

Same Title???

Study of Colliding Flows and Feedback in Star Formation

Adam Frank, Jonathan J. Carroll-Nellenback, Baowei Liu, Martín Huarte-Espinosa, Shule Li, Erica L. Kaminski, Edward Hansen

ABSTRACT

The following two paragraphs need to be updated with TSF stuff

[—

Observational evidence from local star-forming regions mandates that star formation occurs shortly after, or even during, molecular cloud formation. Models of the formation of molecular clouds in large-scale converging flows have identified the physical mechanisms driving the necessary rapid fragmentation. They also point to global gravitational collapse driving supersonic turbulence in molecular clouds. Previous cloud formation models have focused on turbulence generation, gravitational collapse, magnetic fields, and feedback. In this work we propose exploring the effect of structure in the flow on the resulting clouds and the ensuing gravitational collapse. We seek to compare two situations with different flow properties and processes (flow shear, geometry, magnetic fields, feedback). We have already found that structured inflows lead to a delay of local gravitational collapse (“star formation”). Thus, more gas has time to accumulate, eventually leading to a strong global collapse, and thus to a high star formation rate. Uniform inflows fragment hydrodynamically early on, leading to the rapid onset of local gravitational collapse and an overall low sink formation rate.

Based on these results obtained with our previous XSEDE allocations (TG-AST120029) we propose to carry out a series of three-dimensional Adaptive Mesh Refinement (AMR) parallel numerical simulations to study the cloud collapse and cluster formation under various conditions. This will include colliding flows with different shear angles, different magnetic field strengths and orientation angles with respect to the flow.

—]

We are in full production mode with AstroBEAR, but the powerful numerical resources of the XSEDE that we are requesting are needed to make further progress on unraveling the mysteries of feedback process in star formation. To achieve this exciting goal we request the support of **4.0 million SU’s** on Stampede at TACC, **2.1 million SU’s** on Gordon at SDSC and **2.0 million SU’s** on SuperMIC at LSU.

1. Introduction

This section needs to be updated with TSF!!

[—

The concept of flow-driven cloud formation (Vázquez-Semadeni et al. 1995; Ballesteros-Paredes et al. 1999; Hartmann et al. 2001) can explain two observational constraints on how molecular clouds form stars, derived from local star-forming regions: first, all local molecular clouds form stars,

and second, the stellar age spreads are on the order of 1 – 2 Myr, several times shorter than cloud crossing times (Hartmann et al. 2001; Ballesteros-Paredes & Hartmann 2007, see summary in). The first constraint suggests that star formation sets in immediately (or even during) molecular cloud formation, and that the second constraint is trivially fulfilled in a scenario where the clouds themselves form in large-scale “converging” flows. The immediate (“rapid”) onset of star formation

in the forming clouds and the fact that the star formation efficiency is only a few percent (Evans et al. 2009) mandates that the clouds are highly structured: local collapse must set in before global collapse can overwhelm the dynamics.

The notion of cloud formation in converging flows has led to a series of numerical experiments investigating the physical processes relevant for the rapid fragmentation and for the control of the star formation efficiency. There is agreement across the models (despite different initial and boundary conditions) on the following results: (1) Rapid fragmentation is induced by strong radiative losses during the flow collision (possibly even by thermal instability if the clouds form from atomic gas), and by dynamical instabilities (Hueckstaedt 2003; Audit & Hennebelle 2005; Vázquez-Semadeni et al. 2006; Heitsch et al. 2008b) (2) Turbulence in molecular clouds is a natural result of the dynamical instabilities during the cloud formation, and is driven by global gravitational collapse at later stages of the cloud evolution (Vázquez-Semadeni et al. 2007). Some evidence for this is seen in the filamentary nature of some clouds (Fig 1). (3) Strong, non-linear density contrasts can also be driven by self-gravity in finite clouds, due to geometry (or “edge”) effects (Burkert & Hartmann 2004). (4) Although the rapid fragmentation can keep the star formation efficiency low, eventually, feedback or cloud dispersal is needed to prevent a large percentage of the gas to participate in gravitational collapse (Vázquez-Semadeni et al. 2010).

