Performance and Scaling for Athena

Since ATHENA uses an unsplit, explicit update, parallelization with MPI using domain decomposition is straightforward and efficient. ATHENA has been carefully benchmarked on a variety of large-scale computing systems, including a Cray XT-3 at Sandia National Lab (RED STORM), a Cray XT-5 at NICS (KRAKEN), a Sun Magnum at TACC (RANGER), and two Dell Clusters at TACC: STAMPEDE and STAMPEDE 2. In all cases, it has shown nearly perfect scaling out to large numbers of cores, for standard 3D MHD test problems (e.g., propagation of magneto-sonic waves in three dimensions).

Here, we perform scaling tests for the numerical setup specific to the problem we propose to solve, MHD turbulence calculated in the shearing box approximation, with non-ideal MHD terms that result from low-ionization (see main document). Our setup includes specialized shearing-periodic boundary conditions along one dimension to account for Keplerian shear external to the domain, an orbital advection algorithm to solve the MHD equations separately from the Keplerian shear, a mass-conservation algorithm that invokes a global exchange of information between all processors, and a super-time-stepping approach to minimize the impact of some of the non-ideal terms on reducing the time step. We have run this particular configuration on both the Skylake (SKX) nodes and the Knights Landing (KNL) nodes on STAMPEDE 2 as we now describe.

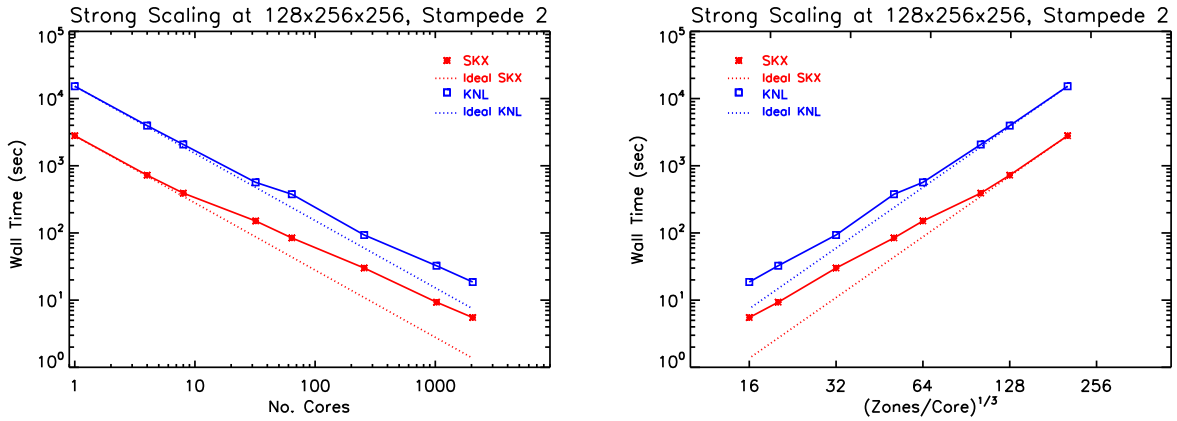

Fig. 1.— Strong scaling of ATHENA on STAMPEDE 2 for the MHD shearing box problem run at a resolution of $128 \times 256 \times 256$ zones. The left plot shows the wall time versus the number of cores, whereas the right plot shows the same data, but plotted as wall time versus zones per core. The strong scaling is close to ideal (dotted line) for all domain decompositions, departing from ideal around 32 cores for both the KNL and SKX nodes. The performance of ATHENA is better on SKX than KNL by a factor of 3–4 for ≥ 32 cores.

The strong scaling for this problem with $128 \times 256 \times 256$ total zones is shown in Fig. 1. In setting the number of cores per node, we chose the maximum available number of cores per node for all but one node (in the case that the total number processors is not evenly divisible by this maximum core number).[1] ATHENA exhibits scaling close to ideal (dotted line) for this particular

setup on both SKX and KNL nodes. The code begins to depart from ideal scaling around 32 cores on both the KNL and SKX nodes. The SKX nodes are a factor of 3–4 faster than the KNL nodes for ≥ 32 cores.

As outlined in the main document, our simulations will be carried out at a higher resolution than presented in Fig. 1. However, this resolution ($256 \times 512 \times 512$) is too large for the KNL nodes; we encountered memory errors when attempting to run this resolution on these nodes. The SKX nodes, on the other hand, have more memory than the KNL nodes and can handle this resolution. The strong scaling of this same setup for the higher resolution is shown in Fig. 2.

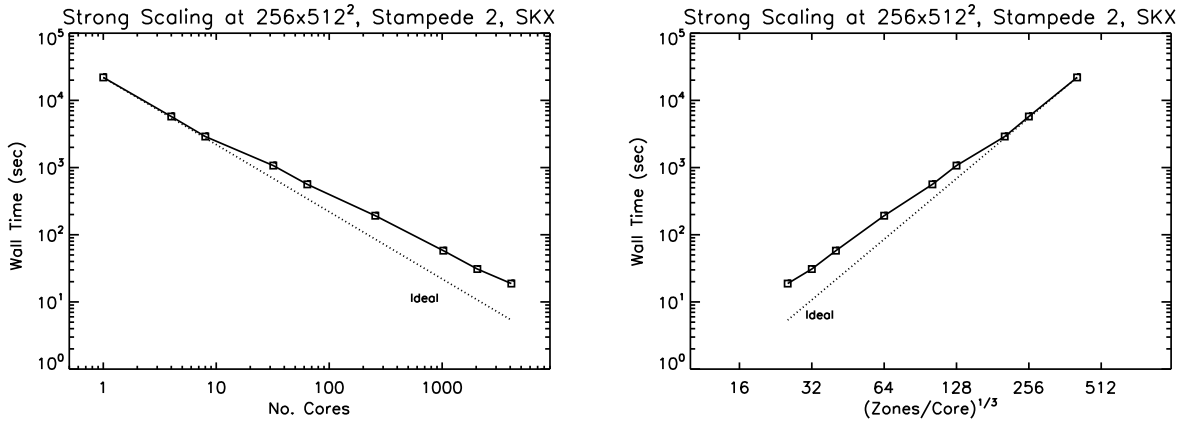

Fig. 2.— Strong scaling of Athena on Stampede 2 SKX nodes for the MHD shearing box problem run at a resolution of $256 \times 512 \times 512$ zones. The left plot shows the wall time versus the number of cores, whereas the right plot shows the same data, but plotted as wall time versus zones per core. The strong scaling is close to ideal (dotted line) for all domain decompositions, departing from ideal around 32 cores.

From this strong scaling, we take a nominal value of 16×32^2 zones per core (i.e., the point with the minimum wall time). While the strong scaling has departed from ideal scaling at this value, this is the number of zones per core in our proposed simulations, which has been chosen to increase the throughput of our runs on Stampede 2 at the expense of a slight efficiency hit. The weak scaling is shown out to 4,096 cores on SKX nodes in Fig. 3. For reference, we also include this data in the Table. Athena maintains scaling above 85% out to 4,096 cores; this excellent scaling makes Athena an ideal tool with which to carry out the high resolution simulations proposed here. Given this excellent scaling, in addition to the faster performance of Athena on SKX nodes and the higher memory of these nodes, we propose to carry out all of our simulations on the SKX nodes.

