Code Performance of AstroBEAR2.0

Adam Frank, Jonathan J. Carroll-Nellenback, Baowei Liu,
Eric Blackman, Jason Nordhaus, Shule Li

June 27, 2015

1 Abstract

To be updated by Baowei .... Strong and weak scaling behavior of the current AstroBEAR2.0 code were tested on Stampede of TACC. To have an accurate estimate the resource and optimal core numbers for our production runs, we use the same multiphysics features (with self-gravity and HYPRE library) as our production runs when running the scaling tests. For the strong scaling test with the Colliding Flows module and the 3D Triggered Star Formation module,we got good scaling result on Stampede (with slope -0.83 and -0.84 on 1024 cores respectively, and slope -0.68 and -0.69 respectively on 4096 cores while the perfect scaling slope is -1). For the weak scaling with fixed grid, we found the running efficiency drops about 17% for up to 1024 cores on Stampede and 27% on 4096 cores due to the communication increase between the processors. The size of our typical run jobs will be 1024 cores.

2 Colliding Flows Module

Figure (1) shows the strong scaling test result of AstroBEAR with the 3D Colliding Flows module on Stampede. The strong scaling tests are done with the same resolution as our proposed production runs (40^3 + 5 levels AMR but for much shorter final time (0.1% of one frame while the production runs need 200 frames) and without IOs. The current code with the Colliding Flows modules shows good scaling on Stampede (with slope$=-0.83$ while the perfect scaling has a slope -1) up to 1024 cores. The behavior drops when running on more cores (with slope$=-0.68$). We plan to run our 3D Colliding Flows calculations on 1024 cores but can also run on 4096 cores to save the walltime. The SUs we request are based on the estimates on 1024 cores (See Table 1 in the main document).

In Figure (1) we show the weak scaling test result of AstroBEAR with 3D Colliding Flows module on Stampede. In these tests we use the same resolution as the strong scaling test case but with fixed grid. The number of cells on each core are keep constant (64^3 per core). Weak scaling is good to test data communication part of the code since the work load on each processor are about

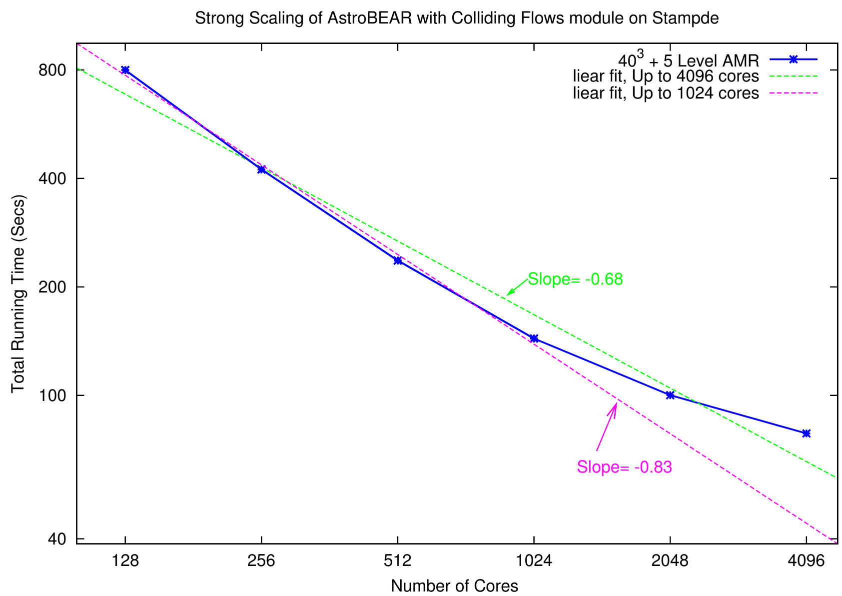

Figure 1: Current strong scaling behavior for AstroBEAR with 3D Colliding Flows module on Stampede. For these tests we run with with the same resolution as our proposed production runs (40^3 + 5 levels AMR for Colliding Flows), but for shorter final time (0.1% of one frame while the production runs need 200 frames) and without IO. The current code with the Colliding Flows modules shows good scaling on Stampede (with slope equals -0.83 up to 1024 cores and -0.68 up to 4096 cores. Perfect scaling has a slope -1). We are planning to run our productions runs on 1024 or 4096 cores. The SUs we request are based on the estimate time on 1024 cores.

