



Performance-Portable Numerical Relativity with AthenaK and Applications



Hengrui Zhu
Princeton Gravity Initiative



Outline

- **Numerical Relativity at Exascale**
- **Introducing Athenak Z4c Module (Zhu+ 2409.10383)**
- **Athenak BBH Tutorial and GW150914 (Radice+ 2506.06838)**
- **Ultra-relativistic black hole encounter (Zhu+ 2604.26253)**
- **Mini-lecture on Gauge Conditions (Zhu and Truong in prep)**

Emergence of Exascale

1,000,000,000,000,000,000 FLOPS



Emergence of Exascale

1,000,000,000,000,000,000 FLOPS




ZetaFLOPS coming soon...

Numerical Relativity at Exascale

Head-on Collision with $q=20$; ~20 M/hr 64 GPUs



- Large Mass Ratio

- Producing binary black hole waveforms for next-generation detectors (CE, ET, LISA, etc.)
- Comparison and calibration for perturbation theory (self-force)




Numerical Relativity at Exascale

- Large Mass Ratio
- Near Extremal Spin
 - Initial data from Idsolve, SpECTRE, etc.
 - Very High Resolution Required at the Punctures see, e.g. Liu et al. 2010 (1001.4077)



(NRSUR) Posterior distribution of GW231123

Numerical Relativity at Exascale




Consistent with Maximally Spinning?

Numerical Relativity at Exascale

Targeted NR BBH survey for GW231123; Zhu with Dittmann, Gamba, Siemonsen, Fields, Radice (in prep)


- Large Mass Ratio
- Near Extremal Spin
- Scanning Large Parameter Space
 - Many science goals limited by the number of available BBH simulation:
 - building and calibrating waveform models
 - targeted simulation for observations



Numerical Relativity at Exascale


Video courtesy by Jacob Fields

- Large Mass Ratio
- Near Extremal Spin
- Scanning Large Parameter Space
- **Multi-physics**
 - MHD (Fields + 2409.10384)
 - Neutrino/Radiation Transport
 - Yi + in prep (M1)
 - Zhang + in prep (multi-frequency/multi-group)
 - Gutiérrez + in prep (Neutrino Opacity)
 - Zhu + in prep (Boltzmann)
 - Nuclear Burning (Alex's talk)



Numerical Relativity at Exascale

- Large Mass Ratio
- Near Extremal Spin
- Scanning Large Parameter Space
- Multi-physics
- **Science Targets**
 - Binary Compact object merger
 - Supernovae
 - Collapsar
 - TDE
- **GR theory**



Zhu, Pretorius, and Stone 2604.26253

Outline

- **Numerical Relativity at Exascale**
- **Introducing Athenak Z4c Module (Zhu+ 2409.10383)**
- **Athenak BBH Tutorial and GW150914 (Radice+ 2506.06838)**
- **Ultra-relativistic black hole encounter (Zhu+ 2604.26253)**
- **Mini-lecture on Gauge Conditions (Zhu and Truong in prep)**


Numerical Integration of Einstein Equations with Z4c Formalism

- Properties of Z4c

- 3+1 decomposition of **augmented** Einstein equation
- No zero-speed mode and explicit constraint damping scheme
- Conformally rescaled (compatible with puncture evolution no need for excision!)

$$G_{ab} = 8\pi T_{ab} - \nabla_a Z_b - \nabla_b Z_a + g_{ab} \nabla_c Z^c + \kappa_1 [n_a Z_b + n_b Z_a + \kappa_2 g_{ab} n_c Z^c],$$


Figure from Baumgarte and Shapiro



Numerical Integration of Einstein Equations with Z4c Formalism


- Properties of Z4c

- 3+1 decomposition of augmented Einstein equation
- No zero-speed mode and explicit constraint damping scheme
- Conformally rescaled (compatible with puncture evolution no need for excision!)




