Detailed project description: Astrophysical turbulence and dynamo action

Axel Brandenburg (Nordita, Stockholm)

October 21, 2014

Abstract

In addition to our activity on the "Formation of active regions in the Sun" (VR breakthrough research grant) we now perform studies of "Particle transport and clustering in turbulent flows" (Research Council of Norway, FRINATEK research grant).

1 Background

Our research group now consists of The following people will be involved:

Mr Xiang-Yu Li (PhD student)
Ms Sarah Jabbari (PhD student, Licentiate 19 May 2014)
Ms Illa R. Losada (PhD student, Licentiate 5 December 2014)
Dr Nishant Singh (Post-doc)
Dr Bidya B. Karak (Nordita fellow)
Dr Lars Mattsson (Nordita fellow)
Dr Dhrubaditya Mitra (assistant professor)

2 Scientific content

Our research covers a broad range of topics ranging from coagulation of particules in turbulent flows (work with Li, Mattsson, and Mitra; supported by the Norwegian Research Council and the Wallenberg foundation), to magnetic flux concentrations in the Sun (Jabbari, Losada; supported through two VR grants) to helioseismology (Singh; also supported through VR). The work so far has been highly successful as is evidenced by the papers published in just the last 12 months, and promises to be even more vigorous owing to the expansion of the research to include coagulation studies.

3 Code and test case

For all tests, the Pencil Code was used. It is publicly available on `http://pencil-code.googlecode.com/`, where also detailed documentation is available. The code uses explicit sixth order finite differences. The time step is third-order. In this sample, we run isothermal magnetohydrodynamics in a periodic domain[1]. Power spectra are computed during the run, but our current parallelization of the Fourier transform requires that the meshpoint number is an integer

[1]Run directories are available on `http://norlx51.nordita.org/~brandenb/pencil-code/timings/bforced/`

multiple of the product of processor numbers in the y and z directions and the product of processor numbers in the x and y directions. In addition, the number of processors in one direction should not be so large that the number of mesh points per processor becomes comparable to or less than the number of ghost zones (which is 6).

4 Running the code

To run the code, get one of the sample run directories, e.g., `http://norlx51.nordita.org/~brandenb/pencil-code/timings/bforced/512_4x16x32`. The relevant file to be changed is `src/cparam.local`

```
ncpus=2048,nprocx=4,nprocy=16,nprocz=ncpus/(nprocx*nprocy)
nxgrid=512,nygrid=nxgrid,nzgrid=nxgrid
```

in particular the values of ncpus, nprocx, nprocy, and nxgrid. Once they are chosen, say make, and submit start_run.csh.

On Triolith, strong scaling tests have been performed for three mesh sizes. The time per time step and mesh point is given for different processor numbers and layouts. Generally, it is advantageous to keep the number of processors in the x direction small.