The most obvious difference in the simulation results is the morphology of the forming clouds. All models use some mechanism to break the symmetry in the colliding flows – otherwise, the flow collision would result in a plane-parallel shock. Models with small-scale perturbations (“noise”) in the velocities tend to lead to extremely flattened clouds with a strong ring due to the gravitational edge effect (Burkert & Hartmann 2004; Hartmann & Burkert 2007) in finite, sheet-like clouds. If the velocity perturbations are imposed on larger scales (e.g as a turbulent power spectrum), or if the collision interface between the two flows is perturbed, turbulent structures form that are not necessarily coherent when viewed from different directions (Heitsch et al. 2009).

2. Previous Accomplishments

new section required for the renewal !!

2.1. Summary of Scientific Discoveries

new section required for the renewal !!

2.2. Details of Previous Accomplishments

new section for the renewal !!

3. Accomplishments and the Way Forward

From the old proposal, need update with TSF or tailor to the section above




Fig. 1.— Comparison of “Clumpy” run (top) and “Smooth” run (bottom) showing the projected column density along the flow direction (right) and perpendicular to the flow direction (left)

To understand better the effect of initial conditions on the clouds forming in the flow collisions, we have already completed a study (Carroll et al 2013) comparing two (otherwise identical) cloud formation simulations, one with a smooth inflow, and one with a clumpy inflow of identical mass flux. Our studies were motivated by the analysis of Pringle et al. (2001), suggesting that cloud formation out of warm atomic gas would require

time and length scales too large to be feasible (see also McKee & Ostriker (2007)). This problem is solved by the realization that the flow collision is three-dimensional, allowing gravitational collapse and accretion along the dimensions perpendicular to the flows, and thus circumventing the one-dimensional limit on column density accumulation (Heitsch et al. 2008a). Yet, Pringle et al. (2001) suggest that clumpy flows could reduce the accumulation and molecule formation time scale (especially if the clumps are already molecular).

The models were identical except for the physical conditions of the inflowing gas. One model (“Smooth”) began with a completely uniform flow. The other model (“Clumpy”) used the same mass inflow rate, but with the gas distributed in dense clumplets with a filling factor of 5%. In both cases, the thermal and dynamical fragmentation of the shocked collision region leads to turbulence, at a level of 10 to 20% of the inflowing kinetic energy. This is consistent with earlier results. The Clumpy run showed a somewhat higher velocity dispersion initially, since the clumps are less easily stopped in a flow collision. (Fig 2)

Due to the lower compression factor in the Clumpy model, less gas is being cooled to higher densities than in the Smooth run. Thus, the kinetic energy of the inflow is less efficiently dissipated. Together with a non-contiguous distribution of cold, sub-jeans, fragments, this leads to a suppression of local collapse for nearly 20 Myr after the initial flow collision. At that point, sufficient mass has assembled to induce global collapse of the whole region, resulting in a “star burst” (more appropriately, “sink burst”) at a high sink formation rate. In contrast, the Smooth run shows local collapse already after 10 Myr, at less than half the sink formation rate of the Clumpy run. Due to the local nature of the thermal fragmentation, more fragmentation will occur with increasing resolution (Hennebelle & Audit 2007), thus, the times quoted here are *upper* limits for the onset of local collapse. Nevertheless, structured flows can delay the onset of (substantial) local collapse. Global collapse is only clearly visible in the Clumpy run.

The differences between Clumpy and Smooth inflows extend to the mixing efficiencies. Somewhat counter to a naive expectation, the Smooth initial conditions result in less well mixed material

(and cores). This is primarily due to the Nonlinear Thin Shell Instability (NTSI) funneling material preferentially into the troughs located far into the opposing inflow. For the Clumpy run, the global collapse of the accumulated clumps erases all memory of the initial inflow direction.

