Wednesday July 8 -- Microphysics

Discussion Leads: Scott Noble & Jeremy Schnittman

Brainstorming Q&A:

1. What EOSs should we be using and why?
    a. Cipo: The new implementation of the Spritz code is currently being tested with the LS220 EOS. In particular, this EOS has been extensively used in the past, both for BNS and SN calculations. How much this EOS is interesting to the eyes of the community?
    b. Jeremy: Should constrain list of EOSs to ones consistent with NICER and LIGO results. How much variation should we expect in outflow mass and kilonova properties from the list of acceptable EOSs?
    c. Francois: Currently there is so much uncertainty in our predictions that it is hard to say right now regarding kilonova signatures.
    d. Luke: Need to transition from tabular EOSs to another one after > 1 sec. of evolution; Nuclear burning occurs out to a ~1 minute, then nuclear decay sets in. Post-processing is a decent approach to handle the > 1 sec Regime and can ignore neutrinos. Decay may contribute to extra liberation of mass, but this is likely a smaller effect than other aspects.
    e. Jeremy: focus on LIGO/kilonova for the moment. Can Francois elaborate on issues?
    f. J: shouldn't expect closer systems than 170817, but will see more in future, and at different orientations. Detector improvements should help constrain these, and lock onto systems more quickly. More simulations will add understanding of appearance of differently oriented systems (see Bayesian investigations). Lots of interest in new small EM satellite detectors to catch these systems; but uncertainty in best energy bands.
    g. Manuela: subject of IAU workshop in Cape Town in February; White Paper in progress. https://kavlitransients2020.saao.ac.za/program/
    h. Federico C.: is there an EOS/set of EOSes associated with NICER research? Could use this for battery of tests, pipeline of results to compare with observation.
    i. Jeremy: NICER mostly associated with compactness constraints. Perhaps 2 decent R/M measurements? (check with Cole Miller for latest constraints)

2. Do others have experience evolving hydrostatic tori? If so, were there any "tricks" to get it to work with tabulated EOSs? -- Francois, David, Josh?
    a. Josh D: did look at it; not trivial to get it to work properly. Don't recall eventual resolution (have to go back and look); atmosphere is big pain
    b. Scott: just don't want to introduce huge transients
    c. Josh D: Jonah Miller was the main worker on dealing with that issue (Ari: will contact him) Section 3.5 of nubhlight paper discussed it.
    d. Scott: Siegel & Metzger started with torus …
3. Most codes use methods that stem from Ruffert++1996; is this method sufficiently accurate (at least in the context of leakage schemes) ? What are their limitations? How accurate are they? Should the community perform "apples to apples" comparisons?
    a. Jeremy: Where is the uncertainty coming from?
    b. Ari: rates are determined from quantum mechanics calculations and depend on the chemical potential, which depends on the EOS.
    c. Luke: There are 2nd-order effects such as opacities, which are uncertain.
    d. Francois: maybe not in the dynamical time scales;
    e. David: opacity may be important in the polar regions in rotationally supported cores, in core collapse. Because the geometry is different from non-rotating core-collapse, i.e. there is an evacuated polar jet region. This could change the composition of the material ejected outward, that is of great significance to the kilonova signal.
    f. Scott: how are people validating rates?
    g. David: just compare to Evan O'Connor's nulib code. Voila!
4. What is the "best" optical depth calculator out there?
    a. Francois Foucart: We switched to Neilsen++2014's method quickly as it is much simpler, requiring less infrastructure than the discrete ray approach used previously. Remember that the two methods yielded results to within a factor of ~2. In arxiv:1405.1121, Sec II.D, we measure that the two methods produce differences in neutrino luminosity at the 20% level, rather insignificant compared to other approximations used in the simulations.
    b. Scott: leakage isn't cutting-edge right now, so perhaps this isn't a pressing issue? Is Nielsen OK? Is it a problem that everyone's converging to the same method -- how will we know if there's a problem?
    c. Phil Chang: why are these choices made? Is leakage used because we're dealing with an optically thin medium?
    d. Scott: partly sociological/historical (what advisor/peers using). Physically, it's easy to deal with to begin with, make incremental improvements. Ultimate goal is transport. Perhaps MC is the best approximation to what "Nature does", but obviously can't capture real numbers. Different methods used (e.g. Vlasov-Boltzmann) to simplify, easier to couple with other code components.
    e. Manuela: are there other examples (besides BNS) where we think we know what's the best method/approximation? Something simple, but not too many symmetries? Any insights from experimental data -- not necessarily our area?