[1]We found a slight improvement in performance when choosing a constant number of cores per node. However, if we choose a constant number of cores per node (e.g., 32 cores per node on the SKX nodes) this will ultimately require more nodes to be used in our production runs (e.g., 128 SKX nodes in the case of 32 cores per node, as opposed to 86 SKX nodes when maximizing the number of cores per node). Thus, to reduce the number of SUs requested, we choose to use the maximum number of cores per node, as described here.

Table: Weak Scaling Data for MHD Shearing Box Simulations on the SKX Nodes

No. Cores	No. Nodes	Total Zone-Cycles / Second	Zone-Cycles / CPU-Second	Efficiency (Relative to 64 Cores on 2 Nodes)
64	2	6.74×10^6	1.05×10^5	1.0
256	6	2.47×10^7	0.97×10^5	0.92
512	11	5.01×10^7	0.98×10^5	0.93
1,024	22	9.78×10^7	0.95×10^5	0.91
2,048	43	1.89×10^8	0.92×10^5	0.88
4,096	86	3.67×10^8	0.90×10^5	0.85

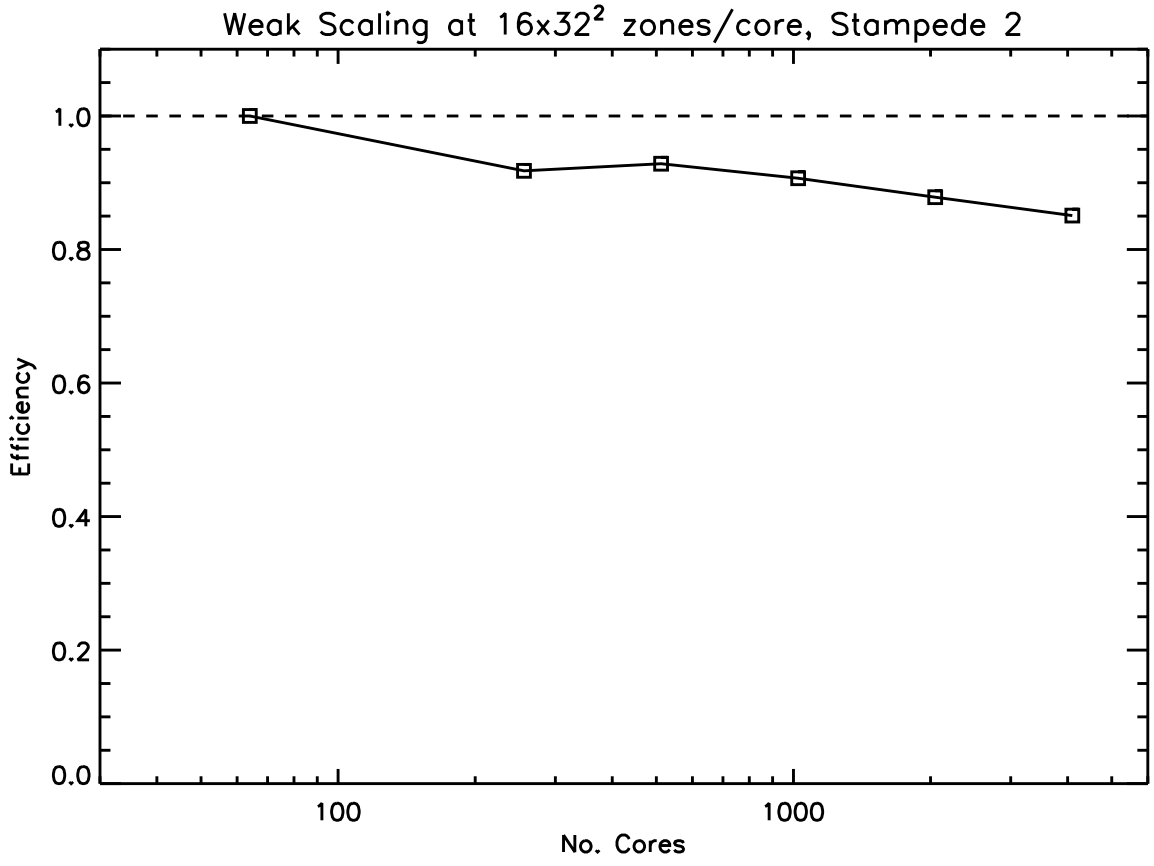

Fig. 3.— Weak scaling of Athena on Stampede 2 SKX nodes for the MHD shearing box problem at the fiducial 16×32^2 zones per core. The efficiency is measured relative to 64 cores (2 nodes) to remove the unavoidable performance hit associated with the introduction of internode communication at > 1 node. The scaling is very good, remaining better than 85% out to 4,096 cores.

Progress Report for TG-AST120062

In the sixth year (2018) of this allocation, simulations of protoplanetary disks were carried out along the two major themes: the nature of gas turbulence in protoplanetary disks as driven by magnetic fields and the formation of planetesimals via the streaming instability.

1. Turbulence in Protoplanetary Disks

The first major theme of our research in the previous allocation period was the study of turbulence in protoplanetary disks, as driven by the magnetorotational instability (MRI; Balbus & Hawley 1998). Our primary focus was to refine our understanding of the properties of this turbulence at large distances from the central star, which has direct relevance to observations made by the Atacama Large Millimeter/sub-millimeter Array (ALMA).

In previous allocations, we have made a series of predictions for the amplitude of this turbulence (Simon et al. 2013a,b, 2015a,b) and its spatial structure and compared it directly with turbulent broadening of molecular lines by ALMA (Flaherty et al. 2015; Flaherty et al. 2017, 2018). We were astonished and intrigued to find that observations constrained the turbulence to be significantly weaker at large vertical extents from the disk midplane when compared to theory; more quantitatively, observations suggest that turbulence is reduced by a factor of at least 20 in the surface layers of the outer disks in several systems.

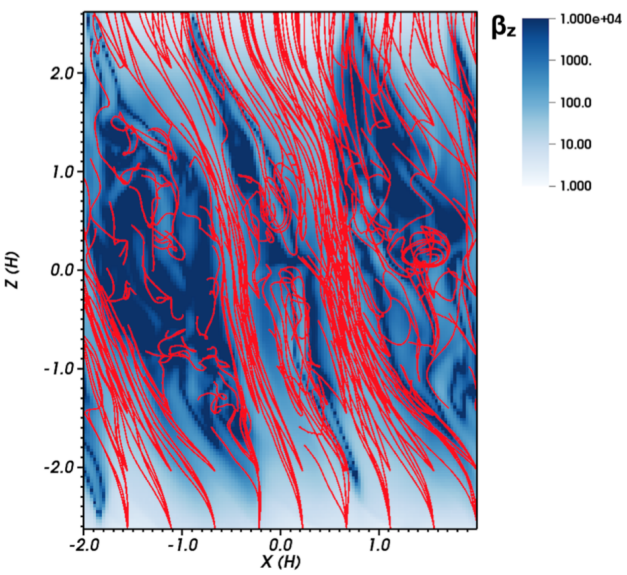

Fig. 1.— Color map of $\beta_z \equiv 2P/B_z^2$ with magnetic field lines superimposed (red lines) for a poloidal (xz) slice of a shearing box simulation. There are large scale, mostly axisymmetric, spatial variations in β_z. In large β_z regions, the magnetic field is tangled due to turbulence, whereas for $\beta_z < 10^3$, the field is relatively laminar.

