

Selecting Full-Disk Ambiguity Resolution Parameters for the HMI Pipeline

© Springer ●●●●

The target for disambiguating the full disk every 12 minutes is to be able to keep up 90% of the time using a single 12-core node in the a-queue. Thus the objective here is to determine the schedule for the annealing which results in the minimum energy ($E \propto \sum(|J_z| + |\nabla \cdot \mathbf{B}|)$) within that time constraint. Specifically, a full-disk disambiguation with the 90th percentile of pixels to anneal can take up to $12 \cdot 720 \text{ s} = 8640 \text{ CPU} \cdot \text{s}$.

First, determine the 90th percentile of the number of pixels to anneal. Using typically one good quality image per day from 2011.07.01 to 2013.06.30 the number of pixels to anneal as a function of time is shown in Figure 1. The 90th percentile of the number of pixels to anneal occurs at approximately 8.7×10^5 pixel, while the median occurs at approximately 7.2×10^5 pixel.

Five dates have been run at both Stanford and NWRA: 2011.02.01 at 00:00 TAI, 2012.01.01 at 00:00 TAI, 2012.06.01 at 00:00 TAI, 2012.07.21 at 19:00 TAI and 2013.05.16 at 18:48 TAI. To estimate the run time at NWRA for the 90th percentile of pixels to anneal, a comparison of these dates run with the same annealing schedule at Stanford and at NWRA is shown in Figure 2. Given the limited sample of comparisons made thus far, it appears that multiplying by a factor of 5 is a reasonable approximation to estimate the run time at NWRA.

To make a first estimate of the optimal parameters, consider 2012.01.01, which is very close to the 90th percentile in the number of pixels to anneal, and thus should run at Stanford in approximately $8.6 \times 10^3 \text{ s}$, corresponding to a run time at NWRA of approximately $4.3 \times 10^4 \text{ s}$. The final energy from the annealing algorithm as a function of run time for a variety of annealing schedules is shown in Figure 3 for this date. The figure also shows the energy interpolated to the targeted run time at NWRA as a function of the annealing parameter AMBNEQ. For the specified run time, the minimum energy occurs at approximately $\text{AMBNEQ} \approx 100$, with a corresponding $\text{AMBTFCR} \approx 0.978$.

The scaling of the run times at NWRA and Stanford with number of pixels to anneal for an annealing schedule close to this ($\text{AMBNEQ} = 100$, $\text{AMBTFCR} = 0.98$) are shown in Figure 4. Over the range of interest, they are approximately linear and given by

$$\begin{aligned} \text{time} &= 2.31 \times 10^4 + 0.0264 n_{\text{pix}} && \text{NWRA} \\ \text{time} &= 2.56 \times 10^3 + 0.0082 n_{\text{pix}} && \text{Stanford} \end{aligned} \quad (1)$$




Figure 1. Number of pixels to anneal as a function of time. The dotted horizontal line marks the 90th percentile which falls at a value of 8.7×10^5 pixel; the dashed horizontal line marks the 50th percentile (median) which falls at a value of 7.2×10^5 pixel.

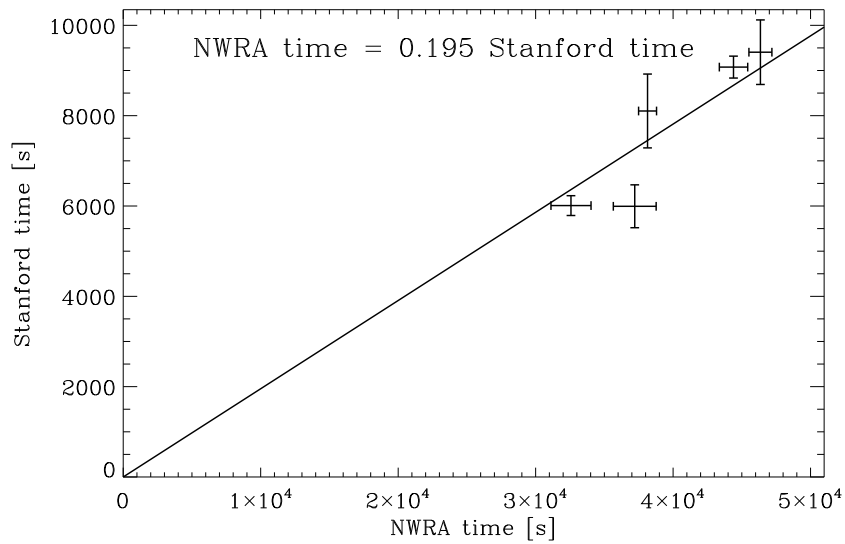


Figure 2. Comparison of run time at NWRA and Stanford for $AMBNEQ = 100$, $AMBTFCR = 0.98$. The dashed line has a slope of 0.195. Error bars are the standard deviation over at least four random number seeds.




Figure 3. Annealing energy for 2012.01.01 at 00:00:00 TAI. *Top:* the final energy from the annealing as a function of the run time for a range of annealing schedules. Each color corresponds to a different value of the annealing parameter $AMBNEQ$; each point within a given color corresponds to a different value of the annealing parameter $AMBTFCTR$. Each point is the mean over different random number seeds (typically 4), and the error bar is the standard error (standard deviation divided by square root of number of random number seeds). The vertical dashed line marks the targeted run time of 4.3×10^4 s. *Bottom:* the final energy from the annealing linearly interpolated to the targeted run time as a function of the annealing parameter $AMBNEQ$. Each point is labeled with the interpolated value of the annealing parameter $AMBTFCTR$. The minimum energy occurs at $AMBNEQ \approx 100$, with a corresponding $AMBTFCTR \approx 0.978$.




