy constraints become critical in computing.

3.2. Type II: Floquet Geodesics (Periodic Scheduling)

Floquet geodesics exploit periodic control to potentially achieve better cycle-averaged performance than static schedules, introducing temporal structure as a fundamental algorithmic resource.

Conjecture 2 (Floquet Optimality). *For computational problems with periodic cost structures, there exist periodic scheduling strategies that achieve strictly better cycle-averaged performance than any static schedule under identical peak resource constraints.*

Theoretical Foundation: Floquet theory, originally developed for periodic differential equations, reveals that periodic driving can stabilize otherwise unstable systems and create effective dynamics

that differ qualitatively from static systems. In computational contexts, this suggests that periodic scheduling of resources might achieve better average performance than time-invariant strategies.

Mathematical Formulation: The Floquet Lagrangian has periodic structure:

$$L_{\text{Floquet}}(\tau) = \sum_{n=-\infty}^{\infty} L_n e^{in\omega\tau}$$

where L_n are Fourier coefficients and $\omega = 2\pi/\Theta$ with period Θ. The effective metric is time-averaged:

$$g_{\mu\nu}^* = \frac{1}{\Theta} \int_0^{\Theta} g_{\mu\nu}(\tau)d\tau$$

Geodesic Solution: The Floquet geodesic equation becomes:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\nu\rho}^{*\ \mu} \frac{dx^\nu}{d\tau} \frac{dx^\rho}{d\tau} = 0$$

where Γ^* are Christoffel symbols computed from the effective metric g^*.

Conjecture 3 (Effective Geodesic Length). *Under suitable regularity and ergodicity conditions on the periodic cost tensor (including bounded variation and appropriate averaging assumptions), the effective geodesic length may satisfy:*

$$\int_0^T \sqrt{g_{\mu\nu}^* \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau}} d\tau \le \min_{static} \int_0^T \sqrt{g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau}} d\tau$$

with equality only in degenerate cases. This conjecture could be validated through averaging lemmas for periodic Lagrangians and experimental measurement of cycle-averaged performance.

Implementation Strategy: Practical Floquet schedules alternate between compute phases, I/O phases, and refresh phases with carefully tuned periods. For matrix multiplication, this might involve periodic tiling with compute-refresh-I/O cycles that could achieve better cache efficiency than static blocking strategies.

Why This Matters: This archetype has potential significance because it introduces temporal structure as a fundamental algorithmic resource. If Floquet principles can be validated across algorithm classes, it would suggest that periodicity should be a first-class optimization target in algorithm design, potentially revolutionizing how we think about scheduling and resource management.

3.3. Type III: Hybrid Optimization Geodesics

Hybrid optimization geodesics propose systematic principles for optimal transitions between computational platforms, analogous to optimization principles in physics.

Conjecture 4 (Computational Optimization Principle). *For hybrid computational systems, there exist universal optimization principles governing platform transitions that can be expressed as variational conditions analogous to Snell's law in optics.*

Theoretical Foundation: Different computational platforms (CPU, GPU, quantum, photonic, neuromorphic) have different cost structures for computational resources. The hypothesis is that optimal computational trajectories should follow systematic optimization principles when transitioning between platforms, similar to how physical systems follow variational principles.

Mathematical Formulation: The proposed computational optimization principle states:

$$\frac{\sin\theta_A}{\sin\theta_B} = \frac{\eta_B}{\eta_A}$$

where $\eta = \sqrt{w_T/w_S}$ is the computational refractive index and θ is the angle between the trajectory and the normal to the platform boundary.

Boundary conditions require continuity of trajectory and normal component of computational flux:

$$\mathbf{x}_A(\tau_b) = \mathbf{x}_B(\tau_b) \tag{6}$$

$$\eta_A \frac{d\mathbf{x}_A}{d\tau} \cdot \hat{\mathbf{n}} = \eta_B \frac{d\mathbf{x}_B}{d\tau} \cdot \hat{\mathbf{n}} \tag{7}$$

Implementation Strategy: Practical implementation requires measuring platform-specific cost tensors to determine refractive indices, then applying the optimization principle to choose transfer points for hybrid algorithms.

Why This Matters: This archetype has potential significance because it proposes systematic principles for heterogeneous computing. If optimization laws can be validated across diverse platform combinations, it could transform hybrid algorithm design from ad-hoc heuristics to principled optimization based on geometric principles.