same. We run these tests with self-gravity and HYPRE library. The runtime efficiency of AstroBEAR drops only 17% on up to 1096 processors and 27% up to 4096 processors which shows a very good weak scaling considering the multiphysics features in these runs.

3 3D Triggered Star Formation

Figure (3) shows the strong scaling behavior for AstroBEAR with 3D Triggered Star Formation module on Stampede. The base dimensions for these tests are same as our production runs but with fewer levels of AMRs (3 levels AMR while the planed production runs has 4 levels). The run is restarted from a point where stars start to form (Figure 4) and this yields a good estimate of run time. We estimate the run time with 4 levels of AMR will be 8 times longer than the runs with 3 levels of AMR, although this depends on the filling fraction which varies at different times. For these tests we run the code for 20% of one

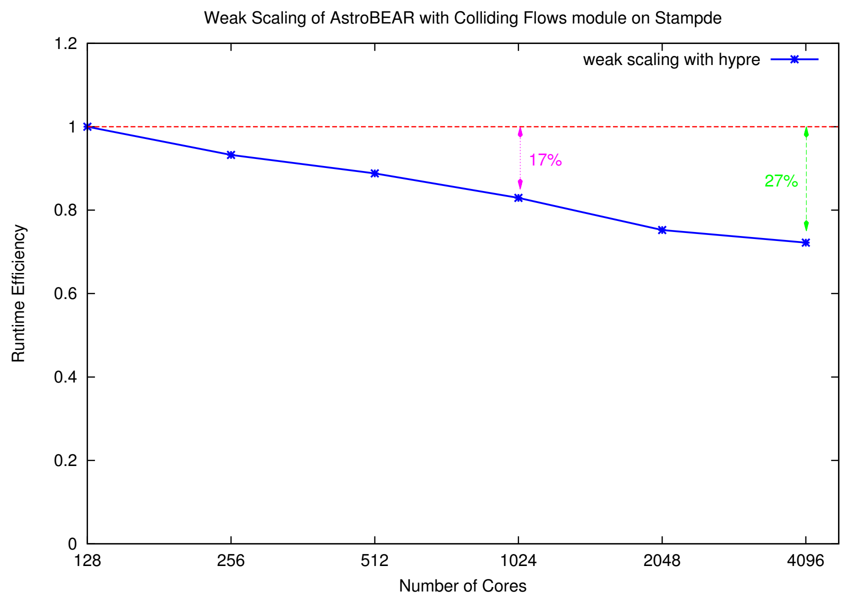

Figure 2: Current weak scaling behavior for AstroBEAR with 3D Colliding Flows module. The test is done with fixed grid and same work load on each processor(each processor has 64^3 cells). We run these tests with self-gravity and using the HYPRE library. The runtime efficiency of AstroBEAR drops only 17% on up to 1096 processors and 27% up to 4096 processors which shows a very good weak scaling considering the multiphysics features involved in these tests. Our typical production runs will be on 1024 cores.

frame (the production runs need 100 frames) and without IO. The current code with the Triggered Star Formation module shows good scaling on Stampede with slope$=-0.84$ on 1024 cores and -0.69 on 4096 cores.

4 Summary

Strong and weak scaling behavior of the current AstroBEAR2.0 code were tested on Stampede of TACC. For the strong scaling test with the Colliding Flows module and the 3D Triggered Star Formation module, we got good scaling result on stampede. We expect the same behavior on other systems (for example, Gordon at SDSC) with similar architecture to Stampede.