Numerical Integration of Einstein Equations with Z4c Formalism

- Properties of Z4c
- Numerical Implementations
 - 2nd, 4th, and 6th order finite differencing
 - RK2, 3, 4 time integrator (low storage variant)
 - Adaptive Mesh Refinement with the Athena++ framework
 - Task-based parallelism




Numerical Integration of Einstein Equations with Z4c Formalism

- Properties of Z4c
- Numerical Implementations
 - 2nd, 4th, and 6th order finite differencing
 - RK2, 3, 4 time integrator (low storage variant)
 - Adaptive Mesh Refinement with the Athena++ framework
 - Task-based parallelism
 - Cell-Centered instead of Vertex-Centered




Accuracy and Convergence

- BAM calibration run (TwoPuncture ID)
- Equal-mass and non-spinning binary merging within 3 orbits




Accuracy and Convergence

- BAM calibration run
- Equal-mass and non-spinning binary merging within 3 orbits
- Demonstrate 4th order convergence



Accuracy and Convergence

- BAM calibration run
- Equal-mass and non-spinning binary merging within 3 orbits
- Demonstrate 4th order convergence
- Consistent with previous codes (GR-Athena++)



Balancing Accuracy and Efficiency with AMR?

- What is the **optimal** grid structure for binary black hole?
 - p-refinement vs h-refinement





Figure by Deus Timbobong


Balancing Accuracy and Efficiency with AMR?

- What is the optimal grid structure for binary black hole?
- BAM calibration run, revisited
 - Low resolution with 10 levels of AMR




Balancing Accuracy and Efficiency with AMR?

- What is the optimal grid structure for binary black hole?
- BAM calibration run, revisited
 - Low resolution with 10 levels of AMR
 - P-refinement: mid resolution, increase number of points within each MeshBlock.




Balancing Accuracy and Efficiency with AMR?

- What is the optimal grid structure for binary black hole?
- BAM calibration run, revisited
 - Low resolution with 10 levels of AMR
 - P-refinement: mid resolution, increase number of points within each MeshBlock.
 - **H-refinement: low with 11 levels of AMR**



Balancing Accuracy and Efficiency with AMR?

- What is the optimal grid structure for binary black hole?
- BAM calibration run, revisited
- Optimal Strategy:
 - H-refinement during **inspiral** (saves computational cost)
 - P-refinement during **merger-ringdown**



Scaling and Performance Portability


- Performance Portability

- Tested with Nvidia, AMD, Intel GPUs; ARM, and X86 CPUs

Configuration	ZCPS	1/Speed-up
NVIDIA A100	1.36×10^7	1
NVIDIA Grace Hopper	2.46×10^7	0.55
NVIDIA RTX 3070 Ti	3.90×10^6	3.48
AMD MI250X (single die)	5.08×10^6	2.68
AMD 7950x Zen 3 (single core)	4.53×10^4	300.3
Intel Cascade Lake (single core)	1.71×10^4	793.65
Apple M3 (single core; ARM)	6.53×10^4	208.33


Scaling and Performance Portability

- Performance Portability
- Weak Scaling
 - Evolving a single puncture with 10 levels of mesh refinement
 - 80% efficiency up to 65,536 GPUs on Frontier



Scaling and Performance Portability

- Performance Portability
- Weak Scaling
- Strong Scaling
 - ~80% efficiency when increasing the number of GPU by 32 times on both Frontier and Perlmutter



Outline


- **Numerical Relativity at Exascale**
- **Introducing Athenak Z4c Module (Zhu+ 2409.10383)**
- **Athenak BBH Tutorial and GW150914 (Radice+ 2506.06838)**
- **Ultra-relativistic black hole encounter (Zhu+ 2604.26253)**
- **Mini-lecture on Gauge Conditions (Zhu and Truong in prep)**

Tutorial: Running Binary Black hole with Athenak

Radice + 2506.06838


- Reproduce GW150914-like BBH (Lovelace et al. 1607.05377)
- Consistency check with SXS and RIT
- Provide a comprehensive walk-through of the BBH functionality

<https://github.com/dradice/athenak-tutorial-gw150914>




Obtaining waveform Through CCE

- Two Puncture **Initial Data** (same parameters as the RIT LazEV runs)
- Evolution in Athenak
- Dump ADM variables on a **World Tube**
- **Cauchy Characteristic Extraction** with SpECTRE CCE
- Align BMS frame at $\mathcal{I}^+$




Obtaining waveform Through CCE



- Dominant spin-weighted spherical harmonic modes
- Memory effect!



