Leveraging Performance Portability for High-Fidelity Simulations of Black Hole Accretion

Vedant Ketan Dhruv*; Advisor: Charles F. Gammie

Department of Physics, University of Illinois at Urbana-Champaign

*vdhruv2@illinois.edu

Introduction

General Relativistic Magnetohydrodynamic (GRMHD) simulations are essential for interpreting observations from the Event Horizon Telescope (EHT). By leveraging high-performance and high-throughput computing resources, these models not only bridge theory and observation but also improve our fundamental understanding of plasma physics and gravity in the extreme environment around a black hole.

Theoretical models suggest that the plasma surrounding EHT sources is effectively collisionless, which may have significant implications for what we observe. Furthermore, the immense gravity of supermassive black holes provides a unique laboratory to test the limits of Einstein's theory of General Relativity. This work addresses two key questions at these frontiers:

1. How does including the physics of low-collisionality plasma affect synthetic observables compared to standard models?
2. How well can EHT observations constrain potential modifications to General Relativity?

To answer these questions, we modify KHARMA—a performance-portable code designed to run efficiently on the evolving, heterogeneous landscape of modern supercomputers.

Messier 87*

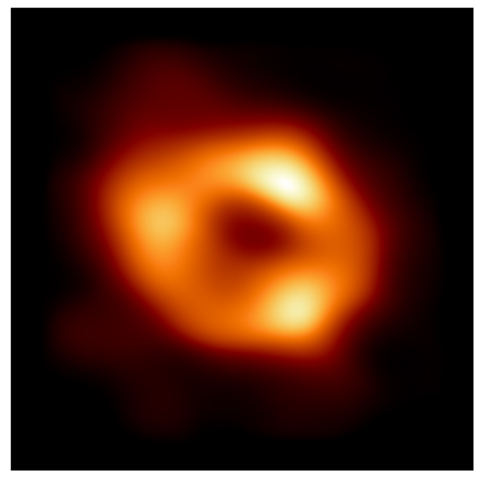

Sagittarius A* (Galactic center)

Link to EHT M87* Theoretical Interpretation

Link to EHT SgrA* Theoretical Interpretation

KHARMA: Performance-Portable GRMHD

KHARMA is an open-source black hole accretion code based on the HARM algorithm, originally written for ideal GRMHD. It is one of the first codes in the field designed to be vendor-agnostic, focusing on performance, extensibility, and adaptability for academic domain scientists. KHARMA's design for requirements are,

1. **Performance portability**: KHARMA leverages the Kokkos C++ library to efficiently run a single codebase on compute and accelerator architectures from any vendor. The underlying Parthenon framework further boosts performance by keeping simulation data in device memory, minimizing data transfers to the CPU. Combined with asynchronous, GPU-aware MPI calls, this design achieves excellent scaling on leading supercomputers, as shown in the Figure below.
2. **Modular and extensible design**: KHARMA inherits its modularity from the Parthenon AMR framework, which enables a 'plug-and-play' interface. Parthenon's modular "packages" let physics modules register variables with clear scopes and callbacks, so components compose cleanly without cross-coupling. This modularity is orchestrated by a dynamic task scheduler, which replaces a hard-coded execution order. This also boosts performance by enabling MPI communication to overlap with computation.
3. **Adaptability in academia**: KHARMA prioritizes simplicity of interaction over strict minimalism. Stable package APIs and clear task/variable registration patterns let newcomers extend the code without touching core internals. Modularity localizes changes, curbs regressions, and makes adding, trimming, or re-introducing features routine for small, time-limited teams.

