

Article

Not peer-reviewed version

P vs NP in Spacetime: Proper-Time Complexity, Curvature-Dependent Trade-Offs, and Conditional Separation Conjectures

[Michael Rey](#) *

Posted Date: 8 September 2025

doi: [10.20944/preprints202509.0628.v1](https://doi.org/10.20944/preprints202509.0628.v1)

Keywords: P vs NP; proper-time complexity; spacetime; general relativity; quantum speed limits; Bekenstein bound; black holes; computational complexity



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

P vs NP in Spacetime: Proper-Time Complexity, Curvature-Dependent Trade-offs, and Conditional Separation Conjectures

Michael Rey 

Octonion Group, Hong Kong; contact@octoniongroup.com

Abstract

Classical complexity theory studies the resources required by algorithms on abstract machines, measuring time by the number of elementary steps. In physical reality, computations are executed by devices embedded in spacetime, with resources bounded by energy, entropy, geometry, and causal structure. We develop a spacetime-aware framework for decision-problem complexity that measures cost in an observer's proper time and couples it to physically motivated bounds on space (memory), energy, and communication. Within this framework we define relativized complexity classes for polynomial-time and nondeterministic polynomial-time problems that depend on the background spacetime geometry and a resource-mapping function that translates physical resources to logical computational steps. We prove an explicit curvature-independent time-space trade-off inequality by combining the Bekenstein entropy bound with quantum speed limits, show an isometry covariance property of our definitions, and formulate two testable conjectures: (i) a Frame-Dependence Conjecture asserting that, for reasonable families of resource mappings, membership in our proper-time polynomial class can differ between non-coincident observers, and (ii) a Gravitational Acceleration Threshold identifying when polynomial-time solvability measured in a distant coordinate clock emerges from extreme redshift while remaining exponential in local proper time. We contrast these statements with classical complexity theory and with results on closed timelike curves and Malament-Hogarth spacetimes. The resulting program does not claim an absolute resolution of the classical polynomial versus nondeterministic polynomial problem; rather, it proposes a physically explicit reformulation and a suite of falsifiable hypotheses linking computation to spacetime.

Keywords: P vs NP; proper-time complexity; spacetime; general relativity; quantum speed limits; Bekenstein bound; black holes; computational complexity

Scope and Guarantees

This work proposes a physicalized reformulation of computational complexity, where resource costs are measured in proper time and constrained by physical laws. The conjectures presented are hypotheses about the interplay of spacetime geometry and computation, intended to be falsifiable through further theoretical and experimental investigation. We do not claim to resolve the classical P vs NP problem, but rather to offer a new perspective grounded in physical reality.

1. Introduction

The classical P vs NP problem asks whether every language decidable by a nondeterministic Turing machine (NTM) in polynomially many steps is also decidable by a deterministic Turing machine (DTM) in polynomially many steps. Since the seminal work of Cook and Karp, the P vs NP problem has been formalized in a model-robust way that counts *discrete steps* rather than wall-clock time [6,11]. The Cobham-Edmonds thesis connects polynomial time with the notion of feasible computation under standard machine models. This abstraction, while powerful, deliberately ignores the physical spacetime in which any real-world computation must take place.

In physics, however, time is not an absolute quantity. The theory of general relativity teaches us that the local proper time experienced by an observer depends on their worldline and the curvature of spacetime. Furthermore, any physical computing device is subject to fundamental limitations. The rate of computation is bounded by quantum speed limits, and the amount of information that can be stored is constrained by entropy bounds, such as the Bekenstein bound [4,12]. These physical realities motivate a complementary, *physicalized* perspective on computational complexity. This paper develops such a framework, reformulating resource bounds in terms of proper time and physical constraints, and then investigates how the resulting complexity classes relate to the classical, step-counting classes.

