


"BH formation and associated explosive events" & "A new GPU-accelerated numerical relativity code: ANRe-M1"



• Takami Kuroda (MPI Potsdam, Germany)




Max-Planck-Institut
für Gravitationsphysik
(Albert-Einstein-Institut)

Brainstorming workshop - Multimessenger astrophysics with core-collapse
supernovae and binary neutron star mergers


Wrocław University of Science and Technology, Wrocław, 9. - 11. September 2026

Diverse scenario of massive stellar death




Intermediate mass SN
progenitors:
MZAMS $\sim 30\text{-}50 M_{\text{sun}}$

Various compact star formation: NS, HS(quark+baryonic) & BH




TK, Fischer+, '21

Various compact star formation:
NS, HS(quark+baryonic) & BH



TK, Fischer+, '21


BH formation




Beyond BH formation:
Rahman+22, Sykes+23, TK&Shibata23,24,
Burrows+'25, Goni+25

BH formation

Non rotating 70Msun in 2D




TK & Shibata, '23




see also Rahman+22, Sykes+23


SN simulation with magnetic fields and relativistic jet formation from BH




SN simulation with magnetic fields and relativistic jet formation from BH




SN simulation with magnetic fields and relativistic jet formation from BH




SN simulation with magnetic fields and relativistic jet formation from BH




SN simulation with magnetic fields and relativistic jet formation from BH



Multi-messenger signals from BH formation



Takeaway and next steps



Relativity (stellar mass/compactness ?)


- Next generation SN simulation will focus on heavier side of progenitor mass. —> Beyond BH formation with numerical relativity
- Relevant physics, e.g., nuclear physics, stellar evolution theory, (beyond-GR theory), is another key to uncover the diverse SN scenarios.

GPU numerical relativity has reached exascale


Recent primary benchmarks show both the opportunity and the scaling constraints

AthenaK

80% weak-scaling efficiency at 65,536 MI250X GPUs



GRACE



ANRe-M1 adds the missing CCSN combination
dynamical full GR + multi-species, multi-energy M1 neutrino transport

Why ANRe-M1: complete physics over seconds

Long-duration three-dimensional core collapse multiplies the work in every cell and every time step

Dynamical spacetime

BSSN–Z4c evolves the metric through collapse, bounce and possible black-hole formation.

GR hydrodynamics

Relativistic fluid dynamics couples shocks, accretion and turbulent post-shock flow to gravity.

Spectral neutrinos

M1 transport follows three species in 12 energy groups, including stiff matter interactions.

$$3 \times N_{\epsilon}$$

neutrino species $\times$ energy groups

$$4 \times 3 \times N_{\epsilon}$$

radiation-moment components per cell

$$(5 + 4 \times 3 \times N_{\epsilon})^2$$

local Jacobian entries in each implicit solve

The production target

Higher resolution, progenitor surveys and post-bounce evolution lasting several seconds

Published GPU NR codes generally use vacuum, M0 or gray transport. ANRe-M1 couples spectral M1 transport to a dynamical full-GR spacetime.

ANRe-M1 architecture

One C++/Kokkos implementation keeps the full update resident on the accelerator

C++

Kokkos

CPU · GPU · APU

BSSN-Z4c

Dynamical spacetime
4 Runge–Kutta stages
Constraint
propagation

GR hydrodynamics

PPM reconstruction
Numerical flux +
refluxing
Primitive recovery

Multi-energy M1

3 species, arbitrary N_E
M1 closure + opacities
Implicit matter coupling

Kokkos::View

flattened RangePolicy


GPU-aware MPI

fixed mesh refinement


Field components used by one work are contiguous in memory. Ghost buffers, EOS tables and weak-interaction tables remain on the device.

Validation from unit tests to relativistic collapse

The port reproduces designed convergence, exact solutions and conserved global quantities




$p = 2.023$



Exact Riemann structure

Validation from unit tests to relativistic collapse

The port reproduces designed convergence, exact solutions and conserved global quantities




$\approx 0.2\%$

ADM-mass error in
adiabatic collapse


Representative 3D result: 20 $M_{\odot}$ collapse with spectral neutrinos

ANRe-M1 follows collapse, bounce and the first 100 ms of post-bounce evolution



Performance on AMD MI300A

The benchmark advances full spacetime, hydrodynamics and 144 radiation moments per cell



2 → 128

MI300A APUs

39.7×

Speed up in comparison to CPU

62%

weak-scaling efficiency

7.4072×10^7 zone-cycles s^{-1} at 128 APUs

The practical step time is about an order of magnitude shorter than the established Fortran CPU workflow

Takeaway and next steps

Physics

Dynamical full GR and multi-species, multi-energy M1 now run together on accelerators.

Validation

Convergence, exact-solution tests and a 3D collapse comparison check the coupled implementation.

Performance

Throughput grows 39.7× from 2 to 128 MI300A APUs while retaining 62% weak-scaling efficiency.

Next

Larger active blocks

Communication overlap

Lower radiation-memory footprint

Longer, higher-resolution 3D runs

ANRe-M1 makes long-term, full-GR supernova calculations practical on heterogeneous machines → Application to Boltzmann transport