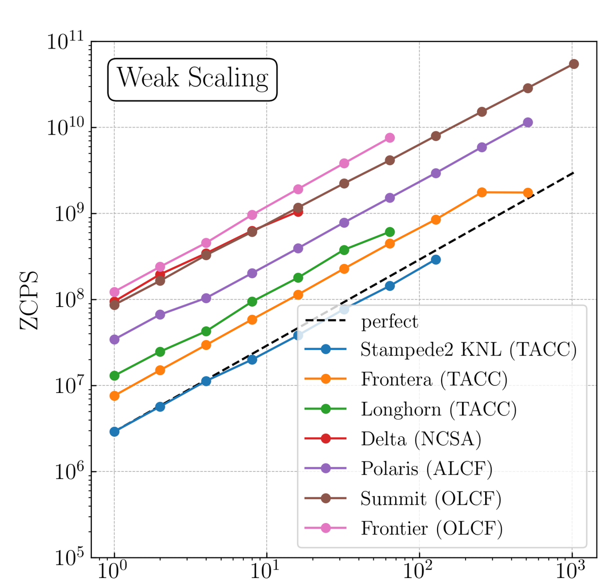

Weak scaling results for KHARMA

KHARMA code paper

KHARMA GitHub repo

Weakly-collisional black hole accretion

- Inferring black hole properties from EHT data requires extensive parameter studies that involve running large libraries of long-duration simulations [1]. We utilize a weakly-collisional fluid model as a crucial intermediate, which offers more realistic plasma physics than ideal GRMHD at a computational cost far below that of global Particle-in-Cell (PIC) models.
- We employ an Extended GRMHD (EGRMHD) model that incorporates leading-order corrections for long particle mean free paths, namely anisotropic pressure (ΔP) and heat flux (q). The semi-implicit nature of this scheme is roughly 10x more expensive per step than ideal GRMHD. However, KHARMA's performance portability makes these demanding simulations practical by efficiently leveraging powerful GPU architectures.

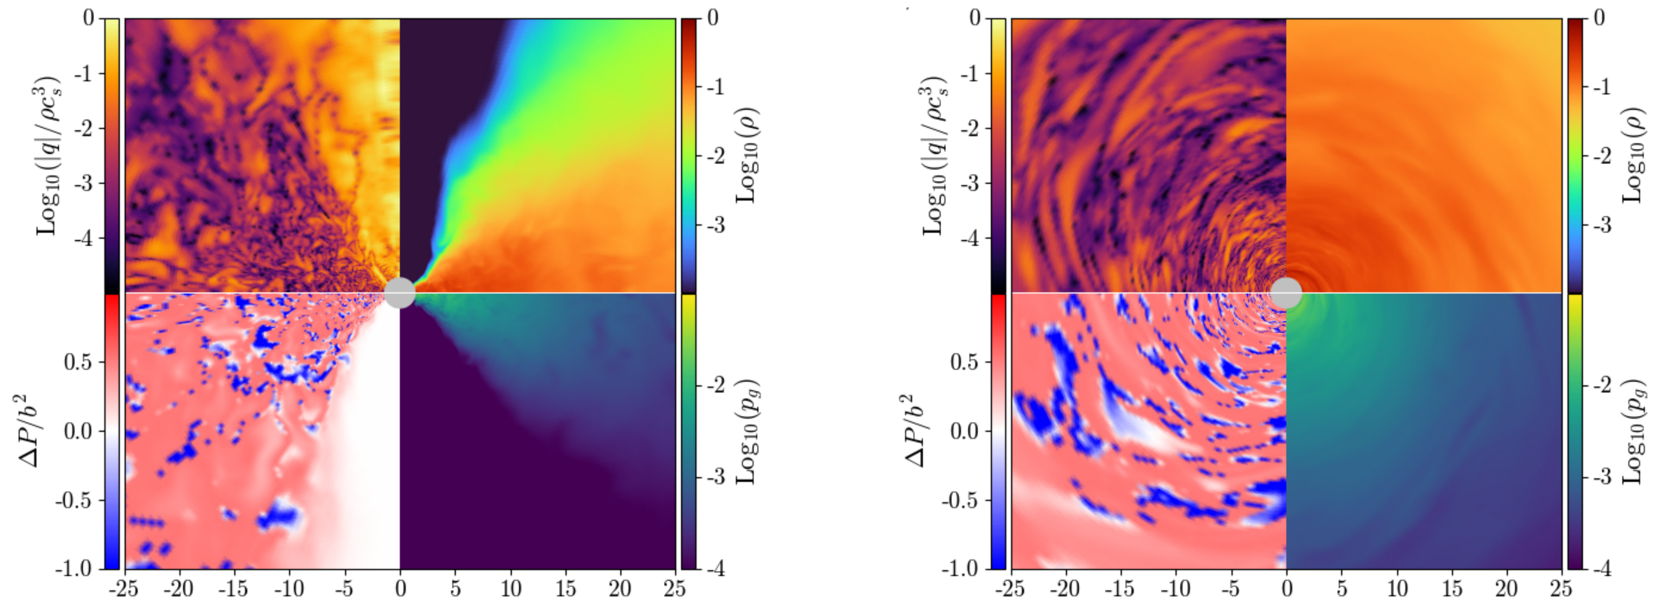

Snapshot of a weakly-collisional black hole accretion simulation. Left: poloidal (r, θ) and Right: midplane (r, ϕ) slice of various fluid quantities. q and ΔP represent new physics associated with low collisionality in these systems.