In (Carroll et al 2013) we chose two extremes as our initial conditions. It is more likely that the inflows themselves will contain turbulent velocity and density structures that are coherent in space. Spatial coherence leads to stronger shocks in the collision region, and thus to more efficient energy dissipation. In that sense, our Clumpy run is overestimating the effect of structured inflows. In addition magnetic fields are likely to play a strong role in shaping the evolution of the flows. The challenge ahead is to use a wider array of initial conditions and include more physical processes (such as magnetic fields) to explore the evolution of colliding flows and compare predictions from these cases with observations

4. Research Objectives

the follow is from the old proposal Need to be updated !!! We will focus on simulations and observational implications of colliding flows under a variety of conditions. This will include (1) colliding flows interacting with various degrees of shear and (2) colliding flows with different magnetic field strengths (3) colliding flows with different magnetic field orientations. Our computations will be carried out using AstroBEAR.

5. Computational Approach

Updated by Baowei 4/10/2014!! AstroBEAR is an Adaptive Mesh Refinement (AMR), multi-physics code for astrophysics. AMR remains at the cutting edge of computational astrophysics. AMR simulations adaptively change resolution within a computational domain to ensure that the most important features of the dynamics are simulated with highest accuracy. By allowing quiescent regions to evolve with low resolution, AMR simulations achieve order of magnitude increases in computational speed. After a decade of development only a handful of AMR-MHD codes exist for astrophysics: (e.g. FLASH, ENZO RAMSES, ORION, CASTRO).




Fig. 2.— Slice showing log density from 3D Colliding Flows simulation performed with AstroBEAR. Also shown is the AMR mesh which contains 4 additional levels of refinement

The UR astrophysics group successfully constructed and tested AstroBEAR, a fully parallelized, multi-dimensional AMR MHD code. The success of this effort is evidenced both in the code’s completion (Cunningham et al 2009) and the papers published using AstroBEAR as it was developed through its radiation-hydrodynamic and MHD versions (a partial list includes: Poludnenko et al 2004ab; Cunningham et al 2005; 2006ab, Hartigan et al 2007, Dennis et al 2008, Yirak 2009, 2010, Li et al 2012, Huarte-Espinosa et al 2012).

The multi-physics capabilities of AstroBEAR have been significantly expanded by including solvers for elliptic and parabolic equations. Adapting the linear system solver HYPRE, we now routinely simulate systems in which self-gravity, heat conduction and magnetic resistivity are important. Radiation transfer in the diffusive limit is currently being added. In addition, AstroBEAR can treat gravitationally interacting point particles which accrete mass.

5.1. AstroBEAR Scaling

AstroBEAR is designed for 2D and 3D adaptive mesh refinement (AMR) simulations which require algorithms that are highly parallelized and manage memory efficiently. AstroBEAR uses a hierar-

chical approach to parallelization suitable for multicore architectures in which large-scale patches of data are distributed to nodes using MPI and the work for an individual patch is distributed across the cores on a node using OpenMP directives. AstroBEAR also employs new techniques such as load balancing by threading the grid advances on each level with preference going to the finer level grids.

Here we present strong scaling and weak scaling results for AstroBEAR. In Figure 5.1, we report scaling test results on Stampede at TACC. The resolution we used for these test are $40^3 + 5$ level AMR which is same as the computation we are planning to do. The strong scaling test plot of the current code shows a slope -0.83 on up to 1024 cores and -0.69 on up to 4096 cores while the slope for perfect scaling is -1 . This shows AstroBEAR has an excellent scaling on Stampede. For the weak testing results, we keep the work load on each core same (each has 64^3 cells). So the weak scaling result measures the efficiency of data communication between the processors. Our test results show the communication efficiency drops only 15% running on 1024 cores and 27% on 4096 cores. More results and details about the performance of AstroBEAR can be found in the report of "Code Performance of AstroBEAR2.0".

6. Resource Request

Tables and Texts are added by Baowei on 4/10. Text and Numbers need to be revisited!!!