Starting with the 2016-2017 allocation period and ending in the previous allocation period, we have carried out high resolution numerical simulations of MRI-driven turbulence, focusing only on the outer disk in a model for a protoplanetary disk (Simon et al. 2018). These simulations were done in the local, shearing box approximation (i.e., a co-rotating Cartesian patch of an accretion disk). Through an extensive parameter survey, we found two key results. First, for sufficiently strong background vertical magnetic fields, the gas flow remains quite turbulent, despite previous results (Simon et al. 2013b) that suggested that the flow may become laminar and that accretion be driven entirely by a magnetically launched wind. The reason for this added complexity is that while much of the magnetic field is indeed laminar, there are highly localized regions of twisted, turbulent field (see Fig. 1). Thus, even if a magnetic wind is present, so should be turbulence; thus, disks with weak turbulence should also harbor no wind-driven accretion, at least at large radial distances from the central star.

Our second key result is that only in the case of a very weak magnetic field *and* very low

ionization can there be sufficiently weak turbulence levels to be consistent with observations. The implications of these results are profound, as they suggest that the outer regions of weakly turbulent disks are likely not accreting (or not accreting very rapidly) onto the central star. This work has been written up in a paper, Simon et al. (2018), which has been published in the Astrophysical Journal.

A follow-up to this project was carried out by graduate student, Daniel Gole. In this project, we analyzed another series of MHD simulations, which again focused on varying the ionization level, but in a manner that allowed us to probe deeper into the physics of protoplanetary disk turbulence. The primary result from this work is that when far ultraviolet (FUV) photons are present, the gas dynamics at the mid-plane is largely controlled by the FUV-ionized surface layers. In particular, a large scale toroidal magnetic field is transported from the active layers to the mid-plane, which serves to shut off the MRI. This result countered our previous understanding that the mid-plane region displayed relatively weak turbulence because of the lower ionization levels there.

As a result of these works, we are now building a coherent understanding of turbulence in the outer regions of protoplanetary disks. However, to push towards a more complete and accurate model of these disks, we must understand the presence or absence of turbulence in the inner regions (at radial distances less than the orbit of Neptune from the central star). A few years ago, we carried out a preliminary study to address this question (see Simon et al. 2015b), and in the main document, we propose a higher resolution set of simulations to complete our understanding of gas dynamics at relatively small radial distances.

2. Planetesimal Formation

A large portion of the time available through our 2018 allocation was devoted to carrying out a series of very high-resolution numerical simulations of planetesimal formation. As with the MHD turbulence simulations described above, these simulations were done in the local, shearing box approximation. Additionally, these particular calculations neglected magnetic fields, but included particles and accounted for momentum exchange between gas and solid particles. Coupled with orbital motion, this momentum exchange leads to a runaway effect known as the streaming instability (Youdin & Goodman 2005); local enhancements in solid densities make it aerodynamically favorable for such enhancements to continue to grow. This process, which is reminiscent of a traffic jam on a highway, causes solid particles to continually clump together and ultimately reach a critical density after which gravitational collapse occurs and planetesimals form.

The paper that resulted from this project was the third in a series of papers within a single project, the primary goal of which has been to characterize the mass distribution of planetesimals formed via this process. In our first paper (Simon et al. 2016), we quantified the primordial mass and size distribution of planetesimals as a function of numerical resolution, initial conditions, and the relative strength of tidal forces to particle self-gravity. In the second paper (Simon et al. 2017), we changed the initial sizes of the particles and the total concentration of the particles relative to the gas (the solid-to-gas ratio).

In both cases, we found that the mass distribution can be fit with a single power law of the form $\mathrm{d}N/\mathrm{d}M_p \propto M_p^{-p}$ with $p \approx 1.6 \pm 0.1$ regardless of the numerical and physical properties

studied up to this point.

In the simulations carried out in the 2018 allocation, we expanded on this parameter survey to study the effect of the background pressure gradient of the gas (which is the primary driver of the instability, as this pressure gradient leads to sub-Keplerian orbiting gas which then interacts with the particles orbiting at Keplerian velocities to produce the streaming instability). This work, which was primarily carried out by my undergraduate, Charles Abod, has again verified that a power law can be fit to the data with $p \approx 1.6$. Furthermore, the maximum planetesimal mass is also only weakly dependent (and possibly independent) of the pressure gradient, which rules out an earlier hypothesis made by several works (Youdin & Goodman 2005; Taki et al. 2016).

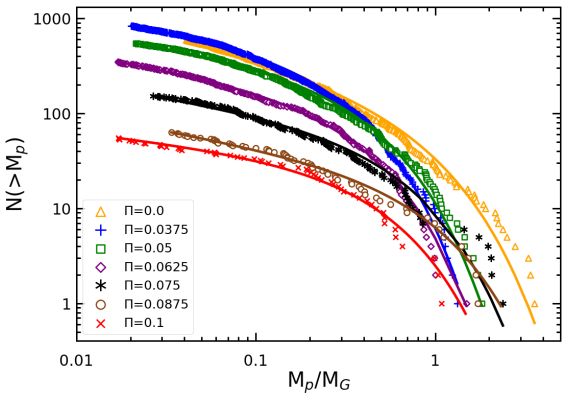

Fig. 2.— Cumulative mass function of planetesimals formed via the streaming instability for different dimensionless pressure gradients ($\Pi \equiv -0.5(\partial \ln P/\partial \ln r)(H/r)$; P is the gas pressure, H is the vertical gas scale height, and r is the radial cylindrical coordinate). Each data point represents a planetesimal from our simulation, whereas the lines of corresponding color are the best fit to an exponentially truncated power law. M_G is a mass determined by solving Toomre's instability condition while assuming the particles have zero velocity dispersion.

We also used a new clump-finding algorithm, PLAN, developed by collaborator Rixin Li to vastly improve our planetesimal statistics. This improvement led us to realize that while a single power law can indeed represent the data, a better fit to the data comes from assuming an exponentially truncated power law (as shown in the cumulative mass distribution; Fig. 2). Regardless of the exact form of the fit, the primary result from all of these studies is that the shape of the mass distribution is at best very weakly dependent on the physical and numerical parameters of the system. This important result suggests that planetesimals formed with the same distribution regardless of location and conditions in the early protosolar disk.

REFERENCES

Balbus, S. A., & Hawley, J. F. 1998, Reviews of Modern Physics, 70, 1
Flaherty, K. M., Hughes, A. M., Rosenfeld, K. A., et al. 2015, The Astrophysical Journal, 813, 99
Flaherty, K. M., Hughes, A. M., Teague, R., et al. 2018, ApJ, 856, 117
Flaherty, K. M., Hughes, A. M., Rose, S. C., et al. 2017, ApJ, 843, 150
Simon, J. B., Armitage, P. J., Li, R., & Youdin, A. N. 2016, The Astrophysical Journal, 822, 55
Simon, J. B., Armitage, P. J., Youdin, A. N., & Li, R. 2017, The Astrophysical Journal Letters, 847, L12
Simon, J. B., Bai, X.-N., Armitage, P. J., Stone, J. M., & Beckwith, K. 2013b, The Astrophysical Journal, 775, 73
Simon, J. B., Bai, X.-N., Flaherty, K. M., & Hughes, A. M. 2018, ApJ, 865, 10
Simon, J. B., Bai, X.-N., Stone, J. M., Armitage, P. J., & Beckwith, K. 2013a, The Astrophysical Journal, 764, 66
Simon, J. B., Hughes, A. M., Flaherty, K. M., Bai, X.-N., & Armitage, P. J. 2015a, The Astrophysical Journal, 808, 180
Simon, J. B., Lesur, G., Kunz, M. W., & Armitage, P. J. 2015b, Monthly Notices of the Royal Astronomical Society, 454, 1117
Taki, T., Fujimoto, M., & Ida, S. 2016, Astronomy and Astrophysics, 591, A86
Youdin, A. N., & Goodman, J. 2005, The Astrophysical Journal, 620, 459