3.4. Type IV: Coherence-Maintenance Geodesics

Coherence-maintenance geodesics exploit quantum coherence as a computational resource, maintaining coherence above critical thresholds through geometric control.

Conjecture 5 (Coherence-Maintenance Optimality). *For quantum algorithms, there exist computational trajectories that maintain coherence above critical thresholds C_{crit} while achieving polynomial improvements in specific problem classes.*

Mathematical Formulation: The coherence-maintenance Lagrangian includes a penalty term for coherence loss:

$$L_{\text{coherence}} = \alpha\left(\frac{dS}{d\tau}\right)^2 + \beta\left(\frac{dT}{d\tau}\right)^2 + \gamma\left(\frac{dC}{d\tau}\right)^2 + \lambda\Theta(C_{\text{crit}} - C)$$

where Θ is the Heaviside function that heavily penalizes trajectories below the coherence threshold.

Geodesic Equation:

$$\frac{d^2C}{d\tau^2} + \Gamma^C_{ST}\frac{dS}{d\tau}\frac{dT}{d\tau} = -\frac{\lambda}{\gamma}\delta(C - C_{\text{crit}})$$

Why This Matters: This archetype could enable quantum advantages in previously classical domains by systematically maintaining coherence through geometric optimization principles.

3.5. Type V: Advice-Augmented Geodesics

Advice-augmented geodesics exploit the gap between computational complexity and verification complexity using succinct certificates.

Conjecture 6 (Advice-Augmented Efficiency). *For problems with expensive computation but efficient verification, there exist advice-augmented trajectories that achieve ecosystem-scale optimization through cryptographic amortization.*

Mathematical Formulation: The advice-augmented Lagrangian separates producer and consumer costs:

$$L_{\text{advice}} = \alpha_p \left(\frac{dS_p}{d\tau} \right)^2 + \beta_p \left(\frac{dT_p}{d\tau} \right)^2 + \sum_{i=1}^{N} \left[\alpha_c \left(\frac{dS_{c,i}}{d\tau} \right)^2 + \beta_c \left(\frac{dT_{c,i}}{d\tau} \right)^2 \right]$$

where subscript p denotes producer and c, i denotes the i-th consumer.

Geodesic Equation:

$$\frac{d^2 S_p}{d\tau^2} + \Gamma^{S_p}_{T_p T_p} \left(\frac{dT_p}{d\tau} \right)^2 = \frac{1}{N} \sum_{i=1}^{N} \frac{\partial V_i}{\partial S_p}$$

where V_i represents the verification cost for consumer i.

Why This Matters: This archetype could enable distributed computing scenarios where expensive computations are amortized across many consumers through cryptographic certificates.

3.6. Type VI: Ultra-Streaming Geodesics

Ultra-streaming geodesics explore whether limited-pass streaming can yield nontrivial improvements within known lower bounds.

Conjecture 7 (Ultra-Streaming Efficiency). *For streaming problems, there exist multi-pass trajectories that achieve improved space-pass trade-offs while respecting established lower bounds for exact computation.*

Mathematical Formulation: The streaming Lagrangian includes pass-counting constraints:

$$L_{\text{stream}} = \alpha \left(\frac{dS}{d\tau} \right)^2 + \beta \left(\frac{dH}{d\tau} \right)^2 + \gamma P(\tau)^2$$

where $P(\tau)$ counts the number of passes through the data stream.

Geodesic Equation:

$$\frac{d^2 S}{d\tau^2} + \Gamma^S_{HH} \left(\frac{dH}{d\tau} \right)^2 = -\frac{2\gamma}{\alpha} P(\tau) \frac{dP}{d\tau}$$

Why This Matters: This archetype could advance streaming algorithms by systematically exploring the space-pass trade-off landscape within theoretical constraints.

3.7. Type VII: Entropy-Injection Geodesics

Entropy-injection geodesics use structured randomness to potentially improve convergence through thermodynamic principles.

Conjecture 8 (Entropy-Injection Optimality). *For optimization problems, there exist stochastic trajectories that achieve better convergence rates by injecting entropy at optimal rates determined by local curvature.*

Mathematical Formulation: The entropy-injection Lagrangian includes stochastic terms:

$$L_{\text{entropy}} = \frac{1}{2} g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} + \kappa \sum_i \sigma_i^2(\tau) R_i(\mathbf{x})$$

where $\sigma_i(\tau)$ represents noise injection rates and R_i are curvature components.