For our colliding flows project, we propose to run 5 levels of AMR (instead of 4 levels as in our last proposal) and smaller base dimensions (comparing with our last proposal). This will increase the filling fractions and increase the efficiency of using computing resources. Our estimate CPU-hours for this project is based on the strong scaling tests we’ve done which has the same resolution but running for only 0.1% of one frame (we are planning to obtain 200 frames of the data). For our scaling test on 1024 cores, it takes about 144seconds or 41CPU – hour to run 0.1% of one frame. So the total CPU-hours for one production run is about 820,000 CPU-hours on Stampede. Since there’s no I/Os when doing scaling test and it takes about 20 150 seconds to save one frame of data (in format of chombo files, depending on the




Fig. 3.— Top: Strong scaling behavior of AstroBEAR with Triggered Star Formation module and 3 levels of AMR on Stampede at TACC. Running time is plotted versus the number of cores in log scale. It has a slope -0.84 up to 1024 cores which shows excellent scaling (slope -1 for "perfect scaling"). The code performance drops when running on more cores (slope -0.69 up to 4096 cores) of Stampede. Bottom: Weak scaling behavior of AstroBEAR with Colliding Flows module and fixed grid on Stampede. Runtime efficiency ($1/t$ where t is the running time) is plotted versus the number of cores in log scale. The work load on each core is kept same (each has 64^3 cells) for the weak scaling test. The communication efficiency of AstroBEAR drops only 15% running on 1024 cores and 27% on 4096 cores.

size of data). So we expect ≈ 1000 CPU-hours to save the total 200 frames on 1024 cores for each

run. So the total CPU-hours for each run is about 821,000 on Stampede (See Table(1)).

When doing strong scaling tests for our triggered star formation project, we start from a point where the stars are just formed to get a better estimate. And we use the same base dimensions but lower (3 or 2) AMR levels (due to the availability of starting point at high levels). It took about 361 seconds to get 20% of one frame from the starting points on 1024 cores which is about 103 CPU hours. And we estimate the running time will be 8 times longer when running with 4 levels of AMR. The current data size for one frame of 3-level-AMR is about 3GB and we estimate for our production runs the data size of each frame is about 6-12 GB. Depending on this information we estimate the total CPU-hours we need for each run of our triggered star formation project is 413,000 on Stampede (See Table(1)).

In total we require 4.0 million CPU-hours on Stampede and 2.1 million CPU-hours on Gordon and 2.0 million CPU-hours on SuperMIC, about 98% of which will be used for production runs and 2% for testing runs and continue development of our code. We can use anywhere from 500-5000 cores for a typical production run, depending on queue limitations.

The following 2 paragraphs are from the old proposal !!! — Using AstroBEAR, we have found very interesting new results related to star formation - in particular studies of feedback (Quillen et al 2005, Cunningham et al 2009, Carroll et al 2009, 2010) as well as molecular cloud formation (Carroll et al 2013). We plan to carry out a series of three-dimensional numerical simulations of colliding flows with different shear angles, magnetic field strengths, and magnetic field orientations to study the star-forming properties of a cluster. In particular this will include colliding flows with $\beta = \infty$ and shear angles 0° , 30° and 60° as well as colliding flows with $\beta = 1$ and $\beta = 10$ and shear angles of 0° , and 60° , as well as colliding flows with $\beta = 1$, a shear angle of 0° , and magnetic field orientation angles of 0° and 60° . See Table ?? This will allow us to address the interplay between magnetic field strength, orientation, and shear of the colliding flows with respect to the resulting properties of the molecular clouds as well as the cores that form within them. AstroBEAR is presently performing well

with multiple 6+ levels of AMR refinement. The tractability of a given run then becomes more a question of the number of needed cell updates, which is mainly determined by the fraction of the volume where mesh refinement is employed. And the volume filling fraction depends on the specific problem/simulation. The computing resources we require are based on our previous runs on XSEDE machines. For example, Our hydro calculation with 0° shear and with resolution $128^3 + 4$ AMR took about 6,000 SU's for 1 frame. In Tables ?? we summary the computing resources we require. In total we require 7.1 million CPU-hours, 99.3% of which will be used for production runs and 0.7% for testing runs and continue development of our code. We can use anywhere from 500-5000 cores for a typical production run, depending on queue limitations. —

