


Tests with 8 and 10 filters

Sebastien produced calibrated Stokes profiles for some data series taken with 8 and 10 filters (instead of 6).



Keiji changed the wrapper to adjust to this possibility and Sebastien modified his filter computing routine.

I made some changes in VFISV to allow for an 8 or 10 filter case. Basically, I changed:

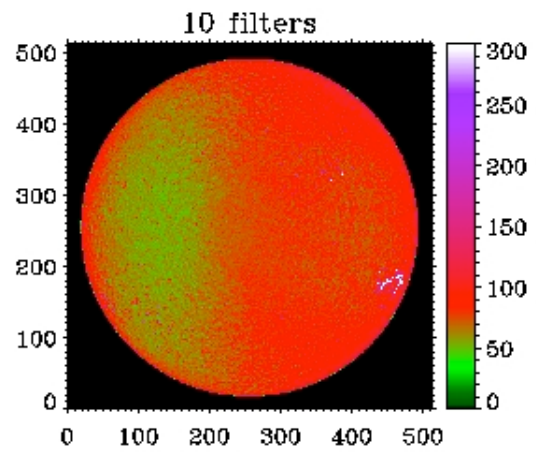
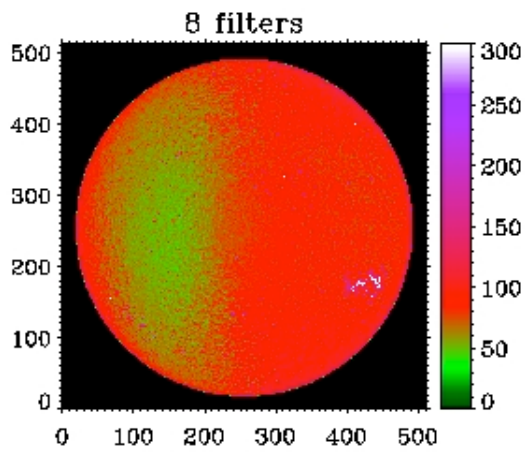
- wfa_guess.f90: to have the tuning positions for the 3 cases (6, 8 and 10 filters) depending on the value of NBINS.

I also changed all occurrences of NTUNE for NBINS.

- filter_init.f90: There were still some remnants of Juanma's filter computation. NTUNE and NBINS can be equal or different, and the combinations of both were dependent on the parameter CONTINUUM. There was an IF statement that checked whether the values of the three parameters were consequential or not, but only checked for the 6 spectral line positions or the 5+continuum. It would abort the run in any other case (so 8 or 10 filters were not considered). I deleted all of this and simplified the routine so that it only allocates memory for the filters.

They are currently running and being stored in: hao_rce.vmagf_test_mar07

We've inverted only one dataset with 8 filters and one with 10, so we won't get the temporal evolution. However, just looking at the magnetic field strength over the full disk, we can see the spatial patterns that were bothering us. The panels show the magnetic field strength saturated at 300 gauss, for the 8 filter case (left) and 10 filter case (right).



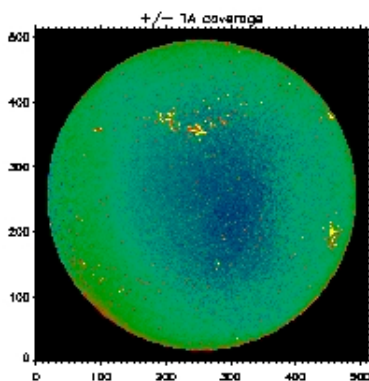



March 21, 2011

Filter coverage:



I'm resuming the analysis of the tests for filter coverage with +/- 1 A, +/- 2 A and +/- 3 A with respect to the center of the spectral line. All the tests were done with a 27 mA sampling, a set of weights of [1,3,3,2] and a convergence criteria of 1d-7.

2010.15.15_04:00:

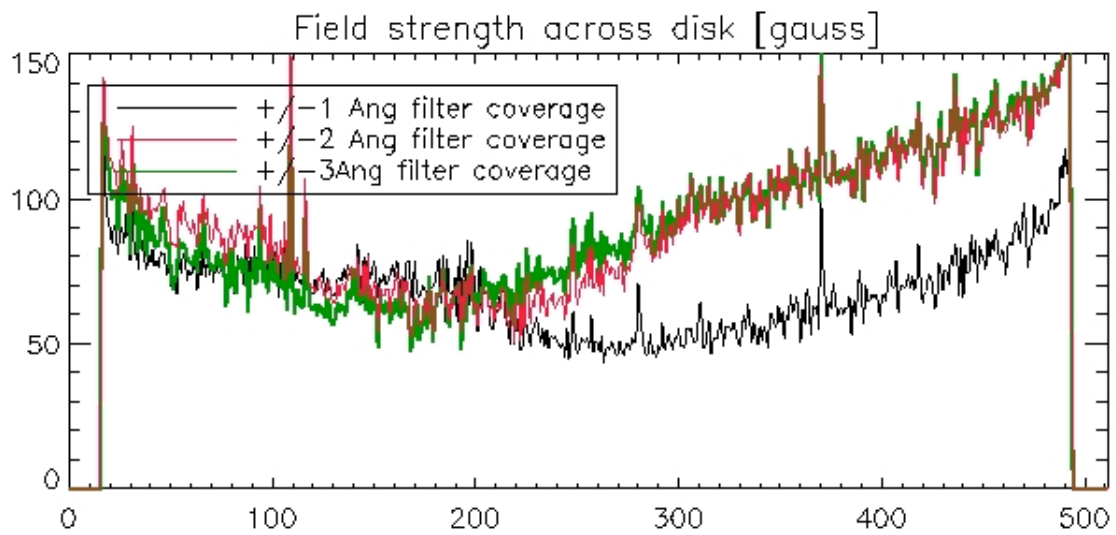





2010.15.10_08:00




2010.15.10_12_00

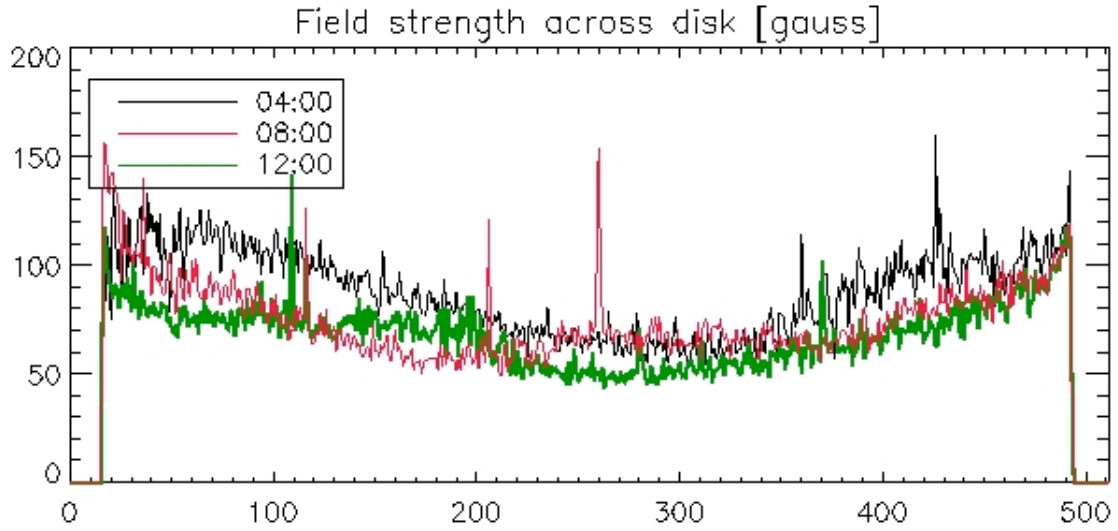


March 22, 2011



Temporal variation: same plots, but showing temporal variation for one coverage range.

For 1 Angstrom coverage and 3 times (04:00, 08:00 and 12:00):





For the 3 Angstrom coverage and 3 times (04:00, 08:00 and 12:00):



Qsub commands:

qstat allows you to see the status of your job in the queue.

-u user (qstat -u rce) gives you the jobs submitted by the user
-pri lists the priority of the jobs
-ext gives a detailed output

qdel allows you to delete a job from the queue (qdel JobID).

Keiji's email on the subject:

Hi, Rebecca

I logged in a cluster machine (c11n033) on which your job is running. What I found is that the machine is processing your job AND Rick's helioseismology modules....

The job queue system can assign one 8-cpu job to each machine exclusively, but under the assumption that 1 cpu job will be short lasting, 1-cpu job can be assigned to the machine that already is taking care of 8-cpu job.

This results in doubling the elapsed time..

Usually, the 8-cpu jobs are assigned from c11n062, c11n061 ... while the 1-cpu jobs are done from c11n008, n009 ..., so, if so many jobs are put queued, such conflict may occur.

Beside that the CPU-conflict, it seems, the standard-output from your job /home/rce/filter_sampling_march2011/dosyninv.sh.o5909145 is saying that yet your run is handling the first full-disk data.

I also took a look at your script, and found it is totally correct. Only one concern is that in your script vfvsv exectubale is called without specifying that ./vfvsv is the one at the current working directory (but I'm skeptical it is the reason why so slow..)

In this way, what I think to recommend you is, split your script into 24 scripts, each of which just run one case. I split yours and put them to the directory; /home/keiji/cvs/JSOC/proj/myproj/inversion_mpi/to_rebecca_11mar22 where you will find 24 script, and one script doqsub.sh which send the 24 script to the queue system. A tweak made in the doqsub.sh is that it send your job only at large-number machine (from c11n045 to c11n062), and wait if these are already used. This strategy will avoid the conflict with 1CPU jobs, so it would work.

Hope this will help.

Keiji

I finally managed to get some of the jobs for the fine spectral sampling test running.