Numerical Simulations of Protoplanetary Disks

Jacob B. Simon (PI), Philip J. Armitage

Summary

We request a total of **476, 372 SUs** on the TACC machine Stampede 2 SKX nodes to run and analyze numerical simulations of magnetohydrodynamic turbulence in the planet forming region of protoplanetary disks. We will use the well-tested Athena code to simulate a series of local, co-rotating radial regions of a model disk. To store the large datasets that will result from these calculations, we request $\mathbf{1.744 \times 10^4}$ **GB** of archival storage on the TACC Ranch system.

1. Introduction and Scientific Background

Protoplanetary disks are the precursors to planetary systems, whether it be our own solar system or the thousands of other planetary systems discovered in the past two decades, and nowhere is there a more important connection between these planetary systems and their parent disks as in the dynamics of the disk gas itself. In particular, how angular momentum is removed from the gas, allowing it to accrete onto the central star, is key to determining the environment in which planets are born, for example through determining the time evolution of the solid and gas surface density profiles.

However, beyond the degree of angular momentum transport, *how* this momentum is transported is also crucially important. It has long been theorized that disk gas will be turbulent due to a powerful instability present in magnetized, orbiting gas: the magnetorotational instability (Balbus & Hawley 1998). Such turbulence will have a significant impact on how dust grains settle to the mid-plane (Fromang & Papaloizou 2006) and grow into larger bodies (Youdin & Lithwick 2007; Birnstiel et al. 2010), and this turbulence can also influence other stages in the planet formation process, such as the concentration of particles to form planetesimals (Cuzzi et al. 2008; Johansen et al. 2009; Simon & Armitage 2014) and the migration of planetary bodies (Nelson & Papaloizou 2004; Lubow & Ida 2011; Baruteau et al. 2011; Paardekooper et al. 2011).

Turbulence clearly plays an important role in the planet formation process, but what is less clear is the amplitude of turbulent fluctuations, which again, will directly impact the planet formation process. Protoplanetary disks, with their relatively cold temperatures, are plagued by low ionization levels that persist throughout the disk, especially near the mid-plane region. In the presence of such weak ionization, the magnetic field can effectively slip through the gas; this occurs both at high densities, relatively close to the central star, where Ohmic diffusion dominates and at low densities, far from the star where ambipolar diffusion dominates. Ultimately these effects can significantly reduce the efficacy of the MRI in generating turbulence (Gammie 1996; Bai & Stone 2011; Simon et al. 2013a,b).

Within the bulk of the planet forming region of the disk, typically from 1 AU (Earth's orbit) to 30 AU (Neptune's orbit), another low-ionization effect comes into play: the Hall effect. While

not necessarily a "diffusive" effect, the Hall effect can drastically alter the strength and nature of angular momentum transport in protoplanetary disks. More specifically, the exact behavior of the gas within the Hall-dominated regime depends on the presence and orientation of any large scale vertical magnetic field (vertical in the cylindrical coordinate sense; the field threads the disk orthogonal to the disk plane). In particular, if the vertical field is aligned with the angular momentum vector of the disk, angular momentum transport is enhanced drastically. Conversely, in the anti-aligned case, transport is either reduced significantly (Simon et al. 2015) or quenched altogether (Kunz & Lesur 2013; Lesur et al. 2014; Bai 2015). Furthermore, in the aligned magnetic field case, the magnetic field becomes highly *laminar*, whereas in the anti-aligned case, it is less laminar and forms structure through intermittent bursts (Simon et al. 2015).

Given these complexities, there remain two big questions to address: is the gas turbulent and if so, what is the amplitude and spatial and temporal structure of that turbulence? Even in the laminar magnetic field case (with the field aligned with the disk angular momentum vector), preliminary results (Bai 2015) have shown that there is some turbulence present, and our previous work indicated (at low-to-moderate resolutions) that in the anti-aligned case, turbulence may be present in an intermittent fashion.

The implications of these results for planet formation could be profound, as they suggest that planet formation itself may crucially depend on the presence and orientation of large scale magnetic fields. If 50% of disks have a field aligned with the disk angular momentum vector, then 50% of exoplanet systems may display a fundamentally different structure than the 50% of disks with an anti-aligned field. Indeed, there has been very preliminary work to suggest that this may be the reason that half of all planetary systems around Sun-like stars have planets orbiting very close to their star (Simon 2016).

It is the primary goal of this proposal to understand and quantify the exact nature of MHD turbulence within the Hall-dominated, planet forming region of protoplanetary disks. In what follows, we we will describe our progress to-date in understanding disk turbulence and how this further motivates the project as we have outlined here. We will then present our computational algorithm and justify the resources that we request.

2. Progress to Date and Motivation for Future Work

In a previous allocation period (2014–2015), we examined the influence of the Hall effect on the gas dynamics and turbulence through a series of low-to-moderate resolution simulations that were centered on several disk radii (see Simon et al. 2015). These simulations were carried out in the shearing box approximation (see Fig. 1), which allowed us to concentrate more grid zones on smaller scales (at the expense of scales significantly larger than the vertical gas scale height). Our primary result, as shown in Fig. 2, was that in the case of a large scale vertical magnetic field aligned with the disk angular momentum vector, large scale highly laminar magnetic fields are generated, which transports angular momentum radially outward (thus allowing the gas to accrete) at an enhanced rate. In the anti-aligned case, a more tangled, turbulent magnetic field is generated intermittently, and angular momentum is transported at a reduced amplitude, yet in a highly variable fashion.

We determined that this behavior was actually the result of an instability, which while marginally related to the MRI, is distinctly different: the Hall-shear instability (HSI; Kunz 2008).

However, one significant limitation of that work was the low resolution that was employed, which at the time was necessary due to performance limitations with the code used at that time (which was not the code we propose to use here). Thus, we require high resolution simulations not only to verify our results, but to resolve HSI unstable modes that may be present (and could potentially generate turbulence), but which were not resolved in these earlier calculations. One piece of evidence in support of this was the finding that the turbulent "bursts" disappeared when ambipolar diffusion was amplified (Simon et al. 2015). Thus, if diffusion can play such a major role, higher resolution calculations, which will have lower *numerical* diffusion, may resolve important scales that will lead to dramatically different behavior.

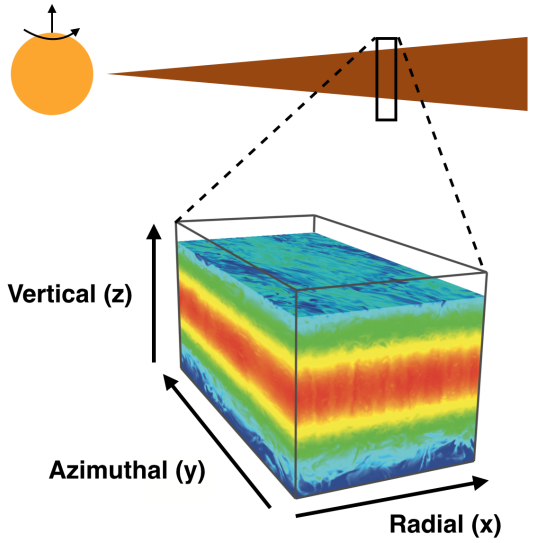

Fig. 1.— Graphical representation of a shearing box, which is a co-rotating disk patch, sufficiently small so as to treat the domain in Cartesian coordinates. A shear profile is set up along the radial dimension to account for differential rotation.