f. Scott: Los Alamos have experience w. Very different plasmas. Supernova has its own experience -- informs much of our expectations.
g. Josh D: Transport & opacities are separate issues. Can't answer definitively. Cooling has definite areas of applicability. Can't derive leakage scheme from transport equations; more an engineering scheme to capture transport expectations; don't really trust leakage now. Methods like M1 are nice because they can be derived from equations -- but closure choice is arbitrary. M1 behaves badly when you'd expect better. Transport community regards MC methods as "gold standard".
h. Vassili: where/how is MC used?
i. Josh: typically, MC is used to sample particles interacting with plasma, with weighting representing density of photons/neutrinos. If sampling is unbiased, then results should be too (if noisy)
j. Julian: how do you handle spectral dependence when sampling? Is the probability of selecting photon energy within a packet related to the ray's intensity? Is this updated afterwards?
k. Josh: each packet is monochromatic initially; if interaction occurs (e.g. inverse Compton), sample from the result?
l. Julian: So with this method, multiple frequencies multiplies computational cost.
m. Josh: "MOCMC" (MOCMC=Method of Characteristics Moment Closure) schemes *don't* have monochromatic packets -- each sample has a spectrum. Scattering is handled differently from true MC (Monte Carlo). See paper for details
n. Julian: danger of handling radiation pressure in low flux/mean intensity regimes. Is MOCMC an improvement here?
o. Josh: explicit MC goes unstable way before this is an issue because emission/absorption times are short compared to time-step; MOCMC definitely helps. It's done through moments (Julian: integrals generally better-behaved)
p. Josh: hit trouble with highly-relativistic shocks; samples end up in a narrow beam, destroying angular information
q. Julian: former student Brooks Kinch (now postdoc at LANL) dealing with implicit update problem with full-Compton
r. How can the Neilsen++2014 scheme be improved? Possibly through the initialization phase?
    i. David Radice: First calculate the optical depth ray by ray, and then use eikonal method for future steps.
    ii. Ian Hawke: Has anyone tried the fast sweeping method for the optical depth? : https://www.ams.org/journals/mcom/2005-74-250/S0025-5718-04-01678-3/home.html DR: I did try that, but there is a difficulty with it when the optical depth is not monotonic (for example if there is a low density pocket of plasma). The issue is that while it is true that the optical depth satisfies the eikonal equation locally, on a global scale there are some differences (it's really a free boundary problem).

s. Jeremy: is this the “right method” -- assigning a single optical depth to a single element? Are we missing important physics?
    i. Scott: Only worry if you’re tied to a leakage scheme. (Jeremy: which we are right now)
    ii. Ari: only relevant in diffusion timescale
    iii. Jeremy: perhaps when setting up testing regime, might be used to see when leakage can be trusted, and when grain of salt required
t. Federico C.: what about modeling neutrino leakage for short evolution (e.g. merger), then moving to a different method?
    i. Francois: like Daniel Siegel did?
    ii. Federico: don’t know specifics of what Daniel did, but interested in the end-stage of this project. If leakage performs badly at long timescales, are we justified in moving to a better-performing scheme later on?
    iii. Josh: can only tell by doing the right thing (i.e. transport) to begin with …
    iv. Manuela: BNS merger requires implementing full scheme in very dynamical situation; post-merger spacetime is (nearly) stationary
    v. Francois: leakage scheme during merger won’t get the instantaneous outflows right (perhaps not even the mass) -- but perhaps OK for other purposes

5. Can we put together a suite of test problems for GR rad-hydro for code comparisons, and new algorithm development?
    a. Jeremy: would be nice to have tests focused on problems of interest to us.
    b. Josh D: we have a number of tests reported, but we should do more.