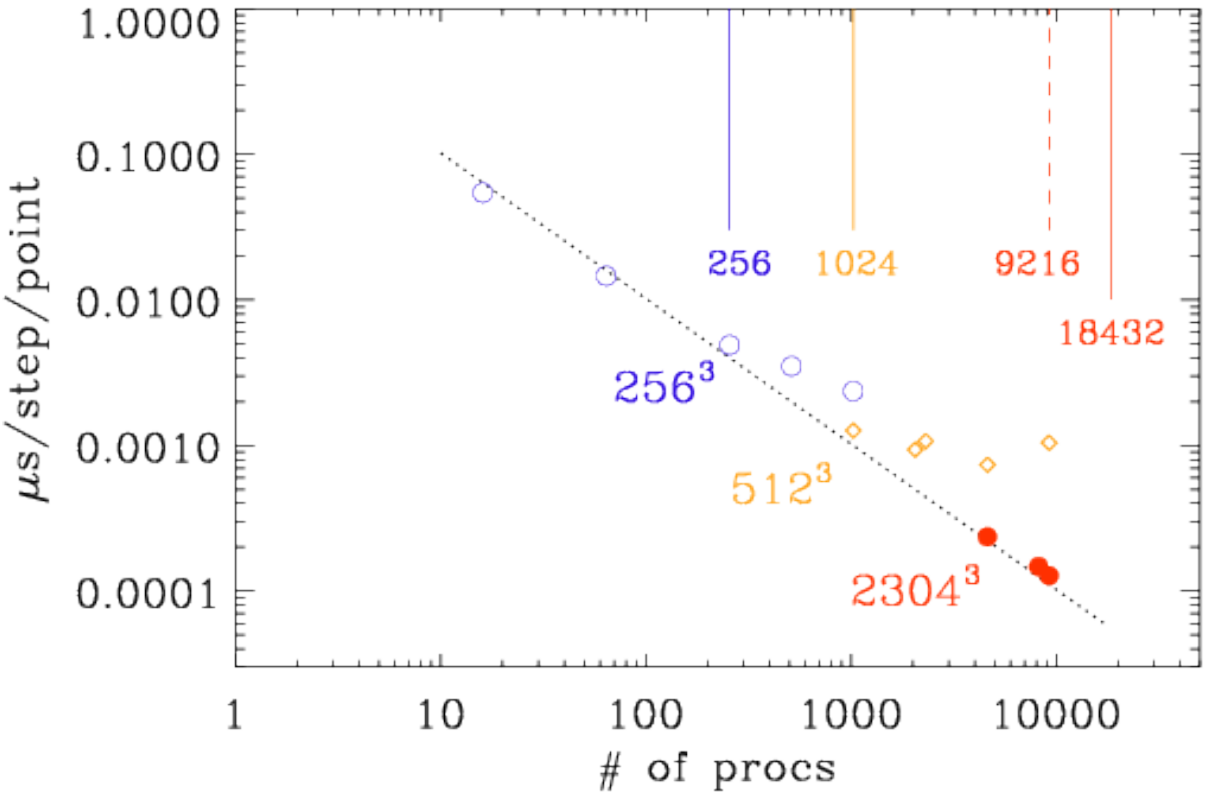

Figure 1: Strong scaling on Triolith.

Table 1: Triolith timings

proc	$\frac{\mu s}{\text{pt step}}$	resol.	layout
16	0.075	128^3	2x2x4
16	0.065	128^3	1x4x4
16	0.0544	256^3	1x4x4
64	0.0146	256^3	1x8x8
64	0.0164	256^3	2x4x8
256	0.0049	256^3	1x16x16
512	0.0035	256^3	2x16x16
1024	0.00236	256^3	2x16x32
1024	0.00127	512^3	2x16x32
1024	0.00129	512^3	4x16x16
2048	9.34×10^{-4}	512^3	4x16x32
2304	0.00107	576^3	4x18x32
4096	3.6×10^{-4}	1024^3	4x32x32
4096	3.8×10^{-4}	1024^3	8x16x32
4096	4.2×10^{-4}	1024^3	4x16x64
4608	7.38×10^{-4}	576^3	8x18x32
4608	2.66×10^{-4}	1152^3	4x32x36
4608	3.03×10^{-4}	1152^3	4x36x32
4608	3.12×10^{-4}	1152^3	4x18x64
4608	2.36×10^{-4}	2304^3	2x32x72
8192	1.475×10^{-4}	2048^3	4x32x64
9216	0.00104	576^3	16x18x32
9216	1.276×10^{-4}	2304^3	4x36x64
9216	1.30×10^{-4}	2304^3	4x32x72

5 Requested resources

Almost all the problems described above will principally use the Pencil Code[2], which is hosted by Google–Code since 2008[3]. This is an open-source code developed by myself, my current and former coworkers, some of whom are part of this project, as well as others that have been invited to join the effort. The performance of this code has been discussed at several international conferences; see, e.g., `http://www.nordita.org/~brandenb/talks/misc/PencilCode09.ppt`. The code has been optimized over the years and is still being improved in terms of performance and new features are also being added. All of the 21,209 revisions since 2001 are publicly available through our svn repository. We have adapted and optimized this code for spherical polar coordinate system (**?**). This addition to the code is used in several of the problems listed in the previous section. The code runs well on all the different platforms. This time we are applying for resources in the four machines Lindgren, Abisko, Gardar, and Triolith.

On Lindgren, we run production runs with $1024^2 \times 1536$ meshpoints on 6144 cores, while on Abisko, Gardar, and Triolith, most of our production runs tend to have 512^3 meshpoints and

[2] `http://www.nordita.org/software/pencil-code`
[3] `http://pencil-code.googlecode.com`

can require typically 512 processors. A typical run requires at least $500,000$ time steps, but it can sometimes be much more, depending on circumstances. With $4.2\times10^{-4}\mu$s per meshpoint and per timestep on Lindgren, this means 4 days of wallclock time at a cost of 600,000 CPU hours, while with $3.5\times10^{-3}\mu$s per meshpoint and per timestep on Abisko, Gardar, or Lindgren, this means 3 days of wallclock time at a cost of 30,000 CPU hours per run.

To address properly the critical question of the dependence on the magnetic Reynolds number we have to use high resolution runs. As we move from 256^3 and 512^3 to $1024^2\times1536$ mesh points (and correspondingly higher magnetic Reynolds numbers), we see the emergence of small-scale dynamo action at all depth. This does not yet affect the 512^3 runs, where the red line shows still a well-developed maximum of $\overline{B}/B_{\rm eq}\approx1$, but for the $1024^2\times1536$ the maximum is now only one third of that. We expect that this value will not decrease further, and that it will actually become bigger at larger stratification, but this needs to be shown. Note that the last of these runs is for a deeper domain, so as to include more safely the deep parts where it is important to reach values of $\overline{B}/B_{\rm eq}$ below 0.01, but this appears not to be possible due to small-scale dynamo action.

To confirm our ideas and to understand the effects of small-scale dynamo action, we plan to perform about 2 big runs per month on Lindgren, which requires at least 1000 kCPU hours, and about 5 intermediate ones on the other 3 machines, which requires 150 kCPU hours on each of them.

Computationally, all machines are comparable, but there can be unpredictable changes that hamper scientific progress. Most important is the waiting time in the queue and occasional opportunities when jobs start immediately. On Abisko and Triolith, the disk quotas restrict the ease with which we can run, while on Gardar there have been several periods when the machine was not functioning properly.

References

Mitra, D., Tavakol, R., Brandenburg, A., & Moss, D. 2009, ApJ, 697, 923