This paper makes the following contributions:

1. We provide a general definition of *proper-time complexity* classes, denoted as $P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ and $NP_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$, for computations that are realized along worldlines within a fixed spacetime $(\mathcal{M}, \mathbf{g})$ and under a specific resource-mapping \mathcal{R} .
2. We prove an isometry-covariance theorem, which demonstrates that our definitions are invariant under the symmetries of spacetime. However, they can vary across non-isometric embeddings and different worldlines.
3. We derive a universal time-space trade-off inequality from the Bekenstein entropy bound and the Margolus-Levitin quantum speed limit. This inequality couples the number of memory bits and logical operations that can be measured in proper time.
4. We formulate two conjectures that formalize when and how gravitational redshift can lead to *apparent* transitions in complexity class for distant observers, without violating the local step-count measures.
5. We provide a comparative analysis of our framework with the classical P vs NP problem, as well as with computational models that utilize closed timelike curves (CTCs) [2] and Malament-Hogarth spacetimes [7,10].

It is important to emphasize that our framework does not assert that the classical P vs NP problem becomes observer-dependent when measured by *step counts*. Instead, it introduces a physically grounded cost model where *proper-time* and thermodynamic constraints can produce observer-dependent assessments of tractability.

Related Work

Aaronson has surveyed various proposals to physically overcome the limitations of NP-completeness and has cautioned against over-interpreting the speedups that can be achieved through physical means [1]. The quantum speed limits, entropy bounds, and ultimate limits to computation provide a set of resource inequalities that are relevant to our work [4,8,12,13]. Computational models that have access to closed timelike curves (CTCs) have been shown to have the power of PSPACE [2], while Malament-Hogarth spacetimes can enable supertasks that go beyond the capabilities of Turing machines [7,10,16]. Our contribution is to provide an explicit and conservative complexity formalism that is applicable in fixed causal spacetimes, without making the assumption of CTCs or MH spacetimes.

1.1. Classical vs. Spacetime P vs NP

Definition 1.1 (Classical P vs NP Problem). *The Clay Millennium Problem asks: does P = NP, i.e. does every language decidable by a nondeterministic Turing machine in polynomially many steps also admit a deterministic polynomial-time algorithm?*

Conjecture 1 (Spacetime P vs NP). *Fix a globally hyperbolic spacetime $(\mathcal{M}, \mathbf{g})$ and a reasonable resource mapping \mathcal{R} . Then:*

- (i) *For any language L, membership in $P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ is observer-dependent: different worldlines can yield different tractability judgments.*

(ii) There exist universal curvature- and energy-dependent trade-offs (Lemma 1) that exclude entire classes of hypothetical algorithms (those violating $\text{Ops} \times \text{Bits}$ bounds), hence narrowing the plausible regime for P^τ vs NP^τ separations.

Interpretation. While the classical P vs NP is maximally hard—robust across machine models and unresolved after decades—the spacetime-aware version is less hard: its formulation immediately restricts admissible algorithms to those compatible with physical laws (QSL, Bekenstein, Landauer, redshift). Thus one obtains tangible, falsifiable constraints that classical theory cannot provide.

Table 1. Side-by-side comparison of classical P vs NP and Spacetime P vs NP .

	Classical Framework	Spacetime Framework
Resource measure	Step counts	Proper time + physical budgets
Observer dependence	None	Yes (different worldlines)
Excluded algorithms	None (all poly-time allowed)	Those violating $\text{Ops} \times \text{Bits}$, entropy, QSL
Problem hardness	Open, unconstrained	Constrained by physics; narrower space

2. Preliminaries

2.1. Classical Step-Complexity

We recall the standard definitions of the complexity classes P and NP. The class P is the set of all decision problems that can be solved by a deterministic Turing machine in a number of steps that is a polynomial function of the size of the input. The class NP is the set of all decision problems for which a given solution can be verified by a deterministic Turing machine in polynomial time. Formally, we have:

$$P = \bigcup_{k \in \mathbb{N}} \text{DTIME}(n^k) \quad (1)$$

$$NP = \bigcup_{k \in \mathbb{N}} \text{NTIME}(n^k) \quad (2)$$

These definitions are robust to reasonable variations in the underlying machine model, due to the linear-speedup and simulation theorems [3,9]. The time hierarchy theorem allows for the separation of complexity classes that are separated by a superpolynomial function, but it does not provide any insight into the relationship between P and NP [3].