In the allocation prior to the one in which these results were uncovered (i.e., the 2013–2014 allocation period), we carried out a large number of numerical simulations to study the nature of turbulence in the outer regions of protoplanetary disks (i.e., at distances much further than the orbit of Pluto). These calculations were done with the specific goal of predicting the degree of turbulent molecular line broadening to be observed with the Atacama Large Millimeter/sub-millimeter Array (ALMA). Ultimately, ALMA observations found that turbulence in the outer disk is *very* weak compared to these early predictions, which then motivated us to explore parameter space further (done in the 2016-2017 and 2018 allocation periods; see Simon et al. 2018) to determine what was responsible for the discrepancy between our first predictions and observational constraints.

As a result of these outer disk simulations, our picture of protoplanetary disk accretion is beginning to undergo a paradigm shift. However, in order to further complete our picture of protoplanetary disk accretion, we must understand the nature of turbulence in the inner disk and make similar comparisons with observations, as we did with ALMA. Indeed, with the advent of JWST and thirty meter class telescopes, such observations will soon be within reach.

Finally, we have begun (with our previous allocation as briefly described in our progress report) to look into the influence of turbulence on planetesimal formation. In our calculations so far, we have employed a simple forcing function at the largest scales of our domain (i.e., a superposition of randomly phased sinusoidal functions driven at an amplitude controlled by the user); this approach generates turbulence at the largest scales, which then cascades to smaller and smaller scales. However, one important aspect of future work will be to control the amplitude and structure of the turbulence as a function of spatial scale. The high resolution simulations we propose to carry out

here will provide the dynamic range necessary to quantify the power spectra of MHD turbulence within the planet forming regions of disks, which we can then use in future studies to determine the role of turbulence in planetesimal formation.

Thus, for a number of reasons, largely motivated by advances in observational facilities and our previous accomplishments in this field, it is important to carry out high resolution numerical simulations of MHD driven accretion flows in the inner, planet forming regions of protoplanetary disks and examine the properties of the resulting turbulence. These are the simulations we propose for the 2019-2020 allocation period.

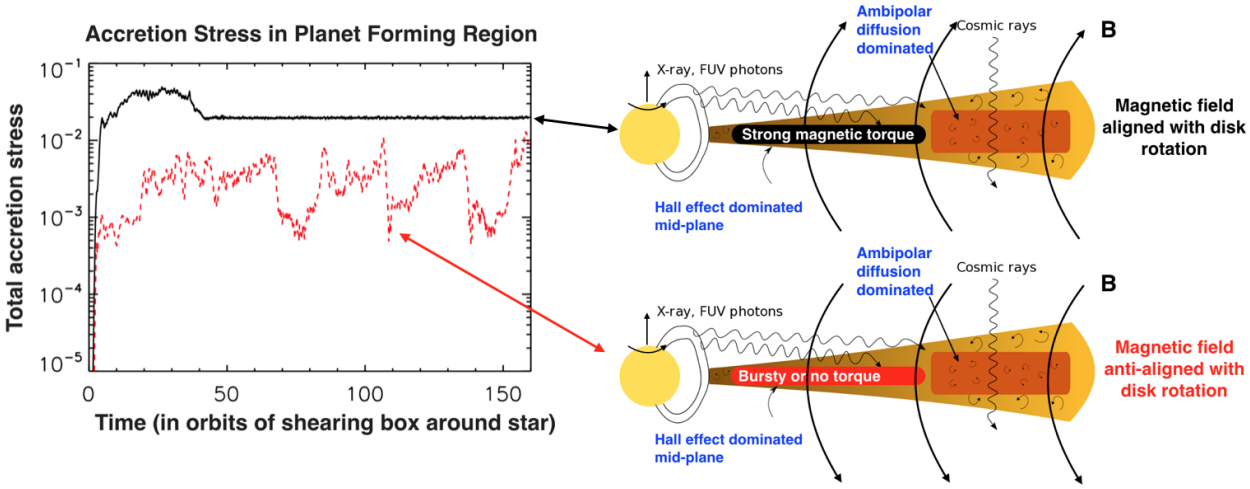

Fig. 2.— Emerging paradigm for magnetically driven accretion in protoplanetary disks, in which Ohmic and ambipolar diffusion, and the Hall effect strongly influence gas dynamics. Within the planet forming region, the Hall effect can create a strong magnetic torque near the mid-plane in the presence of a large scale vertical field that is aligned with the disk rotation. In the anti-aligned case, the Hall effect can induce bursts (for sufficiently weak diffusion). We will use numerical MHD calculations to study the gas dynamics in disks.

3. Proposed Computational Methods

All of our proposed local simulations will be carried out with the publicly available MHD code, Athena[1], which has been extensively tested and used for most of our shearing box calculations done to-date (Gardiner & Stone 2005, 2008; Stone et al. 2008; Stone & Gardiner 2010)[2]. Athena is a second-order accurate Godunov flux-conservative code for solving the equations of MHD; descriptions of the Athena algorithms appear in Gardiner & Stone (2005) and Gardiner & Stone (2008) with a complete description of the code and test problems presented in Stone et al. (2008). In the configuration utilized for our proposed simulations, Athena uses the dimensionally unsplit corner transport upwind (CTU) method of Colella (1990) coupled with the third-order in space piecewise parabolic method (PPM) of Colella & Woodward (1984) and a constrained transport (CT; Evans & Hawley 1988) algorithm for preserving the $\nabla \cdot \boldsymbol{B} = 0$ constraint. The HLLD Riemann solver is

used to calculate the numerical fluxes (Miyoshi & Kusano 2005; Mignone 2007).

For this project, we numerically solve the equations of non-ideal, compressible MHD in the shearing box approximation, in which the domain is a co-rotating disk patch, small enough in size (compared to its radius from the central star) for curvature terms to be neglected (see Fig. 1):

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}) = 0, \tag{1}$$

$$\frac{\partial \rho \boldsymbol{v}}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}\boldsymbol{v} - \boldsymbol{B}\boldsymbol{B}) + \nabla \left(P + \frac{1}{2} B^2 \right) = 2q\rho\Omega^2 \boldsymbol{x} - \rho\Omega^2 \boldsymbol{z} - 2\boldsymbol{\Omega} \times \rho \boldsymbol{v}, \tag{2}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} - \nabla \times (\boldsymbol{v} \times \boldsymbol{B}) = \nabla \times \left[\frac{(\boldsymbol{J} \times \boldsymbol{B}) \times \boldsymbol{B}}{\gamma \rho_i \rho} - \frac{c}{e n_e} (\boldsymbol{J} \times \boldsymbol{B}) - \eta \boldsymbol{J} \right], \tag{3}$$

where ρ is the mass density, $\rho\boldsymbol{v}$ is the momentum density, $\boldsymbol{B}$ is the magnetic field, and P is the gas pressure. The shear parameter is $q = -d\ln\Omega/d\ln R$, where Ω is the angular velocity at radius R away from the disk center. We use $q = 3/2$, appropriate for a Keplerian disk.

