Bart Ripperda, September 20, 2024

The thesis on radiative simulations of ULX's by Fatemeh Kayanikhoo aims to understand the effects of radiative cooling and transfer on the accretion disks of neutron stars, more specifically X-ray pulsars. The thesis contains an introductory chapter where the physics of radiation and plasma flows (magnetohydrodynamics) is presented, and an observational and theoretical overview of neutron stars and their accretion disks is given. The second chapter is a more numerically oriented exploration of how dissipation is handled in ideal and resistive magnetohydrodynamics models. The third chapter is on simulations of neutron star accretion disks with radiative magnetohydrodynamics and the accretion of matter and its rate depending on the radiation and dissipation through magnetic reconnection. In the fourth chapter this is further explored, focusing on the beaming emission from luminous accretion disks for a range of magnetic field strengths. The final chapter has concluding remarks.

The thesis starts with an exploration of magnetohydrodynamics codes and their accuracy in resolving dissipation structures. These structures are important in accretion disks of neutron stars (and black holes) as they are responsible for the heating and acceleration of particles that emit observable radiation. In this first paper, an Orszag-Tang vortex is used as a toy model for turbulence in disks. The vortex develops thin layers of current that, depending on their aspect ratio, are unstable to the plasmoid instability. When the plasmoid instability is active, the rate at which magnetic energy is converted to kinetic energy asymptotes at a high rate, which can potentially explain heating scenarios in disks. This chapter is particularly important to quantify heating through magnetic reconnection in magnetohydrodynamics simulations, which typically goes through a numerical resistivity. To quantify this effect, such "ideal" simulations need to be compared to resistive magnetohydrodynamics where an explicit resistivity can be set and resolved on the grid to control the dissipation rate. Reconnection is often invoked as a mechanism for heating (and resulting emission) in accretion disks, yet typically not resolved on the grid due to the computational expense of capturing the small-scale physics. Therefore, it is essential understand the effects of resolution, and the ideal magnetohydrodynamics assumption prior to using it for global accretion disk simulations. This chapter does exactly that and therefore has a high impact on the field, by setting the standard that every simulator should meet.

A number of questions arose when reading the chapter that may further improve the work.

1. The reconnection rate is higher than 0.1, which seems not possible in magnetohydrodynamics (MHD) reconnection. The 0.1 rate typically occurs in collisionless plasma, e.g., in particle-in-cell simulations. The asymptotic (i.e., when plasmoids formed) reconnection rate is around 0.01, when measured in terms of the outflow velocity. In relativistic plasma that would be v_in / v_out ~ v_in / v_Alfven ~ v_in / c ~ 0.01. A higher rate would indicate that there are no plasmoids, and the "faster" rate is mainly due to a large diffusion. I suspect that the reconnection rate here is either not measured correctly, potentially too far from the current sheet, or that it is not normalized correctly. The author

should explain carefully how the reconnection rate is determined and why it is different than the expected 0.01 rate.

2. In equation 10 in chapter 2, there are a number of terms missing it seems. In the relativistic current that appears on the right hand side of Ampere's law, there is a spatial charge term q*v with q=div(E) that is missing. Also the term (E*v)v in Ohm's law due to a frame transformation resulting in the correct Newtonian limit is missing. The author should explain why these terms are omitted and what the physical effect of omitting them is.
3. Based on Figure 1b, the author claims that the resistivity of 10^-4 is resolved with 512^2 or 512^3 cells. That is not convincing just based on the plot of averaged magnetic energy for different resolutions. The lines may appear close to each other but that does not signify converged results, but does show converging results.
4. A related question arises due to the claim that the numerical resistivity is below 10^-4. I think this is not necessarily correct or proven. Could it instead be the case that the numerical resistivity is larger than 10^-4 and instead the explicit resistivity that is set to 10^-4 is simply not resolved (i.e., that the numerical resistivity is larger than the explicitly resistivity). That is rather what the similar lines in figure 1 seem to indicate. Else, you would expect that eta=10^-4 would dissipate more than eta=0. In other words you show that your 10^-3 < eta_numerical < 1e-4 but not that eta_numerical > 10^-4.
   This would also result in the ideal and the resistive simulations giving similar results, just because the resistive simulation is effectively ideal (because the explicit resistivity is smaller than the numerical resistivity in the resistive simulation). This should be more thoroughly tested in resistive simulations by increasing the grid resolution until the current sheet does not further thin anymore. Based on existing literature on measuring convergence for resistive reconnection, a resistivity of 10^-4 requires resolutions of at least 2000^2 cells, or 10 cells per current sheet thickness. This can be seen from the reconnection rate argument, if plasmoids are observed then v_rec ~ 0.01 ~ eta / delta ~ delta / L where delta is the thickness and L the length of the current sheet. For a current sheet of length L=0.1 L_box where L_box is the box size, this would result in a thickness delta of 0.001 L_box. And then the requirement of having ~ 10 cells over the thickness of the sheet, would give the minimal requirement of 10000 cells over L_box. For larger current sheets of L ~ L_box, this estimate results in at least 1000 cells over the box size. That is well above what the authors claim. It would require additional tests to resolve this issue, but the author should explain and reflect this in their defense.
5. Related to the previous question, how can the numerical resistivity for 4096^2 grid cells be the same as for 512^2 grid cells. I would expect that the numerical resistivity scales (down) with increasing resolution, i.e., that the simulation becomes more and more ideal for increasing resolution?
6. The author should explain how the Alfven speed is measured in the non-relativistic simulations. How can it be of the order 1, which would make it close to the speed of light? In that case Newtonian MHD would not be applicable anymore.
7. Related, how is the magnetization sigma defined in Newtonian MHD. If sigma=B^2/rho then sigma would be the (Newtonian) Alfven speed and not the relativistic one v_A = sqrt(sigma / (sigma+1)) * c. Instead, what is the Newtonian Alfven speed? If that is used, then it becomes invalid for speeds close to the speed of light because v_A = sqrt(B^2/rho) is unbounded. The author should derive Alfven waves from Newtonian and Relativistic