Comparison with SXS and RIT



- Consistent with SXS and RIT simulations
- Difference might be attributed to waveform extraction techniques

CCE vs. Finite Radius


- Junk radiation in CCE causing wiggles in the waveform at $\mathcal{I}^+$



Outline

- **Numerical Relativity at Exascale**
- **Introducing Athenak Z4c Module (Zhu+ 2409.10383)**
- **Athenak BBH Tutorial and GW150914 (Radice+ 2506.06838)**
- **Ultra-relativistic black hole encounter (Zhu+ 2604.26253)**
- **Mini-lecture on Gauge Conditions (Zhu and Truong in prep)**

Surprises from Numerical Relativity




Kip Thorne, unpublished




AI Generated

Surprises from Numerical Relativity



Kip Thorne, unpublished




Pretorius and Khurana 2007

Where goes the nonlinearity in BBH?

GR


$$R_{\mu\nu} \supset (g^{-1}\partial g)^2$$



T. Dietrich and R. Haas; SXS Collaboration


Navier-Stokes

$$\partial_t u^i \supset u^j \partial_j u^i$$



Schaal et al. 2015


GR $\approx$ MHD



Boyeneni, Wu, and Most 2025;
Krynicky, Wu, and Most 2025
Casting GR under Elsasser formulation


Nonlinearity in other settings

ADS (conformal fluid)



Carrasco, Lehner, et al. 2012




Driven (Weak) “Turbulence”



Ma, Lehner, et al. 2025


Nonlinearity in other settings

Inverse Cascades





Nonlinearity in other settings

Inverse Cascades



CONFINEMENT

Sustained High Energy



Ultra-Relativistic Black Hole Encounter

- Kinetic Energy of B.H. is much greater than the rest mass energy $\gamma M_0 \gg M_0$




Figure from Healy et al. 2410.20239

Ultra-Relativistic Black Hole Encounter

- Kinetic Energy of B.H. is much greater than the rest mass energy $\gamma M_0 \gg M_0$
- Circumvent entropy law which bounds the amount of energy radiated in rest-mass dominated systems (**high energy**)




Figure from Healy et al. 2410.20239

Ultra-Relativistic Black Hole Encounter

- Kinetic Energy of B.H. is much greater than the rest mass energy $\gamma M_0 \gg M_0$
- Circumvent entropy law which bounds the amount of energy radiated in rest-mass dominated systems (high energy)
- Kinetic Energy GRAVITATE (**confinement**)








Figure from Healy et al. 2410.20239

Kip's Fantasy Comes True





$$\gamma \approx 5.1$$

Energy Spectrum




Visualizing Strong Field Interaction




Using Tendex-Vortex formalism; See e.g., Owen et al. 2011;
 Zimmerman et al. 2011; Nichols et al. 2011; Zhang et al. 2012;
 Nichols et al. 2012

$$|P| \approx \frac{1}{2} \left| |\Psi_4|^2 - |\Psi_0|^2 \right|$$

Understanding Trapping with Null Integral Curves




$\Gamma = 5.1$





$\Gamma = 1.5$
Small and short lived trapping region


Understanding Trapping with Null Integral Curves




Why Astrophysical Mergers is So Smooth?



Why Astrophysical Mergers is So Smooth?



Wavelength $\gg$ Aperture size:
Diffraction and little confinement :(




Wavelength $\ll$ Aperture size:
Refraction and confinement :)

Why Astrophysical Mergers is So Smooth?




AI Generated

Other Implications




- Strong Absorption For Close Encounter

Other Implications




- Strong Absorption For Close Encounter
- Radiation Efficiency Suppressed Near Merger Threshold (extrapolation to $\gamma \rightarrow \infty$ still unclear)

Other Implications



- Strong Absorption For Close Encounter
- Radiation Efficiency Suppressed Near Merger Threshold
- A New Branch Of Impact Parameter For Maximal Radiation! Inconsistent with previous anticipation/extrapolation

Other Implications




- Strong Absorption For Close Encounter
- Radiation Efficiency Suppressed Near Merger Threshold
- A New Branch Of Impact Parameter For Maximal Radiation! Inconsistent with previous anticipation/extrapolation
- See Frans' talk for details!

Mini Gauge Condition Lecture


(if time permits)

Email: hz0693@princeton.edu

Gauge Condition



Initial Data



Convergence