From left to right, the source terms in equation (2) correspond to radial tidal forces (gravity and centrifugal), vertical gravity, and the Coriolis force. The source term in equation (3) includes ambipolar diffusion (first term in the brackets), the Hall effect (second term) and Ohmic diffusion (third term). $\boldsymbol{J} = \nabla \times \boldsymbol{B}$ is the current density, ρ_i is the mass density of the ions, γ is the drag coefficient between ions and neutral species, c is the speed of light, e is the absolute value of the electron charge, and n_e is the electron number density. The Ohmic diffusion term contains the Ohmic resistivity η, which is determined directly by the gas ionization fraction. Note that in our system of units, the magnetic permeability $\mu = 1$, and a factor of $1/\sqrt{4\pi}$ has been subsumed into the definition of B. We will assume an isothermal equation of state.

Additions to the Athena algorithm relevant to the shearing box approximation can be found in Stone & Gardiner (2010) and the Appendix of Simon et al. (2011). Briefly, we utilize Crank-Nicholson differencing to conserve epicyclic motion exactly and orbital advection to subtract off the background shear flow. This orbital advection vastly improves the accuracy and speed of the code (Stone & Gardiner 2010). The azimuthal boundary conditions are strictly periodic, whereas the radial boundaries are shearing periodic (Hawley et al. 1995; Simon et al. 2011). The vertical boundaries are the modified outflow boundary conditions described in Simon et al. (2013b).

Ohmic and ambipolar diffusion are evolved using an operator split, finite difference integration of each diffusion term coupled with a super-timestepping technique (Choi et al. 2009; Bai 2011; Simon et al. 2013a,b). Since the ambipolar diffusion time step is inversely proportional to the square of the Alfvén speed, the time step can be extremely (and prohibitively) small as a result of magnetic pressure dominating over gas pressure far from the disk mid-plane. This problem has been overcome in our previous calculations with the use of super-timestepping. Super-timestepping works by integrating the diffusion terms (including both Ohmic and ambipolar diffuision) over a

[2]A second generation code, Athena++ has now been written, which is faster than Athena. However, it does not yet include all of the algorithmic additions that are required to do our proposed simulations.

[2]The Athena code and a repository of test problems are maintained online at https://github.com/PrincetonUniversity/Athena-Cversion/.

pre-determined number of sub steps, which start off to be large (certainly larger than the diffusion limited time step) and then become progressively smaller with each iteration. Mathematically, the super time step is

$$\Delta t_{\rm STS} = \sum_{j=1}^{n} \Delta\tau_j, \tag{4}$$

where

$$\Delta\tau_j = \Delta t_{\rm diff} \left[(\nu - 1)\cos\left(\frac{2j-1}{n}\frac{\pi}{2}\right) + \nu + 1 \right]^{-1}. \tag{5}$$

Here, $\Delta\tau_j$ is the jth sub step, n is the number of iterations within a super time step, $\Delta t_{\rm diff}$ is the diffusion time step (most likely set by ambipolar diffusion), and $0 < \nu < 1$ is chosen by the user to maximize optimality and stability of performance. As $\nu \to 0$, $\Delta t_{\rm STS} \to n^2 \Delta t_{\rm diff}$; asymptotically, this approach is n times faster than the standard approach (Choi et al. 2009). We have implemented and tested this method in Athena. It has resulted in a significant speed up for our ambipolar diffusion calculations to-date; the resulting time step is roughly equivalent to the ideal MHD time step (i.e., the time step that would be used if there were no diffusion).

The Hall effect is implemented as described in Bai (2014). Briefly, the Hall term is evolved both in a directionally split and operator split manner. The updates are done on the EMF terms (the terms inside the brackets in Equation 3) so as to maintain the $\nabla \cdot \boldsymbol{B} = 0$ constraint (this is done for all non-ideal terms in general). As with the diffusion terms, the Hall term has been well-tested and employed in a number of studies (e.g., Bai 2014, 2015). Since the Hall effect is a hyperbolic term (and the diffusive terms are parabolic), we cannot use STS to update the Hall term. Therefore, we use the minimum of the ideal MHD time step and the Hall time step, and then if the diffusive time step (described above) is smaller than the minimum of the ideal and Hall time steps, we employ STS.

4. Proposed Simulations and Justification of Requested Resources

We propose to carry out a series of local, shearing box simulations partially (though certainly not exclusively) motivated as an extension to the simulations in Simon et al. (2015), at much needed higher resolution. We thus use those simulations as a basis for calculating required integration times and the required parameter space. In that work, it was found that there are three main parameters controlling the degree of angular momentum transport (also supported by the results shown in Fig. 3):

- The magnetic field strength, characterized by the ratio of the gas pressure to the magnetic pressure associated with the net vertical magnetic field, $\beta_z \equiv 2P/B_z^2$. In Simon et al. (2015), we found that for $\beta_z = 10^3$-10^5, the accretion rates were within observational constraints (Hartmann et al. 1998), but that the accretion rates (and consequently turbulence levels) varied by up to an order of magnitude depending on the β_z value. Thus, in order to capture the

full range in turbulence levels, while maintaining consistency with observational constraints, we propose to run simulations at $\beta_z = 10^3, 10^4, 10^5$.

- The orientation of the net vertical field with respect to the disk angular momentum vector, quantified via an "orientation" function $O_B \equiv \text{sign}(\boldsymbol{\Omega} \cdot \boldsymbol{B})$. The value of O_B can be either -1 or 1, but we must explore each field orientation for every β_z value described above, since, as shown in Simon et al. (2015), both parameters play a large role in determining the strength of turbulence

- The radial location of the center of the shearing box. In Simon et al. (2015), there were three radii chosen interior to the orbit of Neptune: 1 AU, 5 AU, and 10 AU (where AU, the astronomical unit, is the distance from the Sun to the Earth).

As described above, it is essential that for every β_z value, we explore both values of O_B; this combination equates to six simulations. However, to keep parameter space (and the associated expense) to a minimum, we will choose one fiducial value for β_z and O_B and vary the radius over the three values, 1 AU, 5 AU, and 10 AU. This equates to a total of 8 simulations.

For consistency, the setup of each simulation will be very similar to the simulations of Simon et al. (2015). Briefly, the ionization structure of the disk in each simulation will be calculated via a chemical network table (Bai & Stone 2011, 2013), which is accessible to the PI. Each local simulation will be $4H$ radially, $8H$ azimuthally, and $8H$ vertically, where H is the vertical disk scale height. We will choose the radial locations based on the minimum mass solar nebula model (MMSN; Hayashi 1981).[3]

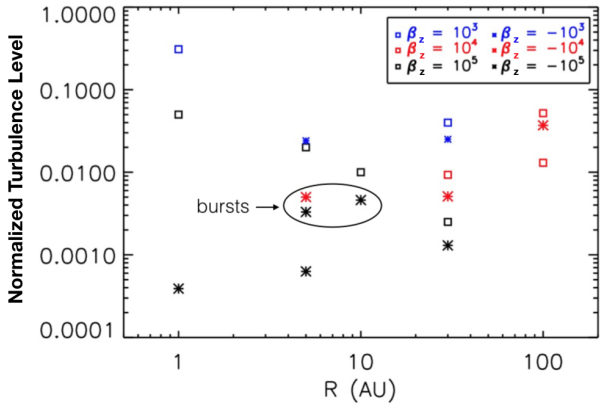

Fig. 3.— Normalized turbulence level (turbulent energy divided by the initial gas pressure at the disk mid-plane) for all of the simulations in Simon et al. (2015). The color denotes the value of the net vertical magnetic field strength, defined via β_z and the symbol denotes the orientation of the vertical field with respect to the disk angular momentum vector (In the main text, this is O_B, but has been attached as a sign on β_z here for simplicity). The turbulence level depends strongly on radial position, strength of the vertical field, and the orientation of that field.