The following section is added by Baowei on 04/10/14

7. Introduction to Local Resources and Research Team

We have a variety of computational platforms available locally for our research. The Center of Integrated Research Computing(CIRC) at the University of Rochester hosts the Blue Gene/Q system called BlueStreak. BlueStreak consists of 1,024 nodes, 16,384 CPU cores, 16 TB of RAM, and 400 TB of storage. Each node consists of a 16-core A2 processor with 32 MB of cache and access to 16 GB of RAM. By reservation our group can have access to the whole rack(Normal reservations are for less than half of the rack, 24 hours per reservation). CIRC also hosts another Linux cluster called BlueHive which is expected to be in production in June of 2014. The new BlueHive system consists of approximately 200 nodes of IBMs iDataPlex architecture with a high-speed, low-latency, InfiniBand interconnect. Each node houses 2 x 12-core Intel Ivy Bridge processors, with a range in memory from 64 GB to 256 GB. 30 of the 64 GB nodes have 2 dedicated NVIDIA K20X (Kepler) GPU (graphics processing unit) cards and 8 of the 64 GB nodes have 2 dedicated Intel Phi 5110P accelerator cards. Our group has access to 170 cores of the BlueHive system and 96 cores of them are owned by and dedicated to our group. Our group also has a dedicated 32 processor equiv-

alent (PE) cluster built for developing and debugging AstroBEAR code.

This part needs more details Our research group which currently has one faculty member, two computational science staffs, two postDocs, four graduate students and two undergraduate students has a very strong backgroup of computaional astrophysics.....

7.1. I/O Requirements, Analysis, and Storage

Based on Table (1), we expect to save 200 frames of data with size 5-20GB for each frame (the size of the each frame depends how many zones trigger the highest level of AMR) for each of the six runs of our colliding flows simulation with 5 levels of AMR. So the total data size for our colliding flows project is about 6 TB - 24 TB. We are planning to run these simulations on Stampede and we expect to need ~ 24 TB of storage on Ranch of TACC. For each of of the four runs of our triggered-star-formation simulations, we expect to save 100 frams of data with size about 6GB, So the total data size for our triggered-star-formation project is about 2.4TB. These simulations will be done on Gordon. So we requestion ~ 3 TB on Gordon. Most of the analysis will be done using parallelized tools built into the AstroBEAR package that can be performed at run time or in post processing.

7.2. Why choosing to run on these machines?

this paragraph is added by Baowei on 04/10/14. **Need revisit !!!** We've been using Kraken and Stampede (before Ranger) for the past two very productive years. With Kraken decommissioned this April, we have to shift most our production runs to Stampede while seeking for opportunities on other machines. AstroBEAR behaves very good on Stampede (see the scaling tests attached). Since SuperMIC and Gordon have similar archetectures to Stampede, we expect the code will scale well on these two machines. For our Colliding flows project, we propose to run with higher level AMRs than before. So it will requires more memory usage. Gordon will be a good machine for these runs since it has big (64GB) RAM per node. Also we are planning to optimize the code

on the Xeon Phi coprocessors of our local Blue-Hive system. So we expect the code will take the advantage of the MIC architecture and run even better on Stampede and SuperMIC.

7.3. Financial Support

waiting for updated information from Sondra. Will be updated by Baowei!! Financial support for this project will come from the Space Telescope Sci Institute grant HST -AR-12128.01-A entitled “STSci - Hubble Telescope - The Reel Deal: Interpreting HST Multi-Epoch Movies of YSO JetsSpace” (PI, A. Frank; 10/1/2010 - 9/30/2013) and the Department of Energy grant DE-SC0001063 entitled “The dynamics of magnetized Astrophysical Jets through Pulsed Power HEDP lab Studies” (PI, A. Frank; 8/15/2012 - 8/14/2015) and the National Science Foundation, NSF AST-1109285 entitled “From Central Engine to Bipolar Outflow: Binaries, MHD and the Evolution of Planetary Nebulae” (PI, A. Frank; 9/1/2011 - 8/31/2014). DOE, Award no. R17081, entitled “Rice - Clumpy Environments & Interacting Shock Waves: Realistic Laboratory Analogs of Astrophysical Flows”, (PI: A. Frank, 2/22/2011 - 2/21/2014) and the Space Telescope Sci Institute grant HST-AR-12832.01-A entitled “Hubble Telescope Cycle 20 - Climbing the Ladder of Star Formation Feedback“, (PI A. Frank, 11/1/2012 - 10/31/2015).