Geodesic Equation:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma^\mu_{\nu\rho} \frac{dx^\nu}{d\tau} \frac{dx^\rho}{d\tau} = \kappa \sum_i \frac{\partial}{\partial x^\mu} \left[\sigma_i^2(\tau) R_i(\mathbf{x}) \right]$$

Why This Matters: This archetype could establish thermodynamic principles for algorithmic optimization, connecting statistical mechanics to computational efficiency.

3.8. Type VIII: Adaptive-Granularity Geodesics

Adaptive-granularity geodesics adjust computational granularity based on local geometric properties of the problem space.

Conjecture 9 (Adaptive-Granularity Optimality). *For hierarchical computational problems, there exist adaptive trajectories that achieve better performance by dynamically adjusting granularity based on local curvature and resource availability.*

Mathematical Formulation: The adaptive Lagrangian includes granularity-dependent terms:

$$L_{\text{adaptive}} = \frac{1}{2} g_{\mu\nu}(\mathbf{x}, \delta(\tau)) \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau}$$

where $\delta(\tau)$ represents the computational granularity parameter.

Geodesic Equation:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma^\mu_{\nu\rho} \frac{dx^\nu}{d\tau} \frac{dx^\rho}{d\tau} = \frac{1}{2} \frac{\partial g_{\alpha\beta}}{\partial \delta} \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} \frac{d\delta}{d\tau} \frac{\partial x^\mu}{\partial \delta}$$

Why This Matters: This archetype could enable real-time optimization of computational paths by adapting to local problem structure and resource constraints.

3.9. Type IX: Bursty-I/O Geodesics

Bursty-I/O geodesics achieve global optimality through locally suboptimal I/O patterns that challenge smooth optimization assumptions.

Conjecture 10 (Bursty-I/O Optimality). *For I/O-bound computations, there exist bursty trajectories that achieve better global performance through locally suboptimal I/O patterns that exploit temporal correlations in storage systems.*

Mathematical Formulation: The bursty Lagrangian includes non-smooth I/O terms:

$$L_{\text{bursty}} = \frac{1}{2} g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} + \lambda \sum_k \left| \frac{dH}{d\tau} - H_k \right|$$

where H_k are preferred I/O burst rates.

Geodesic Equation:

$$\frac{d^2 H}{d\tau^2} + \Gamma^H_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} = \lambda \sum_k \text{sgn}\left(\frac{dH}{d\tau} - H_k\right) \delta\left(\frac{dH}{d\tau} - H_k\right)$$

Why This Matters: This archetype could optimize for modern storage systems where bursty access patterns can be more efficient than smooth I/O due to hardware characteristics.

3.10. Type X: Holographic-Compression Geodesics

Holographic-compression geodesics project computation onto lower-dimensional manifolds while preserving essential algorithmic structure.

Conjecture 11 (Holographic-Compression Efficiency). *For high-dimensional computational problems, there exist holographic trajectories that achieve comparable results by projecting onto lower-dimensional manifolds with controlled information loss.*

Mathematical Formulation: The holographic Lagrangian includes dimensional projection terms:

$$L_{\text{holo}} = \frac{1}{2} \Pi^*_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} + \epsilon \|\mathbf{x} - \Pi(\mathbf{x})\|^2$$

where Π is the projection operator and Π^* is the induced metric on the lower-dimensional manifold.

Geodesic Equation:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma^{*\mu}_{\nu\rho} \frac{dx^\nu}{d\tau} \frac{dx^\rho}{d\tau} = -\frac{2\epsilon}{\alpha} \frac{\partial}{\partial x^\mu} \|\mathbf{x} - \Pi(\mathbf{x})\|^2$$

Why This Matters: This archetype could enable efficient computation on high-dimensional problems by identifying essential lower-dimensional structure.

3.11. Type XI: Redundancy-for-Latency Geodesics

Redundancy-for-latency geodesics trade space and energy for improved tail performance in distributed systems.