One major (and necessary) difference between these proposed simulations and those carried out in Simon et al. (2015) is the numerical resolution. Here, we propose a resolution of 64 grid zones per H. Not only is this a *substantial* increase in resolution over that of the prior work (thus making these calculations truly state-of-the-art), but it is necessary resolution to achieve our goals. Previous works have shown that 36 zones per H is the bare minimum resolution required to achieve some form of convergence (see e.g., Simon et al. 2011). However, even higher resolution is required if we want to determine the properties of

[3]We realize that there are better models than the MMSN. However, in order to maximize the scientific value of our simulations, we wish to maintain consistency with our previous results and change as few variables as possible. Improved disk models will be used in future work.

turbulence on scales $< 0.1H$, which is where planetesimal formation takes place. Since this is one of our primary goals in carrying out these calculations, we need to carry out as high resolution as possible (without asking for an outrageously large number of Node-hours of course). Thus, given our domain size as outlined above and a requirement of 64 zones per H, each of our eight proposed simulations will have a resolution of $N_x \times N_y \times N_z = 256 \times 512 \times 512$.

With this information in hand, along with the benchmarking data presented in the Performance and Scaling Documentation, we now calculate the number of SUs required for each run and thus the total number of SUs. In doing our benchmarking runs, we calculated that the average time step value (in code units) for our production level simulations is $\Delta t = 3.95 \times 10^{-2}$. In units of orbital periods (of the shearing box around the central star), $\Delta t = 6.29 \times 10^{-6}$ orbits. From Simon et al. (2015), it was generally required that at least 100 orbital periods be integrated to capture one or two episodes of low-to-high turbulence transitions (as shown in the left panel of Fig. 2). Even in the case of a turbulent steady state (i.e., no large variability), 100 orbits is a reasonable value over which to time average results and thus reduce statistical noise. Ideally, a longer integration period should be used (Fig. 2 shows a range of ~ 150 orbits), but to minimize computational expense (and number of days in the STAMPEDE 2 queue), we think a minimum of 100 orbits will suffice. We re-emphasize that while these timescales are empirically obtained (as opposed to calculated through analytical calculations), we have done the necessary groundwork to ensure that they are reasonable values by running short test simulations at the proposed resolution and by (re-)examining the simulations in Simon et al. (2015).

Given this information, the total number of time steps is $N_t = 1.59 \times 10^7$. To maximize throughput on STAMPEDE 2, we propose to run our simulations on 4,096 cores. With our chosen resolution of $N_x \times N_y \times N_z = 256 \times 512 \times 512$, and a Zone-Cycles/CPU-Second of 9.0×10^4 (see Performance and Scaling Documentation), we now calculate the time required for each simulation as follows:

$$\text{Run Time (hours)} = \frac{1}{3600} \left(N_t N_x N_y N_z\right) / \left(N_{\text{cores}} \frac{\text{zone-cycle}}{\text{cpu-sec}}\right) \tag{6}$$

$$\#\text{SUs} = \text{Run Time (hours)} \times N_{\text{Nodes}} \tag{7}$$

where N_{cores} is the number of cores, and the number of nodes, $N_{\text{Nodes}} = 86$ is chosen by dividing 4,096 by the number of cores per SKX node (48) and then rounding up to the nearest integer. From these equations, we calculate that **each run will require 804 hours of run time, which equates to $\approx$ 34 days, and 69,147 Node-hours.**

For 8 simulations, the total expense equates to 553,172 SUs. However, as described in more detail below, our group has access to a small amount of computing time associated with one of the grants supporting this proposal. This computing time is designed for a related set of simulations and thus we feel it appropriate that we subtract it off of our request here. This time, while on a NASA machine, is the equivalent of 840,000 CPU-Hours on the *original Stampede.* Taking a factor of 0.32 (based on emails from XSEDE during the transition from STAMPEDE to STAMPEDE 2 KNL nodes) conversion to STAMPEDE 2 and an additional factor of ≈ 3.5 in converting from KNL nodes

to SKX nodes, this extra computing time equates to 76,800 SUs. Subtracting this from 553,172 SUs we arrive at our **total request of 476,372 SUs**. The proposed simulations along with this adjustment are presented in Table 4.

Table 1: Proposed Simulations

R (AU)	β_z	O_B	# zones	# cores	$\frac{\text{zones}}{\text{core}}$	N_t	Days	**SUs**
5	10^3	1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
5	10^3	-1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
5	10^4	1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
5	10^4	-1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
5	10^5	1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
5	10^5	-1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
1	10^4	1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
10	10^4	1	256×512^2	4,096	16×32^2	1.59×10^7	33.5	$\mathbf{69,147}$
							Total:	553,172 SUs
							Other resources:	-76,800 SUs
							Final Total:	**476,372 SUs**

5. Computational Resources

Access to the TACC system Stampede 2 is *essential* for the success of our research project. Locally, our group does have access to a moderate sized high performance computing system, Summit. This machine has 9,120 general computing CPUs (plus GPU capability on other nodes) and is operated through the University of Colorado, Boulder. However, using this machine for the simulations proposed here is problematic for two reasons. First, the number of cores required for a single shearing box simulation is a substantial fraction of the maximum number of cores on Summit. Thus, we could never expect this system to permit regular access to the 4,096 cores required for our calculations. Second, in consulting with the Research Computing division at the University of Colorado, it has become clear that we can acquire at most one million CPU-hours on this machine per year *per research group*, significantly less than what we require to carry out our proposed simulations.

Furthermore, our group does have access to computing time on NASA supercomputers through two grants. One allocation, equating to approximately 1.3 million CPU-hours is allocated for an entirely different project and thus cannot be used here. In the second allocation, which is through the grant titled "Planetesimal Formation in the Protosolar Disk: The Influence of Turbulence", we do have access to 840,000 CPU-hours for a set of runs very similar to those proposed here. As such, we feel that it is only appropriate to subtract this time from our request, and we have done so in calculating the total number of SUs above. However, as demonstrated by the numbers, this small number of CPU-Hours is far too small to account for all of the parameters that we have outlined here, and thus we are making a request for the additional time on Stampede 2.

The proposed simulations will produce rather large datasets. Based upon the runs described in Simon et al. (2018) (which will be the most analogous in terms of code output) the high resolution 256×512^2 cells take up 1.8 GB of disk space for each VTK file data dump. Generally, we will output 1000 VTK files (for sufficient temporal resolution) per simulation, equating to 1.8 TB of

space. The checkpoint files are 3.8 GB in size each. However, we will output these at a factor of 10 less frequent than the VTK files, equating to 380 GB total. We will also output other ASCII formatted files, but the sizes of these files are very small and are thus negligible in calculating the total the storage request. Thus, with eight datasets **the total amount of disk space required for the entire suite of simulations is 1.744×10^4 GB**. This amount of data is too large to transfer to and house on local storage systems. Thus, we request that archival storage is made available to house this data on the TACC archival storage system Ranch.