2.2. Physical Resource Bounds

Any physical computation is subject to a number of fundamental bounds. We will make use of the following constraints in our framework:

- **Quantum speed limits (QSL).** For an isolated system with an average energy of E above the ground state, the minimum time Δt required to transition to an orthogonal state is given by the Margolus-Levitin theorem as $\Delta t \geq \frac{\pi\hbar}{2E}$. This implies that the maximum number of elementary operations that can be performed over a proper time interval of τ is bounded by $\text{ops}(\tau) \leq \frac{2E}{\pi\hbar} \tau$ [8,13].
- **Bekenstein bound.** For a system with a radius of R and a total energy of E , the entropy S is bounded by $S \leq \frac{2\pi k R E}{\hbar c}$. This, in turn, implies a bound on the memory capacity of the system, given by $B \leq \frac{S}{k \ln 2} \leq \frac{2\pi R E}{\hbar c \ln 2}$ bits [4,14].
- **Ultimate limits.** By combining the constraints of relativity and quantum mechanics, it is possible to derive ultimate limits on the total number of operations and the amount of storage for any computing device. An example of this is Lloyd's concept of the "ultimate laptop" [12].

3. Computations in Spacetime

Let us consider a globally hyperbolic spacetime, which is described by a pair $(\mathcal{M}, \mathbf{g})$, where \mathcal{M} is a manifold and \mathbf{g} is a metric. A computing device can be represented by a worldvolume \mathcal{W} that is a

subset of \mathcal{M} . The device has a timelike reference worldline, which we will denote as γ . The proper time along this worldline is given by τ . The average energy of the device above its ground state is a function of the proper time, $E(\tau)$, and the circumscribed radius of the device's worldtube slice at a given proper time is bounded by $R(\tau)$.

Definition 3.1 (Resource mapping). *A resource mapping, denoted by \mathcal{R} , is a function that assigns to a physical implementation the following functions:*

$$\Phi_{\text{ops}}(\tau) := \int_0^\tau \frac{2E(t)}{\pi\hbar} dt \quad (3)$$

$$\Phi_{\text{mem}}(\tau) := \sup_{t \leq \tau} \frac{2\pi R(t)E(t)}{\hbar c \ln 2} \quad (4)$$

These functions are interpreted as upper bounds on the number of logical operations that can be executed and the number of bits that can be reliably stored by a proper time of τ .

The resource mapping \mathcal{R} encapsulates the engineering details of the computing device, such as the overhead associated with error correction, control, and input/output. It can be made more precise by incorporating additional constraints, such as power, cooling, and signal delays. The resource mapping provides an *effective step budget*, $S(\tau) \leq \Phi_{\text{ops}}(\tau)$, and a memory budget, $B(\tau) \leq \Phi_{\text{mem}}(\tau)$, which are available for the simulation of an abstract machine.

Definition 3.2 (Proper-time complexity). *For a given decision problem $L \subseteq \{0,1\}^*$, we say that L is in the complexity class $P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ if there exist constants c and k , and a family of physical implementations in the spacetime $(\mathcal{M}, \mathbf{g})$, such that for any input x , a device following a timelike worldline γ can decide whether $x \in L$ within a proper time of $\tau(x) \leq c|x|^k$, while adhering to the resource budgets specified by \mathcal{R} . The complexity class $NP_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ is defined in a similar manner, using verifiers that, when given a witness of length at most $|x|^k$, can accept in a proper time of at most $c|x|^k$ for yes-instances.*

This definition is analogous to the classical definitions of P and NP, but it measures the *physical* runtime in terms of proper time. The number of logical steps is implicitly limited by the function Φ_{ops} . Therefore, a polynomial bound on the proper time τ implies that there is sufficient physical capacity to simulate the required number of steps within the given spacetime constraints.

3.1. Covariance and Observer Dependence

Proposition 1 (Isometry covariance). *If $\psi : (\mathcal{M}, \mathbf{g}) \rightarrow (\mathcal{M}, \mathbf{g})$ is an isometry and the resource mapping \mathcal{R} is defined by the scalar quantities E and R , and the proper time along the worldline γ , then a decision problem L is in the complexity class $P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ if and only if L is in the complexity class $P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ under the pulled-back implementation $\psi(\mathcal{W})$ and the worldline $\psi \circ \gamma$.*

This proposition shows that the complexity classes are invariant under changes of coordinates. However, they can depend on the specific worldline that is chosen to execute the computation. For example, a worldline that experiences extreme gravitational redshift will have a different complexity class than an inertial observer in an asymptotically flat region of spacetime.