- The implicit solver is designed to be physics-agnostic, allowing new physics modules to reuse the same high-performance kernel. This is achieved by performing the implicit solve only on those variables that are tagged with the `MetadataFlag` "`Implicit`".
- The implicit routine is fundamentally a multidimensional Newton-Raphson scheme which minimizes a vector-valued loss function over the end-of-timestep state primitives $R(P^{n+1})$. This effective replacement of the 1D solver used in ideal GRMHD with a 7D solver makes EGRMHD simulations approximately 10x more expensive.
- The solver begins with an initial guess for the state primitives P^n. In each iteration, it computes the Jacobian and performs a linear solve to find a correction vector δP^{n+1}. This process is repeated until the desired accuracy is reached, with a backtracking line search invoked at each step to improve and ensure convergence.
- The implicit kernel uses hierarchical parallelism to leverage Kokkos Kernels' batched linear routines. The solver operates by first performing a pivoted QR decomposition on the system of equations in each cell, followed by a triangular solve to find the solution.
- A key result of our work is that the EGRMHD simulations are significantly less variable than their ideal GRMHD counterparts [2]. As shown in the comparison plot, this reduction in the 3-hour modulation index brings the simulated radio light curves into better agreement with EHT observations.

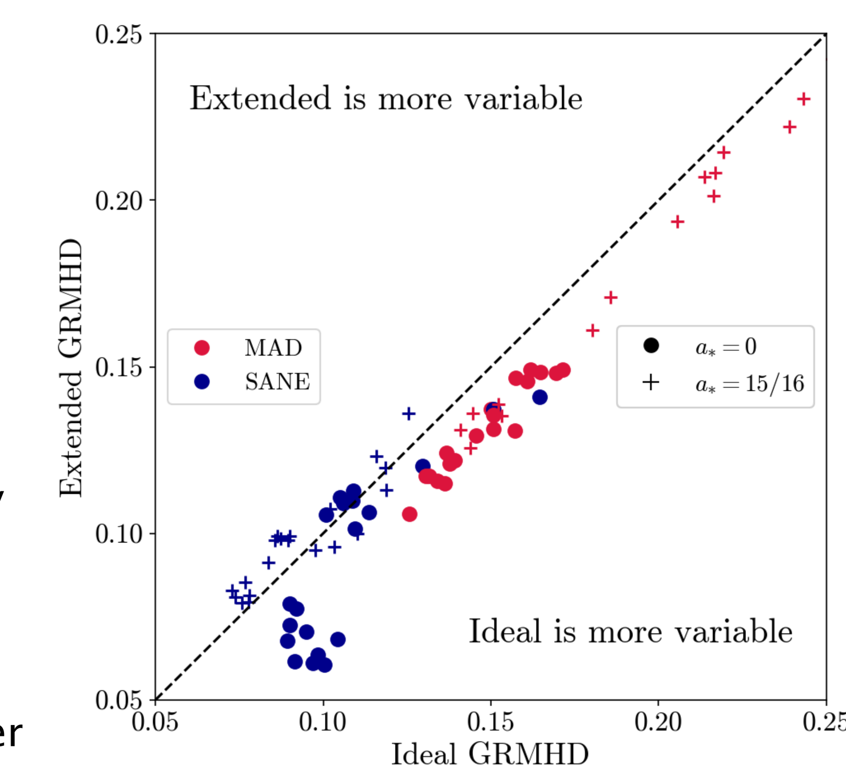

The 3-hour modulation index (rms %) for Extended and Ideal GRMHD models.

Testing Einstein's Theory of Gravity

- The strong gravitational fields near a black hole's event horizon potentially offer a new regime to test the predictions of General Relativity. These observations probe conditions of strong gravitational fields and considerable spacetime curvature that are inaccessible by other astrophysical sources.
- We leverage KHARMA's extensible design to simulate black hole accretion in alternative theories of gravity, specifically dynamical Chern-Simons (dCS) and Einstein-dilaton-Gauss-Bonnet (EdGB). By generating synthetic images from these models, we can constrain potential deviations from General Relativity using EHT observations.
- The resulting time-averaged images from these theories are nearly indistinguishable (see figure). We find the black hole's unknown spin creates larger image variations than the subtle differences between gravitational theories [3]. This degeneracy makes it challenging to use EHT data to constrain deviations from General Relativity.

