

# Code Performance of AstroBEAR2.0 on Stampede

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

April 14, 2014

## 1 Abstract

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 very excellent scaling result on up to 1024 cores on Stampede (with slope  $-0.83$  and  $-0.84$  respectively while the perfect scaling slope is  $-1.0$ ). The behavior of the code with these two modules drops to (with slope  $-0.68$  and  $-0.69$  respectively) on 4096 cores. For the weak scaling with fixed grid, we found the running efficiency drops about 17% 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 job 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 exact 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 excellent 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  $= -0.68$ ). Our typical production runs are on 1024 cores and the SUs we request are based on the estimate time on 1024 cores.

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 fixed grid. The number of cells on each core are keep as constant ( $64^3$  per core). Weak scaling is good to test the data communication part of the code since the work load on each processor are about same. The runtime efficiency of AstroBEAR drops only 17% on up to

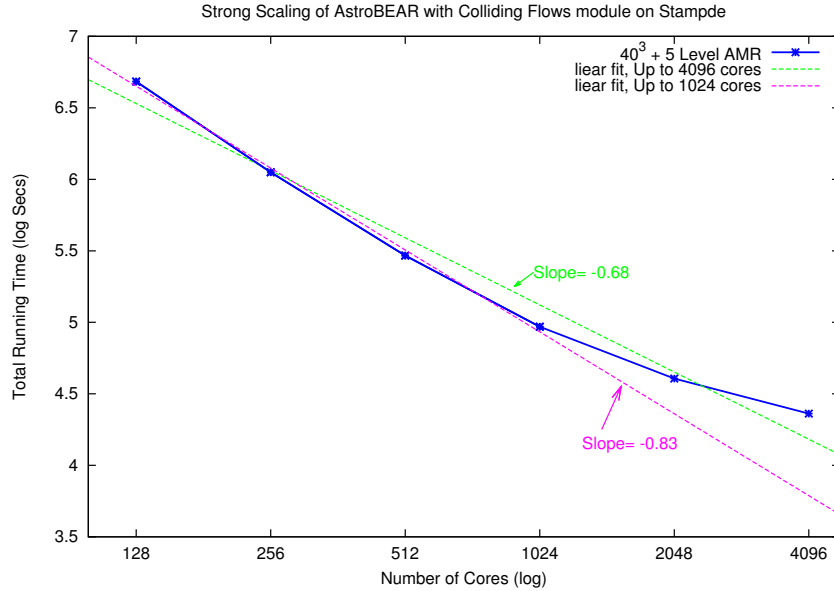


Figure 1: Current strong scaling behavior for AstroBEAR with 3D Colliding Flows module on Stampede. For these tests we run with the exact same resolution as our proposed production runs ( $40^3 + 5$  levels AMR for Colliding Flows), but for much shorter final time (0.1% of one frame while the production runs need 200 frames) and without IOs. The current code with the Colliding Flows modules shows excellent scaling on Stampede (with slope= $-0.83$  while the perfect scaling has a slope  $-1$ ) up to 1024 processors. The behavior drops with this particular module and production set up when running on more processors (with slope= $-0.68$ ). Our typical production runs are on 1024 processors and the SUs we request are based on the estimate time on 1024 processors.

1096 processors and 27% up to 4096 processors which shows a very good weak scaling. Our typical production runs will be on 1024 cores.

### 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 less levels of AMRs (3 levels AMR while the planned production runs has 4 levels). We think if we restart from a point where stars start to form (Figure 4) we will have a better estimate about the running time. But We don't have a similar picture with 4 levels of AMR so we use 3 levels of AMR instead to do our scaling test. We estimate the running time with 4 levels of AMR will be 8 times longer than the runs with 3 levels of AMR, although this is highly depended on the fill fraction ratio which varies

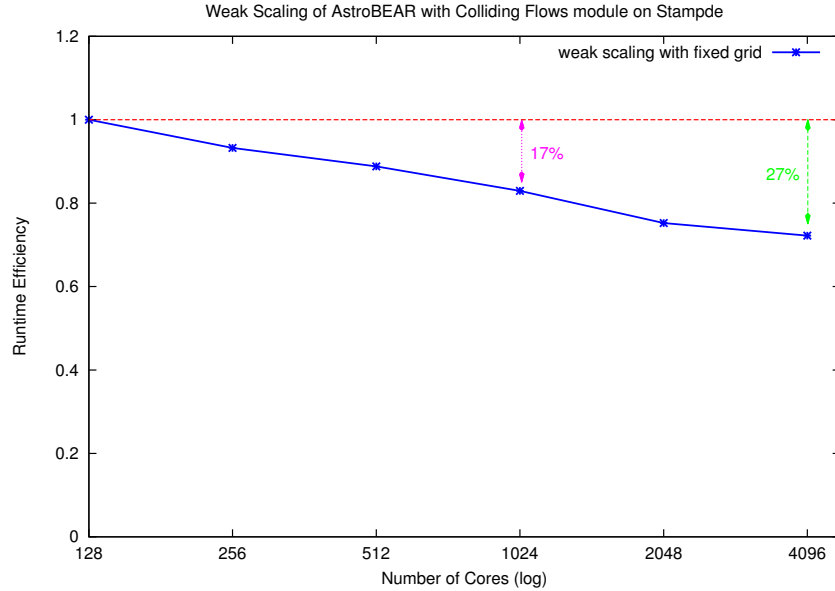


Figure 2: Current weak scaling behavior for AstroBEAR with 3D Colliding Flows module. The test is done with fixed grid and same work load on each processor (each processor has  $64^3$  cells). Weak scaling is good to test the data communication part of the code since the work load on each processor are about same. 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. Our typical production runs will be on 1024 cores.

at different time. For these tests we run the code for 20% of one frame (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 =  $-0.84$  (the perfect scaling has a slope  $-1$ ) up to 1024 processors. The behavior drops when running on more processors (with slope =  $-0.69$  on up to 4096 processors which is consistent with the runs of the Colliding Flows module).

## 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 very good scaling result on up to 1024 cores 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 less levels of AMRs (3 levels AMR while the planned production runs has 4 levels). We restart from a point where the star formations 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 processors. The behavior drops when running on more processors (with slope= $-0.69$  on up to 4096 processors). Our typical production runs are on 1024 processors and the SUs we request are based on the estimate time on 1024 processors.




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).