Proposition 2 (Closure under Karp reductions). *If $A \leq_p B$ via a reduction f computable in $P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ with budgets bounded by $p_1(|x|)$ and $B \in P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ with budgets bounded by $p_2(|x|)$, then $A \in P_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$ with budgets bounded by $\text{poly}(p_1(|x|), p_2(|x|))$.*

4. A Universal Time – Space Trade-Off

By combining the quantum speed limit and the Bekenstein bound, we can derive a fundamental trade-off between the number of achievable operations and the amount of memory that can be stored, as a function of proper time.

Lemma 1 (Ops \times bits bound). *For any computing device with an energy of E and a radius of R (both of which may be time-dependent) that executes for a proper time of τ , the maximum number of operations, $\text{Ops}(\tau)$, and the number of storable bits, $\text{Bits}(\tau)$, are constrained by the following inequality:*

$$\text{Ops}(\tau) \cdot \text{Bits}(\tau) \lesssim \frac{4E^2 \tau R}{\hbar^2 c \ln 2} \quad (5)$$

This bound can only be approached under idealized conditions.

Proof sketch. According to the Margolus-Levitin theorem, the maximum number of operations that can be performed in a proper time of τ is given by $\text{Ops}(\tau) \leq \frac{2}{\pi\hbar} \int_0^\tau E(t) dt$. The Bekenstein bound states that the number of bits that can be stored at a given time t is limited by $\text{Bits}(t) \leq \frac{2\pi R(t) E(t)}{\hbar c \ln 2}$. By taking the suprema and products of these two inequalities, and by absorbing the geometric factors, we arrive at the desired result. \square

Lemma 1 is a significant result because it prevents pathological scenarios in which both time and space resources can be scaled independently without incurring any energy or geometric costs. It provides a formal statement of a trade-off that is enforced by the laws of physics.

5. Gravitational Redshift and Apparent Speedups

Let us consider a spacetime $(\mathcal{M}, \mathbf{g})$ that contains a static region with a timelike Killing field, which we will denote as ∂_t . The redshift factor in this region is given by $\alpha(x) = \sqrt{-\mathbf{g}(\partial_t, \partial_t)}$. For a computing device that is located at a radius of r outside a Schwarzschild black hole, the relationship between the proper time, $d\tau$, and the coordinate time, dt , is given by $d\tau = \alpha(r) dt$, where $\alpha(r) = \sqrt{1 - \frac{2GM}{rc^2}}$. In this scenario, a computation that takes a polynomial amount of proper time, τ , can appear to be exponentially faster when measured in the distant coordinate time, t , as the redshift factor, α , approaches zero. However, this apparent speedup is subject to constraints on energy, stability, and communication.

Definition 5.1 (Coordinate-time complexity). *For a distant observer who uses the coordinate time t , we can define the complexity class $\text{P}_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^t$ in a manner that is analogous to the definition of $\text{P}_{(\mathcal{M}, \mathbf{g}, \mathcal{R})}^\tau$, but with the time being measured at infinity (when this is well-defined).*

Conjecture 2 (Frame-Dependence Conjecture). *Let us fix a spacetime $(\mathcal{M}, \mathbf{g})$ and a reasonable family of resource mappings, \mathcal{F} , which is closed under constant-factor overheads and standard fault-tolerance. There exist decision problems, L , and observers on worldlines, γ_1 and γ_2 , such that L is in the complexity class $\text{P}_{(\mathcal{M}, \mathbf{g}, \mathcal{R}_1)}^\tau$ but not in the complexity class $\text{P}_{(\mathcal{M}, \mathbf{g}, \mathcal{R}_2)}^t$ for some resource mappings \mathcal{R}_1 and \mathcal{R}_2 in the family \mathcal{F} , and vice versa.*

This conjecture captures the idea that judgments about tractability can depend on which clock is used to measure the time (i.e., proper time or coordinate time), even when the step complexity of the computation remains unchanged.