Conjecture 12 (Redundancy-for-Latency Optimality). *For latency-critical distributed computations, there exist redundant trajectories that achieve better tail latency performance by strategic over-provisioning of computational resources.*

Mathematical Formulation: The redundancy Lagrangian includes replication terms:

$$L_{\text{redundancy}} = \frac{1}{2} g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} + \eta R(\tau) \left[\alpha \frac{dS}{d\tau} + \beta \frac{dE}{d\tau}\right]$$

where $R(\tau)$ is the replication factor.

Geodesic Equation:

$$\frac{d^2 S}{d\tau^2} + \Gamma^S_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} = -\eta\alpha \frac{dR}{d\tau}$$

Why This Matters: This archetype could optimize modern distributed systems where tail latency is critical and resources can be traded for performance guarantees.

3.12. Type XII: Proof-Carrying Geodesics

Proof-carrying geodesics amortize computational cost across multiple consumers through cryptographic certificates.

Conjecture 13 (Proof-Carrying Efficiency)**.** *For computations with many consumers, there exist proof-carrying trajectories that achieve better amortized complexity by embedding verification certificates directly in computational results.*

Mathematical Formulation: The proof-carrying Lagrangian includes certificate generation costs:

$$L_{\text{proof}} = L_{\text{compute}} + L_{\text{prove}} + N \cdot L_{\text{verify}}$$

where N is the number of consumers.

Geodesic Equation:

$$\frac{d^2 T_{\text{total}}}{d\tau^2} = \frac{d^2 T_{\text{compute}}}{d\tau^2} + \frac{d^2 T_{\text{prove}}}{d\tau^2} + N \frac{d^2 T_{\text{verify}}}{d\tau^2}$$

with the constraint that $T_{\text{verify}} \ll T_{\text{compute}}$ for large N.

Why This Matters: This archetype could enable ecosystem-scale optimization where computational work is efficiently shared across many consumers with cryptographic guarantees.

4. Why These Geodesics Matter

The significance of extreme geodesics lies not merely in their individual algorithmic contributions but in their potential to advance computational complexity theory. This section analyzes which archetypes possess the greatest potential for theoretical impact, drawing parallels to Williams' 2025 breakthrough.

4.1. The Williams Benchmark: Characteristics of Significant Results

Williams' breakthrough "Simulating Time With Square-Root Space" established new standards for significant results in complexity theory [2]. Rather than optimizing a specific algorithm, Williams revealed a universal principle: space can substitute for time with surprising efficiency across broad problem classes. This result shifted the frontier of achievable trade-offs rather than merely improving individual points on existing curves.

The key characteristics that made Williams' result significant were universality across computational domains, frontier impact that shifted entire Pareto frontiers, theoretical depth connecting to fundamental principles, and testable predictions that could be experimentally validated.

4.2. Potential Impact of Extreme Geodesics

Several of our extreme geodesic archetypes show particular promise for advancing complexity theory:

Floquet/Periodic Scheduling Geodesics could establish temporal structure as a fundamental algorithmic resource, comparable to how Williams established new relationships between space and time. Periodic scheduling could potentially apply to matrix operations, sorting algorithms, graph traversals, machine learning training, and database operations.

Hybrid Optimization Principles could promote ad-hoc offload heuristics to systematic optimization principles, establishing universal approaches for heterogeneous computing. The principle would poten-

tially be platform-agnostic, applying to CPU-GPU, classical-quantum, edge-cloud, and neuromorphic-digital scenarios.

Landauer-Null Trajectories could elevate energy to a fundamental complexity measure alongside space and time. This approach applies to linear algebra, cryptography, signal processing, and computations that can be made reversible, particularly relevant as energy constraints become critical in modern computing.

The remaining archetypes—advice-augmented geodesics, ultra-streaming approaches, coherence-maintenance geodesics, entropy-injection geodesics, adaptive-granularity geodesics, bursty-I/O geodesics, holographic-compression geodesics, and proof-carrying geodesics—represent important algorithmic research directions that could advance specific computational domains.

5. Applications and Algorithmic Design Blueprints

This section demonstrates potential applications of extreme geodesic principles across diverse computational domains, showing how the theoretical framework could translate into practical algorithmic research directions. We focus on the high-significance potential archetypes and provide detailed algorithmic design blueprints that can be experimentally validated.

5.1. Floquet Geodesics in Matrix Multiplication

Matrix multiplication serves as an ideal testbed for Floquet geodesics due to its well-understood complexity landscape [10–12] and the recent breakthroughs that have pushed the exponent to $\omega \leq 2.371552$ [13–15].