MHD equations (by taking second derivatives of the momentum equation to obtain a wave equation) showing the different limits. The author should explain that a Newtonian simulation cannot set or use sigma, because it is a relativistic quantity.

8. In figure 6, can the author explain how magnetization can increase in reconnection. If there is reconnection, by definition the magnetic energy dissipates and the magnetization goes down. Similarly in decaying turbulence. Unless there is a dynamo, but that does not exist in 2D. Or unless the turbulence is driven instead of decaying, but then the author should explain what driving term is used.

The third chapter presents ideal MHD simulations of super-eddington accretion onto neutron stars. The strength of the magnetic field is varied for the dipolar field on the star, maintaining a high accretion rate. Additionally, the magnetic field is kept fixed and the accretion rate is varied. The novelty and impact of this work comes from the inclusion of radiative cooling, which is essential for the accretion disk physics in particular for ultra-luminous sources, but has not been explored in depth before this thesis.
I list a number of questions for the author below:

1. How does the author ensure that the inner region of the dipolar field, i.e., the closed zone, remains clean and without any B^phi component. What boundary conditions for the magnetic field are employed?
2. In 2D simulations, how does the author ensure that magnetic flux is kept in the disk when there is no magnetorotational stability in the toroidal direction (due to axisymmetry)?
3. What does the author expect to change if multiple orbits of a spinning star could be simulated?

The fourth chapter focuses on the beaming emission from ultaluminous sources, as simulated in the third chapter, for a range of magnetic field strengths. The author finds that despite the lower accretion rate for weaker magnetic field strengths, the apparent luminosity is higher and consistent with that of observed ultraluminous sources. This is a particularly nice and impactful finding in the understanding of ULX's. I list a number of questions regarding this topic below:

1. How could the method improve on the m1 assumption. What particularly is not captured with the m1 method and what are the downsides of the method? Which radiative effects are missing currently, and how could they be added?
2. Is m1 applicable in the regime that you study? Would there be multiple light sources/directions?
3. How does the numerical method handle the stiffness of the equations when radiation time scales are much different than Alfven time scales. Does the method use Strang split or an implicit-explicit method? Can the author explain the difference between the methods and the effect of using one or the other. Does the use of the Strang split method limit the parameter space of the radiative transfer?

The doctoral thesis constitutes an original solution to the problem of accreting ultraluminous X-ray sources, in particular accreting neutron stars. The thesis incorporates radiative effects in simulations of accreting neutron stars for the first time The thesis also quantifies the effects of numerical dissipation in reconnection and turbulence, often encountered in accretion disk simulations. The dissertation demonstrates the candidate's general theoretical knowledge in computational physics, astrophysics, neutron stars, and accretion theory. The dissertation demonstrates the candidate's ability to conduct independent scientific work, in particular simulating complex physical systems with massively parallelized codes on supercomputers, handling large data sets, understanding and interpreting the data, quantifying the accuracy of obtained numerical results, and comparison to observations.

Summing up, I consider the doctoral thesis of Fatemeh Kayanikhoo to be a valuable contribution and to meet the criteria prescribed by the law for a doctoral dissertation. Therefore, I request that this dissertation be admitted to a public defense.

Bart Ripperda, September 23, 2024