

Preparing for Exascale Computing with GRTeclyn

In-situ visualization and unified programming across
accelerated architectures

Juliana Kwan, Paul Shellard, James Fergusson, Amelia Drew, Ulrich Sperhake (CTC),
Marek Szuba (DAMTP), Miren Radia (UIS/DiRAC), Kacper Kornet (UIS)
Carson Brownlee, Dave DeMarle, John Pennycook, Maxwell Cai, Jay Mahalingham (Intel)

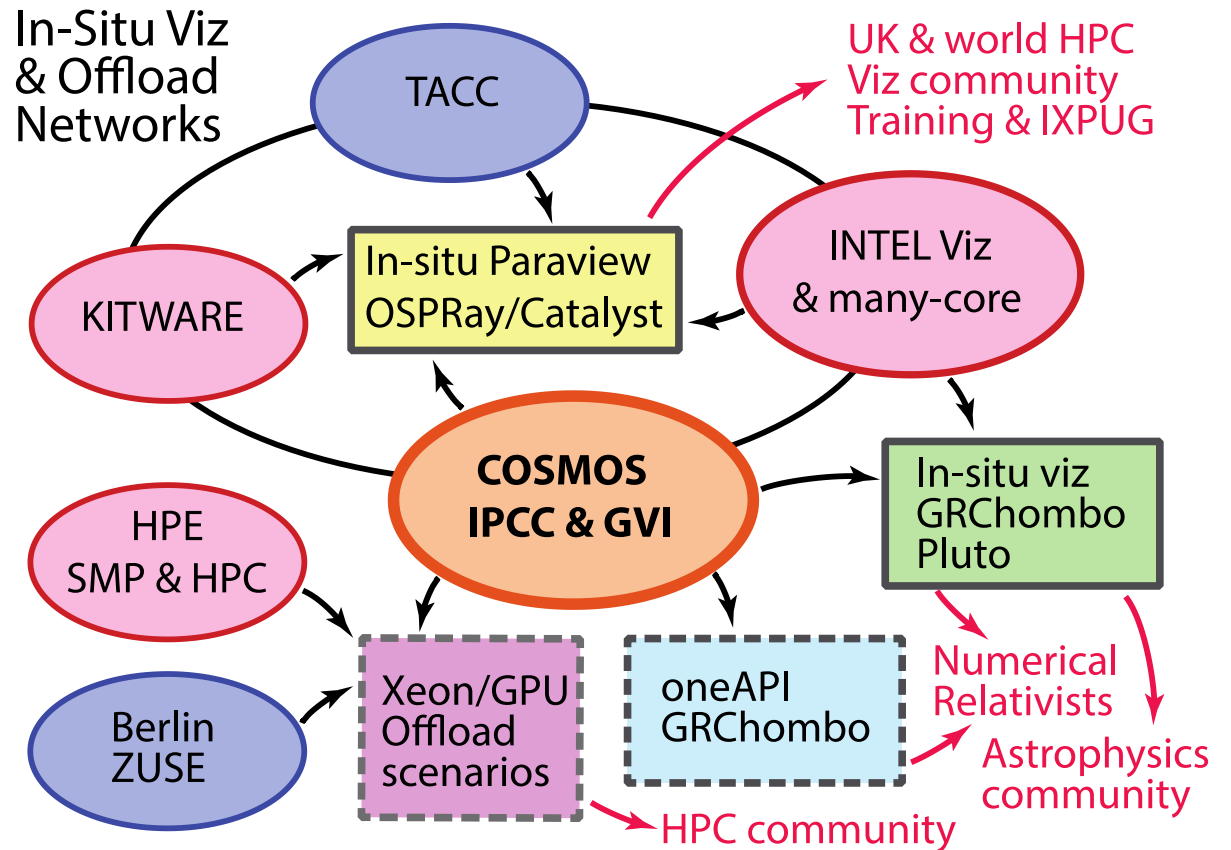


The Stephen Hawking
Centre for Theoretical Cosmology





What is this project about?

- Converting GRChombo to GRTeclyn
- Extend HPC in-situ visualization to fully incorporate AMR (particularly for GRTeclyn)
- Introduce in-situ capabilities to other community codes
- Creation of a Intel GPU testbed (available to all members of the ExCALIBUR community)
- Accelerate ML codes for Intel GPU systems



From GRChombo to GRTEclyn: getting ready for exascale computing

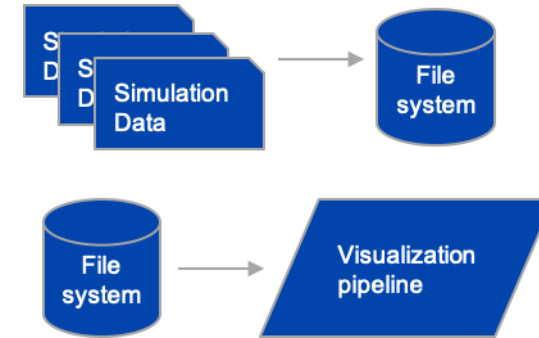
- GRChombo is an open source massively parallel AMR code that solves Einstein's equation.
- Unfortunately, no support for GPUs (or other accelerated devices)
- GRTEclyn was built on top of AMReX to take advantage of GPU capabilities.
- We are currently porting the GRChombo codebase to GRTEclyn, with the black hole binary example already completed.
- GRTEclyn is open source:
<https://github.com/GRTLCollaboration/GRTEclyn>



In-situ visualization

- Storage is costly! Reading is time consuming!
- Why not analyse simulations on the fly?
- Catalyst is an API specification developed for simulations to analyse and visualize data in-situ.
- Simulations invoke Catalyst which will wrap up the data using Conduit to pass on to Paraview.
- Currently working on raytracing with OSPRay – regions must be sliced into locally convex areas for the correct ordering when using multiple AMR levels.

Visualization in postprocessing:



With in-situ visualization:



All I want for Christmas is a...

Intel GPU testbed

- Two nodes each with:
 - 4 x Intel Xeon CPU Max 9480 (56 cores, 64GB HBM, 350W)
 - 1TB DDR5 RAM
- Three nodes each with:
 - 2 x Intel Xeon CPU Max 9480 (56 cores, 64GB HBM, 350W)
 - 4 x Intel Data Center GPU Max 1550 (128GB HBM, 600W)
 - 1TB DDR5 RAM
- 3 x Nvidia A100 40GB
- 1 x HGX node (for ML applications)
- To be made available to all ExCALIBUR researchers