Classical Blocking Analysis: Standard blocked matrix multiplication uses static tile sizes determined by cache hierarchy. For matrices of size $n \times n$ with cache size M, optimal tile size is $b = \sqrt{M/3}$, yielding complexity $O(n^3/\sqrt{M})$ cache misses.

Floquet Schedule Design: We propose a periodic schedule with three phases:

Floquet Matrix Multiplication Schedule

1. **Compute Phase** ($\tau \in [0, \Theta/3)$): Standard GEMM operations with large tiles
2. **Refresh Phase** ($\tau \in [\Theta/3, 2\Theta/3)$): Cache line prefetching and data reorganization
3. **I/O Phase** ($\tau \in [2\Theta/3, \Theta)$): Coordinated memory hierarchy management

---

Testable Predictions:

- Floquet scheduling should achieve measurable improvement in cache efficiency over optimal static blocking
- Optimal period Θ should correlate with cache hierarchy timing characteristics
- Improvements should be larger on NUMA architectures with complex memory hierarchies

Implementation Strategy: Use hardware performance counters to measure cache miss rates and adjust period Θ dynamically. Implement using OpenMP with careful thread synchronization during phase transitions.

5.2. Hybrid Optimization in CPU-GPU Computing

Modern heterogeneous platforms require careful orchestration of computation between CPUs and GPUs. Current approaches use heuristics, but optimization principles could provide systematic approaches.

Cost Tensor Measurement: For a given computational kernel, measure the cost tensors on both platforms:

$$\text{CPU:} \quad g_{\text{ST}}^{\text{CPU}} = \frac{\text{Time per operation}}{\text{Memory per operation}} \tag{8}$$

$$\text{GPU:} \quad g_{\text{ST}}^{\text{GPU}} = \frac{\text{Time per operation}}{\text{Memory per operation}} \tag{9}$$

Optimization Principle Application: For a hybrid algorithm transitioning from CPU to GPU computation, the proposed optimal handoff condition is:

$$\frac{\sin \theta_{\text{CPU}}}{\sin \theta_{\text{GPU}}} = \frac{\eta_{\text{GPU}}}{\eta_{\text{CPU}}}$$

where θ represents the trajectory angle in space-time coordinates.

Experimental Validation: Implement for sparse matrices from SuiteSparse collection, measuring total solution time against heuristic handoff strategies. The optimization principle should be tested against current best practices.

5.3. Landauer-Null Geodesics in Energy-Aware Computing

Energy-aware computation is becoming increasingly important as energy constraints affect all levels of computing from mobile devices to data centers.

Reversible Algorithm Design: Cryptographic operations are candidates for Landauer-null geodesics due to their mathematical structure:

Energy-Aware Reversible Computation

1. **Forward Computation**: Standard algorithm execution with intermediate state storage
2. **Checkpoint Management**: Store minimal state information for reversibility
3. **Reverse Computation**: Uncompute intermediate states to recover space
4. **Energy Accounting**: Track bit erasures and energy dissipation

---

Theoretical Framework: Reversible algorithms trade time and space for energy efficiency. The trade-off relationship follows:

$$\text{Energy} \cdot \text{Time} \cdot \text{Space} \geq \text{const} \cdot \text{Problem Size}$$

Testable Predictions:

- Reversible algorithms should achieve measurable energy reductions
- Energy savings should come at predictable costs in time and space
- Approach to thermodynamic limits should be measurable with appropriate instrumentation

5.4. Advice-Augmented Geodesics in Distributed Computing

Distributed computing workloads often involve expensive computations followed by many verification or inference queries, making them candidates for advice-augmented optimization.

Succinct Verification Framework: Use cryptographic proof systems to create succinct certificates for computational results:

Advice-Augmented Distributed Pipeline

1. **Producer**: Perform expensive computation, generate succinct proof of correctness
2. **Certificate**: Cryptographic proof that computation meets specified criteria
3. **Consumers**: Verify computation quality efficiently, use results for downstream tasks
4. **Amortization**: Computation cost amortized across many consumers

Verification Complexity: Many zk proof systems achieve verification that is polylogarithmic in circuit size for specific models and constructions, where n is the computation size, compared to $O(n^k)$ for recomputation [16].