8. Summary

Text from old proposal, need to updated !!! With the advent of our efficient 3D AMR MHD code AstroBEAR and state of the art HPC facilities, we are on the doorstep of significant breakthroughs in understanding the physics of the feedback process in the clustered star formation. With our previous start-up XSEDE allocation (TG-AST120029), we have already found and are in the process of publishing very important results. We are in full production mode with AstroBEAR, but the powerful numerical resources of the XSEDE that we are requesting are needed to make further progress on unraveling the mysteries of interacting binaries.

If granted this allocation, we will perform a series of 3D Adaptive Mesh Refinement parallel numerical simulations with higher resolution and

longer time-scales than ever before to study the formation, structure and stability and observational implications of inter-acting stellar binaries.

- Study of colliding flows with MHD 4,930,000 SU’s on Stampede.
- Study of triggered star formation with/without MHD 10,000,000 SUs on Stampede.
- Testing and continue development of our code. 7,000 SU’s on Stampede; 0.7% of total request.

To achieve these goals, we request SU’s on the TACC Stampede. bibliography need to be updated !!!

REFERENCES

- Audit, E., & Hennebelle, P. 2005, A&A, 433, 1
- Bally, J., 2011, Computational Star Formation, 270, 247
- John Bally, Nathaniel J. Cunningham, Nickolas Moeckel, Michael G. Burton, Nathan Smith, Adam Frank, and Ake Nordlund, 2011, ApJ, 727, 113
- Bo Reipurth and John Bally, 2001, ARAA, 39, 403
- Bally, J., Licht, D., Smith, N., & Walawender, J., 2006, AJ, 131, 473
- Bally, J., Reipurth, B., and Davis, C., 2006, Protostars and Planets V, University of Arizona Press, Tucson, in press
- Ballesteros-Paredes, J., Hartmann, L., & Vázquez-Semadeni, E. 1999, ApJ, 527, 285
- Ballesteros-Paredes, J., & Hartmann, L. 2007, Revista Mexicana de Astronomia y Astrofisica, 43, 123
- Burkert, A., & Hartmann, L. 2004, ApJ, 616, 288
- Carroll, J.J., Frank, A., & Blackman, E.G. 2010, ApJ, 722, 145
- Carroll, J.J., Frank, A., & Heitsch, F. 2013, “The effects of Inhomogeneities within Colliding Flows on the Formation and Evolution of Molecular Clouds”, in preparation