6. Is MC ready for standard supercomputer usage, or is it still proof-of-principle stage? Is MC ideal? What about Boltzman/Vlasov? What are their limitations? Is one universally better everywhere? What about parallelization for MC schemes (fluid data sharing between nodes?)?
    a. Jeremy: Does MC is trivially parallelizable, but how do you handle the load balancing?
        i. Josh D.: we pass particles between the subdomains of the domain decomposition. But you are not guaranteed that you’ll have equal numbers of particles per subdomain. Thus a load balancing problem arises. Often radiation is the most expensive process, so one should really start with a radiation code and then add MHD to it. Getting radiation methods to scale well is challenging. Often get too many particles at large radii where the action is not happening, so we sometimes stop the radiation transport beyond some radius.
    b. Jeremy: Do most modern supercomputing facilities include GPUs? Can MC codes leverage GPUs well as it seems it could?
    c. Manuela: Lengthy discussion yesterday -- see summary. But many supercomputers *don’t* include GPUs, for cost reasons.

d. Jeremy: GPUs well-suited to a very narrow piece of the entire problem -- but a part responsible for a big chunk of cycles. Thus aggressive optimization might be merited. Josh, how many flops are spent on the geodesic equation?
e. Josh: mostly done in Kerr spacetime, where symmetries simplify the problem. But dynamical spacetimes should shift the need.
f. Francois: the geodesic equation is not actually a big part, compared with overhead.
g. Jeremy: but “noise” of MC methods justify cheating in precision (and thus cost)
h. Josh: MC not as good for GPUs as you’d expect, because of branching. Also a lot of “available flops” on supercomputers are currently tied up in GPUs. Recommend focusing on portability (see kokkos), rather than a particular architecture. Diversity/heterogeneity of computational resources is generally increasing.
i. Jeremy: is easily portable code synonymous with well-written code?
j. Josh: code-cleanup inherent in moving to new platform often increases speed/performance even before using kokkos
k. What's the challenge in moving to MC?
l. Josh: hit a ridiculously huge DoF problem. E.g. making the 6th-order scheme seems obvious, but ends up less efficient/slower to converge than MC.
m. David: trying spectral methods now -- in principle means can get over convergence issue
    i. Julian: how well do spectral schemes work with low optical depth problems?
    ii. David: have a spherical harmonic filter (small amount of artificial scattering) -- seems to work
    iii. Phil: how many spherical harmonics? Davide: lell=3 does well; Phil: for a single beam? David: for shadow tests, a huge number, but generally much less (radiating sphere OK with ell=3)
n. Josh: another benefit of MC: small number of “particles” needed
    i. Francois, 1-5 particles needed per cell
    ii. Josh: for semi-transparent region, ---
    iii. Francois: even in NS interior, can ignore neutrino contribution
o. Phil: summary -- MC good when optical depth not so high?
    i. David: agree for short timescales; similarly leakage is OK then
    ii. Josh: in our disk evolutions, absorption is critical at certain times
    iii. Francois: definitely needed to get Y_e right
p. Phil: MC: how important are true geodesics vs straight-line approximation?
    i. Josh: depends on path length -- not if sub-cell, but yes if it’s a large fraction of comp. Domain
    ii. Redshift may be more important

7. How much effort do you think it would be to port a MC (e.g., bhlight) code into Einstein Toolkit?

a. Josh Dolence: It's tough to say, especially because communication between processing units is necessary for handing off work to each other as the particles/rays move around.
b. Zach E.: I am the author of the particle_tracerET thorn that can evolve populations of particles, e.g., for serving as seeds to generate magnetic field lines. See no performance impact using up to 1e4 particles; it might be able to handle many more. Note that timesteps are often constrained by non GRMHD lengthscales, but rather grid length scales required to resolve dynamical black holes, so one generally need not update every CFL-constrained timestep. Uses RK4 update scheme.
    i. The thorn is part of the ETK, here:
    ii. https://bitbucket.org/zach_etienne/wvuthorns_diagnostics/src/master/particle_tracerET/
    iii. It could in principle be extended to provide live feedback into the GRMHD fields
    iv. Ian Hawke: Even GRHydro inherited particles in this sense from the Whisky implementation back in the day (although I've no idea if it still works): it's done through MoL grid arrays. Scaling to large numbers of particles is always the problem.
c. David Radice: Abdikamalov MC in the ETK: https://inspirehep.net/literature/1093716 https://iopscience.iop.org/article/10.1088/0004-637X/755/2/111/meta#apj436812s8
d. Francois Foucart: Parallelized MC cell-by-cell and found good results, not an egregious impact to the performance.

Key References

- https://github.com/lanl/nubhlight (Monte Carlo neutrino radiation code)
- https://github.com/AFD-Illinois/ebhlight (Monte Carlo electromagnetic radiation code)
- https://github.com/lanl/parthenon (AMR infrastructure)