Testable Predictions:

- Verification time should scale logarithmically with computation size
- Amortization benefits should increase with number of consumers
- Certificate size should be succinct relative to computation complexity

5.5. Performance Predictions and Validation Metrics

For each application domain, we provide specific performance predictions that can be experimentally validated:

Floquet Matrix Multiplication:

- Measurable improvement in cache efficiency over static blocking
- Optimal period Θ correlates with cache hierarchy characteristics
- Larger improvements on complex memory hierarchies

Hybrid Optimization:

- Consistent improvement over heuristic handoff strategies
- Performance improvements correlate with platform cost differences
- Optimization indices correlate with hardware specifications

Landauer-Null Computing:

- Measurable reduction in energy dissipation
- Predictable increases in computation time and memory usage
- Approach to thermodynamic limits with appropriate instrumentation

Advice-Augmented Systems:

- Verification time scales logarithmically with computation size
- Amortization benefits increase with consumer count
- Certificate size remains succinct

These applications demonstrate that extreme geodesic principles could translate into concrete algorithmic research directions with measurable performance characteristics. The next section discusses broader implications for complexity theory and algorithm design.

6. Implications and Future Directions

The extreme geodesic archetypes introduced in this work represent more than algorithmic innovations—they constitute a systematic methodology for algorithm discovery based on geometric principles rather than problem-specific insights. The implications extend across multiple dimensions of computational theory and practice.

6.1. Theoretical Implications

Expansion of Complexity Theory: If validated, the high-significance geodesics could necessitate expansions to computational complexity theory. Energy, quantum coherence, temporal periodicity, and platform heterogeneity could join space and time as fundamental complexity measures. This expansion would parallel historical developments where new computational models (randomized, quantum, parallel) required new complexity classes and analysis techniques.

Geometric Foundations: The success of extreme geodesics would support geometric optimization as a methodology for algorithm design. Rather than optimizing individual resources separately, algorithm designers could work with unified geometric optimization in extended spacetime manifolds. This represents a shift toward more systematic approaches to computational complexity.

Predictive Algorithm Discovery: The geometric framework could enable prediction of algorithmic improvements before their explicit construction. By analyzing the curvature and topology of computational spacetime, researchers could identify promising regions for exploration, similar to how mathematical theories guide scientific discovery.

6.2. Open Problems and Research Directions

Geodesic Completeness: Are the twelve archetypes we have identified complete, or do additional extreme geodesic types exist? Systematic exploration of computational spacetime topology could reveal new archetype classes.

Computational Complexity of Geodesic Optimization: What is the computational complexity of finding optimal geodesics in extended spacetime manifolds? This meta-complexity question could determine the practical feasibility of geometric algorithm design.

Hardware-Geodesic Co-design: How should hardware architectures be designed to optimally support extreme geodesics? Floquet-optimized cache hierarchies, optimization-aware interconnects, and Landauer-null processing units represent new directions for computer architecture research.

Validation Roadmap: The validation of extreme geodesic theory requires coordinated experimental efforts across multiple domains, from proof-of-concept implementations to comprehensive benchmarking to theoretical validation.

7. Conclusion

This work has introduced twelve extreme geodesic archetypes that represent potentially new approaches to navigating computational spacetime. These archetypes suggest novel directions for complexity analysis by exploiting previously unexplored combinations of computational resources—energy reversibility, quantum coherence maintenance, temporal periodicity, cross-platform optimization principles, and information-theoretic decoupling.

Our analysis suggests that several archetypes, particularly Floquet scheduling, hybrid optimization principles, and Landauer-null trajectories, possess potential significance comparable to recent breakthroughs in complexity theory. These geodesics could advance complexity theory by introducing new resource dimensions, establishing systematic optimization principles, and enabling predictive algorithm discovery.

The theoretical framework we have developed provides mathematical foundations for extreme geodesic analysis, including generalized action principles, significance assessment metrics, and stability analysis. The concrete applications across matrix multiplication, cryptography, machine learning, graph analytics, and quantum computing demonstrate that these principles could translate into measurable algorithmic research directions.

The ultimate significance of extreme geodesics lies not in their individual contributions but in their collective suggestion of geometric structure underlying computational efficiency. By systematically exploring computational spacetime, we may identify universal principles that transcend specific problem domains and enable breakthrough algorithmic discoveries.