Since our simulation data will already be on Stampede 2, we will use the SKX processors to visualize and analyze our simulation data sets. This system already contains access to the tools necessary for us to analyze our data, namely IDL and VisIt; both of these tools have served us well previously in analyzing our large datasets and generating reduced forms of data, figures, etc.. In the past year, we have used a non-zero, but quite small number of SUs to analyze our datasets. We expect similar usage for this allocation period, and we thus anticipate a number that is significantly less (and thus within the noise) of our request for running the proposed simulations. Therefore, it will not change our total request for time on Stampede 2. Upon completion of analyzing our simulation data, we will transfer the entire dataset to Ranch.

6. Research Plan

Setting up and executing the proposed simulations will be carried out by the PI, Dr. Jacob Simon. Dr. Simon has extensive experience with the Athena code, with shearing box simulations, and with running large scale computations on high performance computing systems. He has been PI on, co-I on, and involved in various projects utilizing the XSEDE systems. In particular, he was PI on grant TG-AST120062 (for which this proposal is a renewal), co-I on grant TG-AST090106, and has run numerous simulations under the grant TG-MCA95C003 with Prof. John Hawley as PI.

PI Simon will be responsible for running the entire simulation suite and analyzing the results. He will also be responsible for managing the large datasets produced and transferring them to Ranch. Co-I Prof. Philip Armitage will assist in the interpretation of our results and its relevance to issues of planet formation and general properties of protoplanetary disks, areas in which he is an expert.

The first objective of our project will be to run one of the shearing box simulations (e.g., the R = 5AU, $\beta_z = 10^4$, $O_B = 1$ calculation) for one 24 hour period. The resulting data will then be analyzed using the tools available on Stampede 2. At this point, our goal will be to verify that the simulation data looks sensible, test our analysis pipeline on the dataset, and consider any additional diagnostics that may be useful in understanding the streaming instability. After this initial analysis, we will begin the process of running all of our simulations to completion. This will consist of running several simulations in parallel with each other in order to maximize the throughput and efficiently carry out our entire simulation suite. However, as each simulation comes to completion, we will analyze the data and then transfer the data to Ranch for long term storage. We will then write up our results and submit them to peer reviewed journals as appropriate.

7. Grant Support

This research will be supported by two NASA grants. The first is an Astrophysical Theory Program (ATP) grant, "Magnetic Fields and Self-Gravity in Early Protostellar Disks", the PI of which is Armitage. This grant will be used to fund PI Simon, a Senior Research Associate at the University of Colorado, at 50% for the duration of this XSEDE request. One of the primary science objectives of this grant is to study the advection and diffusion of large scale magnetic fields. To understand such processes, a detailed understanding of the turbulence in protoplanetary disks is required, and as such, this grant is relevant to the simulations proposed here. The second grant is an Emerging Worlds grant, "Planetesimal Formation in the Protosolar Disk: The Influence of Turbulence", the PI of which is Simon. This grant will be used to fund Simon at an additional 50% for the duration of this XSEDE request. While this grant has not yet been administered to the University of Colorado (forcing us to enter some requested data as "not yet available" in the online form), the funds will become available later in the spring. The first task of this proposal is to characterize the nature of turbulence in the Hall-effect-dominated planet forming regions of disks, and as such is directly relevant to our proposed simulations.

8. Total Allocation Request

Our total request for the study of MHD turbulence in the planet forming regions of disks is:

476,372 SUs on Stampede 2

1.744×10^4 GB of storage on Ranch

REFERENCES

Bai, X.-N. 2011, The Astrophysical Journal, 739, 50

—. 2014, The Astrophysical Journal, 791, 137

—. 2015, The Astrophysical Journal, 798, 84

Bai, X.-N., & Stone, J. M. 2011, The Astrophysical Journal, 736, 144

—. 2013, The Astrophysical Journal, 769, 76

Balbus, S. A., & Hawley, J. F. 1998, Reviews of Modern Physics, 70, 1

Baruteau, C., Fromang, S., Nelson, R. P., & Masset, F. 2011, Astronomy and Astrophysics, 533, A84

Birnstiel, T., Dullemond, C. P., & Brauer, F. 2010, Astronomy and Astrophysics

Choi, E., Kim, J., & Wiita, P. J. 2009, The Astrophysical Journal Supplement, 181, 413

Colella, P. 1990, JCP, 87, 171

Colella, P., & Woodward, P. R. 1984, JCP, 54, 174

Cuzzi, J. N., Hogan, R. C., & Shariff, K. 2008, The Astrophysical Journal, 687, 1432

Evans, C. R., & Hawley, J. F. 1988, ApJ, 332, 659

Fromang, S., & Papaloizou, J. 2006, A&A, 452, 751

Gammie, C. F. 1996, ApJ, 457, 355

Gardiner, T. A., & Stone, J. M. 2005, JCP, 205, 509

—. 2008, JCP, 227, 4123

Hartmann, L., Calvet, N., Gullbring, E., & D'Alessio, P. 1998, The Astrophysical Journal, 495, 385

Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1995, ApJ, 440, 742

Hayashi, C. 1981, Progress of Theoretical Physics Supplement, 70, 35

Johansen, A., Youdin, A., & Klahr, H. 2009, The Astrophysical Journal, 697, 1269

Kunz, M. W. 2008, Monthly Notices of the Royal Astronomical Society, 385, 1494

Kunz, M. W., & Lesur, G. 2013, Monthly Notices of the Royal Astronomical Society, 434, 2295

Lesur, G., Kunz, M. W., & Fromang, S. 2014, Astronomy and Astrophysics, 566, 56

Lubow, S. H., & Ida, S. 2011, Exoplanets, 347

Mignone, A. 2007, JCP, 225, 1427

Miyoshi, T., & Kusano, K. 2005, JCP, 208, 315

Nelson, R. P., & Papaloizou, J. C. B. 2004, MNRAS, 350, 849

Paardekooper, S. J., Baruteau, C., & Kley, W. 2011, Monthly Notices of the Royal Astronomical Society, 410, 293

Simon, J. B. 2016, The Astrophysical Journal Letters, 827, L37

Simon, J. B., & Armitage, P. J. 2014, The Astrophysical Journal, 784, 15

Simon, J. B., Bai, X.-N., Armitage, P. J., Stone, J. M., & Beckwith, K. 2013b, The Astrophysical Journal, 775, 73

Simon, J. B., Bai, X.-N., Flaherty, K. M., & Hughes, A. M. 2018, ApJ, 865, 10

Simon, J. B., Bai, X.-N., Stone, J. M., Armitage, P. J., & Beckwith, K. 2013a, The Astrophysical Journal, 764, 66

Simon, J. B., Hawley, J. F., & Beckwith, K. 2011, ApJ, 730, 94

Simon, J. B., Lesur, G., Kunz, M. W., & Armitage, P. J. 2015, Monthly Notices of the Royal Astronomical Society, 454, 1117

Stone, J. M., & Gardiner, T. A. 2010, ApJS, 189, 142

Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, The Astrophysical Journal Supplement, 178, 137

Youdin, A. N., & Lithwick, Y. 2007, Icarus, Icarus, 588