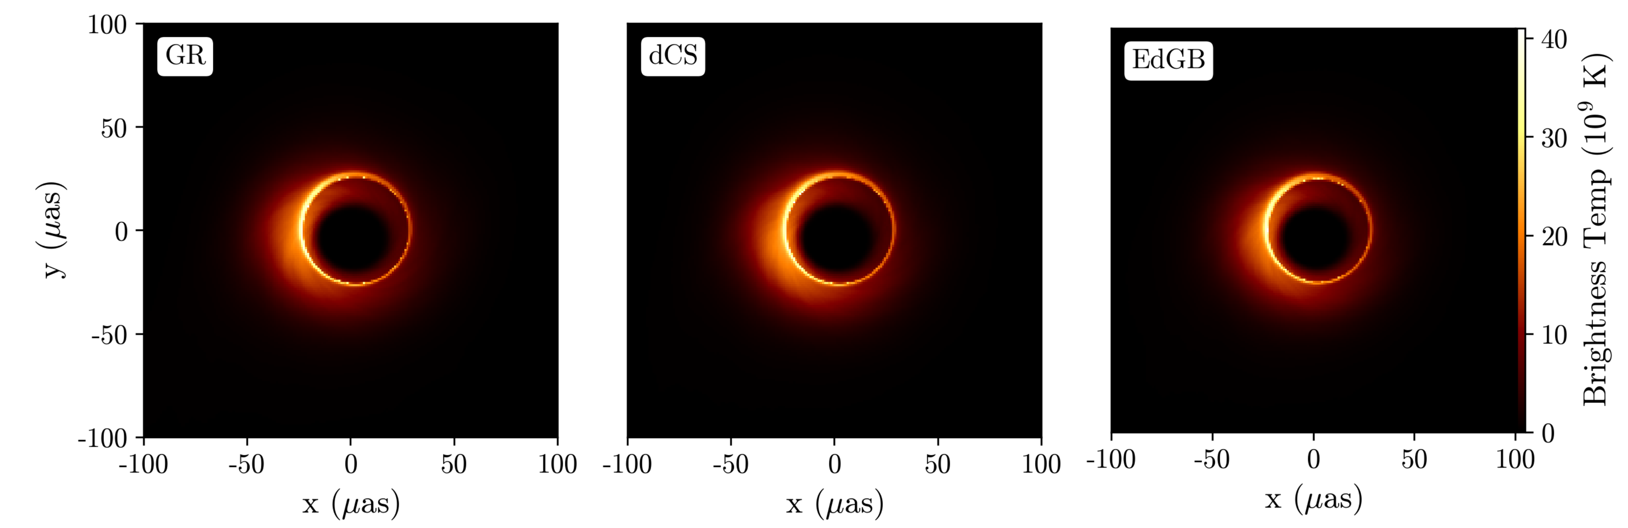

Time-averaged synthetic images for different theories of gravity

Summary

- Performance-portable and extensible GRMHD codes are essential for building high-fidelity models of black hole accretion flows on modern, diverse supercomputing architectures.
- We leveraged the performance-portable code KHARMA to generate a library of simulations with low-collisionality physics, producing synthetic observables that provide an improved fit to EHT observations of Sagittarius A*.
- Our simulations show that alternate theories of gravity produce time-averaged images remarkably similar to General Relativity. We find that deviations from GR are largely degenerate with the black hole's spin, making it challenging to constrain these theories with current EHT observational capabilities.

Acknowledgements

V.K.D. is supported by the Dissertation Completion Fellowship (UIUC), Jones Fellowship (Physics, UIUC), and was supported by the ICASU/NCSA Fellowship. This work was supported by NSF grants AST 17-16327 (horizon), OISE 17-43747, and AST 20-34306. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory and that of the Argonne Leadership Computing Facility. This research is part of the Delta research computing project, which is supported by the National Science Foundation (award OCI 2005572), and the State of Illinois. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), specifically the XSEDE resources Longhorn, Frontera, and Stampede2 at the Texas Advanced Computing Center (TACC) through allocation TG-AST170024.

References

1. **Vedant Dhruv**, Ben Prather, George N. Wong, and Charles F. Gammie. "A Survey of General Relativistic Magnetohydrodynamic Models for Black Hole Accretion Systems". The Astrophysical Journal Supplement Series, 277(1):16, March 2025. doi:10.3847/1538-4365/adaea6.
2. **Vedant Dhruv**, Chandra Mani Joshi Abhishek V. Prather, Ben S., and Charles F. Gammie. "Electromagnetic observables of weakly collisional black hole accretion". Physical Review Letters (submitted), 2025
3. **Vedant Dhruv**, Yiqi Xie, Tegan Thomas, Shreya Majumdar, Abhishek Joshi, Ben S., Nicolas Yunes, and Charles F. Gammie. "Testing Deviations from General Relativity with EHT Observations". In prep.