The geometric approach to computational complexity theory has shown its value through Williams' demonstration of fundamental space-time relationships. The extreme geodesics introduced here represent potential next steps in this geometric approach, suggesting that the full structure of computational spacetime contains unexplored territories with significant research potential.

As we continue to push the boundaries of computational efficiency in an era of energy constraints, quantum computing, and heterogeneous platforms, the geometric approach to algorithm design offers a systematic framework for navigating these challenges. The extreme geodesics we have identified provide concrete research directions for achieving improved performance, while the underlying geometric theory offers a methodology for discovering future breakthroughs.

The exploration of computational spacetime has only just begun. The twelve extreme geodesic archetypes presented here represent initial proposals for this research direction. As our understanding of computational geometry develops, we anticipate the discovery of additional archetypes and the development of increasingly sophisticated tools for geodesic optimization.

The future of algorithm design may lie not in optimizing individual resources but in understanding and exploiting the geometric structure of computational spacetime itself. The extreme geodesics introduced in this work provide a foundation for systematic exploration of this structure and establish a framework for geometrically-guided algorithmic research.

References

1. J. Hopcroft, W. Paul, and L. Valiant, "On Time Versus Space," *Journal of the ACM*, vol. 24, no. 2, pp. 332-337, 1977.
2. R. Williams, "Simulating Time With Square-Root Space," *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*, 2025. Available: https://acm-stoc.org/stoc2025/accepted-papers.html
3. S. Nadis, "For Algorithms, a Little Memory Outweighs a Lot of Time," *Quanta Magazine*, May 21, 2025. Available: https://www.quantamagazine.org/for-algorithms-a-little-memory-outweighs-a-lot-of-time-20250521/
4. "You Need Much Less Memory than Time," *Communications of the ACM Blog*, 2025. Available: https://cacm.acm.org/blogcacm/you-need-much-less-memory-than-time/
5. "For Algorithms, Memory Is a Far More Powerful Resource Than Time," *WIRED*, 2025. Available: https://www.wired.com/story/for-algorithms-a-little-memory-outweighs-a-lot-of-time/
6. M. Rey, "Computational Relativity: A Geometric Theory of Algorithmic Spacetime," *Preprints*, 2025. DOI: 10.20944/preprints202509.0015.v1. Available: https://www.preprints.org/manuscript/202509.0015/v1
7. S. A. Cook, "The Complexity of Theorem-Proving Procedures," *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing*, pp. 151-158, 1971.
8. R. Landauer, "Irreversibility and Heat Generation in the Computing Process," *IBM Journal of Research and Development*, vol. 5, no. 3, pp. 183-191, 1961.
9. C. H. Bennett, "Logical Reversibility of Computation," *IBM Journal of Research and Development*, vol. 17, no. 6, pp. 525-532, 1973.

10. V. Strassen, "Gaussian Elimination is Not Optimal," *Numerische Mathematik*, vol. 13, no. 4, pp. 354-356, 1969.

11. D. Coppersmith and S. Winograd, "Matrix Multiplication via Arithmetic Progressions," *Journal of Symbolic Computation*, vol. 9, no. 3, pp. 251-280, 1990.

12. F. Le Gall, "Powers of Tensors and Fast Matrix Multiplication," *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, pp. 296-303, 2014.

13. V. V. Williams, D. Xu, Y. Xu, and R. Zhou, "New Bounds for Matrix Multiplication: from Alpha to Omega," *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 3792-3835, 2024. arXiv:2307.07970. DOI: Available: https://epubs.siam.org/doi/10.1137/1.9781611977912.134

14. R. Duan, H. Wu, and R. Zhou, "Faster Matrix Multiplication via Asymmetric Hashing," *Proceedings of the 64th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 2129-2138, 2023. arXiv:2210.10173. Available: https://arxiv.org/abs/2210.10173

15. J. Alman and V. V. Williams, "A Refined Laser Method and Faster Matrix Multiplication," *TheoretiCS*, vol. 3, 2024. arXiv:2010.05846. Available: https://arxiv.org/abs/2010.05846

16. D. Kang et al., "A Survey of Zero-Knowledge Proof Based Verifiable Machine Learning," arXiv:2502.18535, 2024. Available: https://arxiv.org/abs/2502.18535