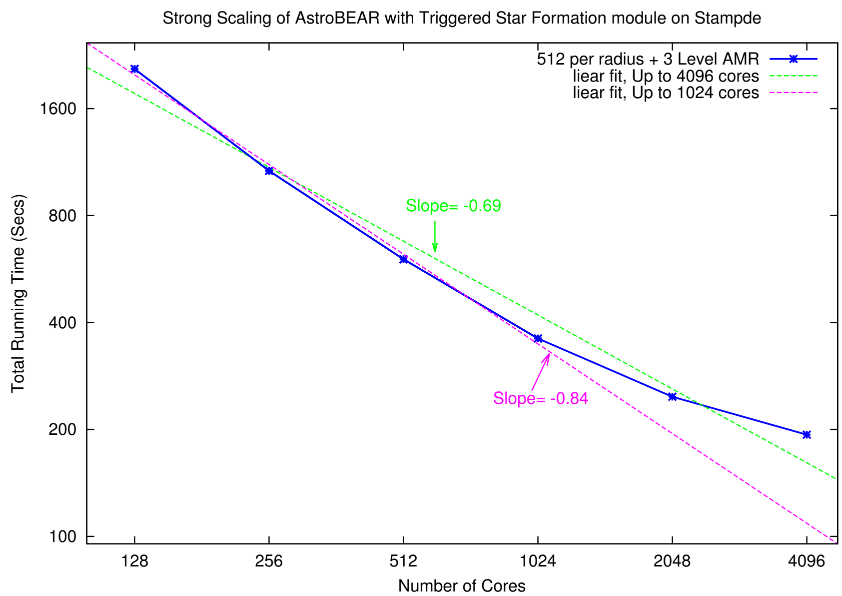

Figure 3: Current strong scaling behavior for AstroBEAR with 3D Triggered Star Formation module on Stampede. The base dimensions for these tests are same as our production runs but with fewer levels of AMRs (3 levels AMR while the planed production runs has 4 levels). We restart from a point where the star formation starts to get a better estimate for much shorter final time (1% of one frame while the production runs need 100 frames) and without IOs. The current code with the Triggered Star Formation module shows excellent scaling on Stampede (with slope$=-0.84$ while the perfect scaling has a slope -1) up to 1024 cores. The behavior drops when running on more cores (with slope$=-0.69$ on up to 4096 cores). Our typical production runs are on 1024 cores and the SUs we request are based on the estimate time on 1024 cores.

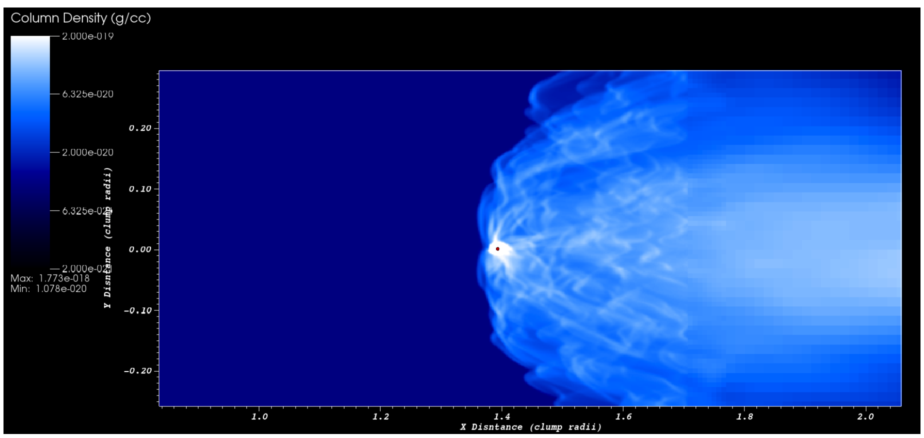

Figure 4: The start point of strong scaling test for AstroBEAR with 3D Triggered Star Formation module. It shows a star formation is triggered at the center of the image. This start point has the same base dimension as our proposed runs and 3 levels of AMR (we propose to run at 4 levels of AMR).