- Carroll, J.J., Frank, A., Blackman, E.G., Cunningham, A.J., & Quillen, A.C. 2009, *ApJ*, 695, 1376
- Carroll, J. Shroyer, B; Frank, A; Ding, C eprint arXiv:1112.1710. 2011, submitted to JCP
- Cunningham, A.J., Frank, A., Carroll, J., Blackman, E.G., & Quillen, A.C. 2009, *ApJ*, 692, 816
- Cunningham, A. Frank, A., Varniere, P., Mitran, S., & Jones, T.W. 2009, *ApJS*, 182, 51
- Elmegreen B., G., 2000, *ApJ*, 539, 342
- Evans, II, N. J., Dunham, M. M., Jørgensen, J. K. et al. 2009, *ApJS*, 181, 321
- Hartmann, L., Ballesteros-Paredes, J., & Bergin, E. A. 2001, *ApJ*, 562, 852
- Hartmann, L., & Burkert, A. 2007, *ApJ*, 654, 988
- Heitsch, F., Naab, T., & Walch, S. 2011, *mnras*, 415, 271
- Heitsch, F., Hartmann, L. W., Slyz, A. D., Devriendt, J. E. G., & Burkert, A. 2008, *ApJ*, 674, 316
- Heitsch, F., Hartmann, L. W., & Burkert, A. 2008, *ApJ*, 683, 786
- Heitsch, F., Ballesteros-Paredes, J., & Hartmann, L. W. 2009, *ApJ*, 704, 1735
- Hennebelle, P., & Audit, E. 2007, *A&A*, 465, 431
- Huarte-Espinosa, M., & Frank, A., 2012, arXiv:1201.4322
- Hueckstaedt, R. M. 2003, *New Astronomy*, 8, 295
- Klein, R., Inutsuka, S., Padoan, P., and Tomisaka, T , 2006, *Protostars and Planets V*, University of Arizona Press, Tucson
- Krumholtz, M., McKee, K., & Klein, R., 2005, *Nature*, 438, 333
- Knee, L.B.G., & Sandell, G., 2000, *aap*, 361, 671
- Nakamura, F., & Li, Z.-Y., 2011, *ApJ*, 740, 36
- Nakamura, F., & Li, Z.-Y., 2007, *ApJ*, 662, 395
- Matzner, C.D., 2002, *ApJ*, 566, 302
- McKee, C. F., & Ostriker, E. C. 2007, *ARA&A*, 45, 565
- Moriarty-Schieven, G.H., Johnstone, D., Bally, J., & Jenness, T. 2006, *ApJ*, 645, 357
- Norman, C., & Silk, J., 1980, *ApJ*, 238, 158
- Pringle, J. E., Allen, R. J., & Lubow, S. H. 2001, *MNRAS*, 327, 663
- Quillen, A.C., Thorndike, S.L., Cunningham, A., Frank, A., Gutermuth, R.A., Blackman, E.G., Pipher, J.L., & Ridge, N., 2005, *ApJ*, 632, 941
- Smith, N., Povich, M.S., Whitney, B.A., et al. 2010, *mnras*, 406, 952
- Smith, N., Bally, J., & Walborn, N.R. 2010, *mnras*, 405, 1153
- Vázquez-Semadeni, E., Passot, T., & Pouquet, A. 1995, *ApJ*, 441, 702
- Vázquez-Semadeni, E., Ryu, D., Passot, T., González, R. F., & Gazol, A. 2006 *ApJ* 643, 245
- Vázquez-Semadeni, E., Gómez, G. C., Jappsen, A. K., Ballesteros-Paredes, J., González, R. F., & Klessen, R. S. 2007, *ApJ*, 657, 870
- Vázquez-Semadeni, E., Colín, P., Gómez, G. C., Ballesteros-Paredes, J., & Watson, A. W. 2010, *ApJ*, 715, 1302
- Warin, S., Castets, A., Langer, W. D., Wilson, R. W., & Pagani, L. 1996, *A&A*, 306, 935
- Wang, P., Li, Z.-Y., Abel, T., & Nakamura, F. 2010, *ApJ*, 709, 27

Table 1: Estimate CPU-hours needed for each run (based on 1024 cores of Stampede)

Project	Colliding Flows	Tiggered Star Formation
AMR levels for (Scaling VS Production)	5 VS 5	3/2 VS 4
Num of frames (Scaling VS Production)	0.1% VS 200	20% VS 100
WallTime/CPU-hour for Strong Scaling(1024 cores)	144 Secs/41	361 Secs/103
Data size of each frame	5-20 GB	3GB (6-12GB with 4 AMR)
Time needed to Save one frame	20-150 Seconds	22-120 Seconds
Estimate CPU-hours for computing	820,000	412,000
Estimate CPU-hours for I/Os	1,000	800
Total Estimate CPU-hours	821,000	413,000

Table 2: Expected CPU-hours for Continuing Colliding Flows 3D Simulations

Shearing Angle(θ)	Resolution	Magnetic Field (β)	Expected SUs on Stampede
0	$40^3 + 5$ AMR	1	821,000
15	$40^3 + 5$ AMR	1	821,000
30	$40^3 + 5$ AMR	1	821,000
0	$40^3 + 5$ AMR	1	821,000
15	$40^3 + 5$ AMR	1	821,000
20	$40^3 + 5$ AMR	1	821,000
30	$40^3 + 5$ AMR	1	821,000
Total		4,926,000	

Table 3: Expected CPU-hours for 3D Triggered Star Formation Simulations

Hydro/MHD	Resolution	Rotation	Expected SUs on Stampede
Hydro	$320 \times 912 \times 192 + 4$ AMR	Y	413,000
Hydro	$320 \times 912 \times 192 + 4$ AMR	Y	413,000
MHD $\beta = 1$, Toroidal	$320 \times 912 \times 192 + 4$ AMR	Y	413,000
MHD $\beta = 1$, Poloidal	$320 \times 912 \times 192 + 4$ AMR	Y	413,000
Total		1,652,000	