Conjecture 3 (Gravitational Acceleration Threshold). *There exist families of problems and ranges of parameters (for example, near-horizon redshift with bounded tidal forces and sustainable energy and radius) where, for typical NP-complete instances of size n , a physically realizable verifier that runs in a polynomial amount of proper time will have a coordinate-time scaling of $t \in n^{O(1)} \cdot \alpha(r)$. As the redshift factor, α , approaches zero, the apparent runtime in terms of the coordinate time, t , will cross certain thresholds (for example, from super-polynomial to polynomial), without any change in the local step complexity.*

Conjecture 3 reframes the speedups that are induced by redshift as phenomena that are *observer-relative*. These phenomena are constrained by the trade-off that is described in Lemma 1, as well as by the costs that are associated with moving the computing device and exfiltrating the outputs.

Caveats

It is important to note that these conjectures do not imply that $P = NP$ in the classical, step-counting sense. Instead, they articulate the differences between the cost models of *physical-time* and *step-time*. Furthermore, scenarios that involve closed timelike curves (CTCs) or Malament-Hogarth spacetimes can lead to qualitative changes in computability and complexity (such as PSPACE power or capabilities beyond those of a Turing machine) [2,7,10]. These scenarios are outside the scope of our fixed-causal framework.

Author Contributions: Sole author: conceptualization, formal analysis, writing.

Funding: None.

Data Availability Statement: No data were analyzed; all results are theoretical.

Acknowledgments: The author acknowledges the use of AI assistance in developing and refining the mathematical formulations and computational validations presented in this work. All theoretical results, proofs, and interpretations remain the responsibility of the author.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. S. Aaronson, "NP-complete problems and physical reality," *ACM SIGACT News*, vol. 36, no. 1, pp. 30-52, 2005.
2. S. Aaronson and J. Watrous, "Closed timelike curves make quantum and classical computing equivalent," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 465, no. 2102, pp. 631-647, 2009.
3. S. Arora and B. Barak, *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.
4. J. D. Bekenstein, "A universal upper bound on the entropy-to-energy ratio for bounded systems," *Physical Review D*, vol. 23, no. 2, pp. 287-298, 1981.
5. A. Cobham, "The intrinsic computational difficulty of functions," *Logic, Methodology and Philosophy of Science: Proceedings of the 1964 International Congress*, pp. 24-30, 1965.
6. S. A. Cook, "The complexity of theorem-proving procedures," *Proceedings of the third annual ACM symposium on Theory of computing*, pp. 151-158, 1971.
7. G. Etesi and I. Németi, "Non-Turing computations via Malament-Hogarth space-times," *International Journal of Theoretical Physics*, vol. 41, no. 2, pp. 341-370, 2002.
8. V. Giovannetti, S. Lloyd, and L. Maccone, "Quantum limits to dynamical evolution," *Physical Review A*, vol. 67, no. 5, p. 052109, 2003.
9. J. Hartmanis and R. E. Stearns, "On the computational complexity of algorithms," *Transactions of the American Mathematical Society*, vol. 117, pp. 285-306, 1965.
10. M. L. Hogarth, "Does general relativity allow supertasks?", *Philosophy of Science*, vol. 59, no. 1, pp. 116-123, 1992.
11. R. M. Karp, "Reducibility among combinatorial problems," *Complexity of computer computations*, pp. 85-103, 1972.
12. S. Lloyd, "Ultimate physical limits to computation," *Nature*, vol. 406, no. 6799, pp. 1047-1054, 2000.
13. N. Margolus and L. B. Levitin, "The maximum speed of dynamical evolution," *Physica D: Nonlinear Phenomena*, vol. 120, no. 1-2, pp. 188-195, 1998.
14. D. N. Page, "Comment on 'A universal upper bound on the entropy-to-energy ratio for bounded systems'," *Physical Review D*, vol. 97, no. 10, p. 108501, 2018.
15. B. A. Trakhtenbrot, "A survey of Russian approaches to *perebor* (brute-force search) algorithms," *Annals of the History of Computing*, vol. 6, no. 4, pp. 384-400, 1984.
16. P. D. Welch, "The extent of computation in Malament-Hogarth spacetimes," *The British Journal for the Philosophy of Science*, vol. 59, no. 4, pp. 659-674, 2008.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.