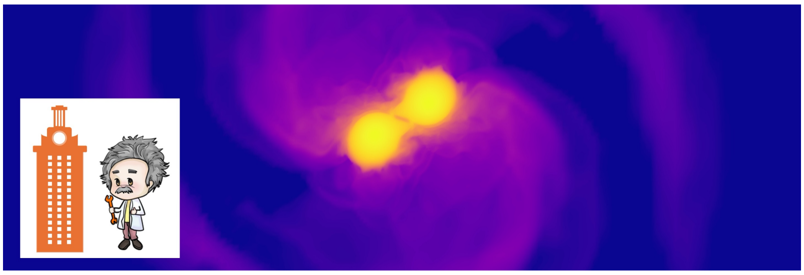

Contribution ID: **4** Type: **not specified**

(virtual) GPU-accelerated Subcycling Time Integration with the Einstein Toolkit

Tuesday 10 June 2025 10:30 (1 hour)

Adaptive Mesh Refinement (AMR) with subcycling in time enables different grid levels to advance using their own time steps, ensuring finer grids employ smaller steps for accuracy while coarser grids take larger steps to improve computational efficiency. We present the development, validation, and performance analysis of a subcycling in time algorithm implemented within the CarpetX driver in the Einstein Toolkit framework. This new approach significantly improves upon the previous subcycling implementation in the Carpet driver by achieving higher-order convergence—fourth order in time instead of second order—and enhanced scaling performance. The key innovation lies in optimizing the exchange of ghost points at refinement boundaries, limiting it to the same number as those at inter-process boundaries using dense output from coarser levels, thereby reducing computational and communication overhead compared to the implementation in Carpet, which required a larger number of buffer zones. To validate the algorithm, we first demonstrate its fourth-order convergence using a scalar wave test. We then apply the algorithm to binary black hole (BBH) simulations, confirming its robustness and accuracy in a realistic astrophysical scenario. The results show excellent agreement with the well-established LazEv code. Scaling tests on CPU (Frontera) and GPU (Vista) clusters reveal significant performance gains, with the new implementation achieving improved speed and scalability compared to the Carpet-based version.

Presenter: JI, Liwei (Rochester Institute of Technology)