Figure 4. Scaling of run time with number of pixels to anneal for $\text{AMBNEQ} = 100$, $\text{AMBTFCR} = 0.98$ at NWRA (left) and at Stanford (right). The solid line shows the best linear fit, while the dashed line shows the best power law fit. Over the range of interest, the linear fit is a reasonable approximation.

Table 1. Summary of selected dates.

date	# pixel ($c \geq 60$)	# pixel ($c=90$)	Stanford time	NWRA time
2012.07.21_19:00:00	4.23×10^5	0.42×10^5	5.5×10^3 s	2.8×10^4 s
2011.02.01.00:00:00	5.22×10^5	0.43×10^5	6.0×10^3 s	3.0×10^4 s
2012.06.01.00:00:00	6.72×10^5	0.96×10^5	6.8×10^3 s	3.4×10^4 s
2012.01.01.00:00:00	8.70×10^5	1.36×10^5	7.8×10^3 s	3.9×10^4 s
2013.05.16_18:48:00	8.71×10^5	1.67×10^5	7.8×10^3 s	3.9×10^4 s

from which we can estimate the run time for other dates. A selection of other dates is given in Table 1, along with the number of pixels to anneal and the estimated run times at Stanford and NWRA, as determined from the average of the fits for the two case with $\text{AMBTFCR} = 0.97, 0.98$. A value intermediate between these two is likely appropriate to keep up with the 90th percentile, but these are the closest cases that have been run. It also appears that a value of AMBNEQ between 100 and 200 might be optimal, but none have yet been run.

The final energy from the annealing algorithm as a function of run time is shown in Figure 5 for some of the additional dates listed in Table 1. For each date, the energy has been linearly interpolated to the targeted run time. In each case, to the extent which it can be determined at this point, the minimum in energy for the targeted run time occurs in the vicinity of $\text{AMBNEQ} = 100$, $\text{AMBTFCR} = 0.975$, thus it appears that this annealing schedule is close to optimal for all the cases considered, and hopefully therefore for all cases.

As another measure of how well we can keep up with the processing, use the expression in equation 1 to estimate the run time at Stanford for the same dates used to estimate the 90th percentile in number of pixels to anneal. Figure 6 shows the cumulative run time from this expression. Although there are intervals of a

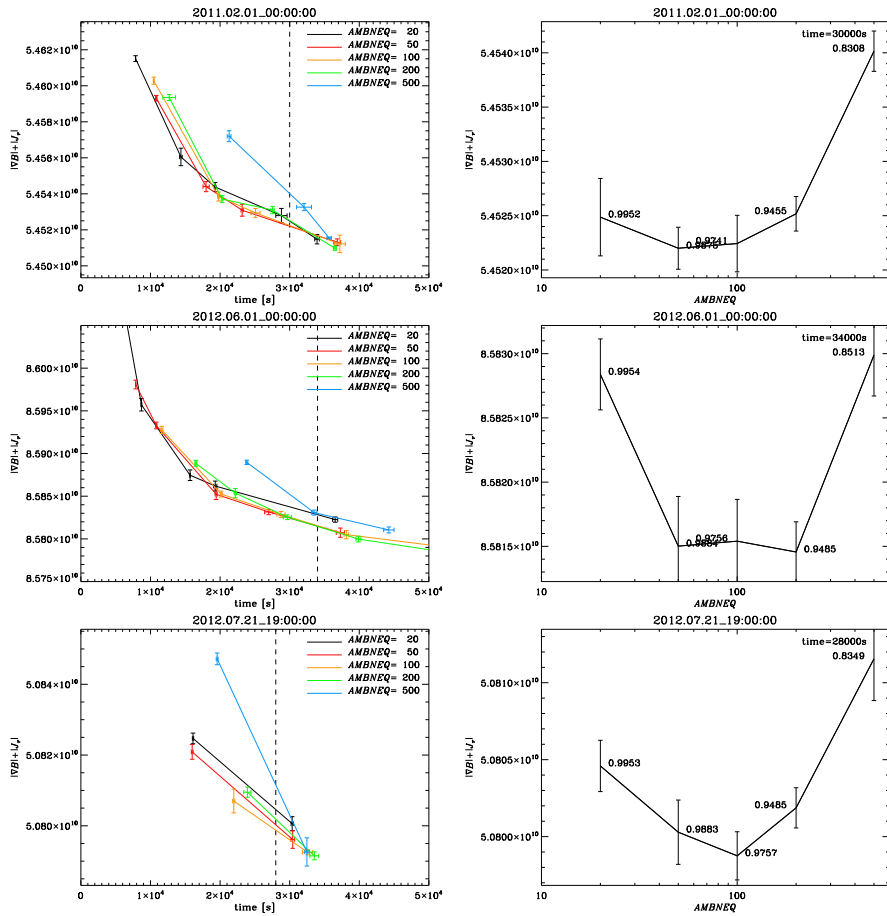


Figure 5. Annealing energy for 2011.02.01 at 00:00:00 TAI, (top), 2012.06.01 at 00:00:00 TAI, (middle) and 2012.07.21 at 19:00:00 TAI, (bottom), in the same format as Fig. 3. In the right panels, the time is that given in Table 1, as inferred from the linear fit to the run time.

few weeks when the processing falls behind, on average it is able to keep up, and generally is slightly ahead.

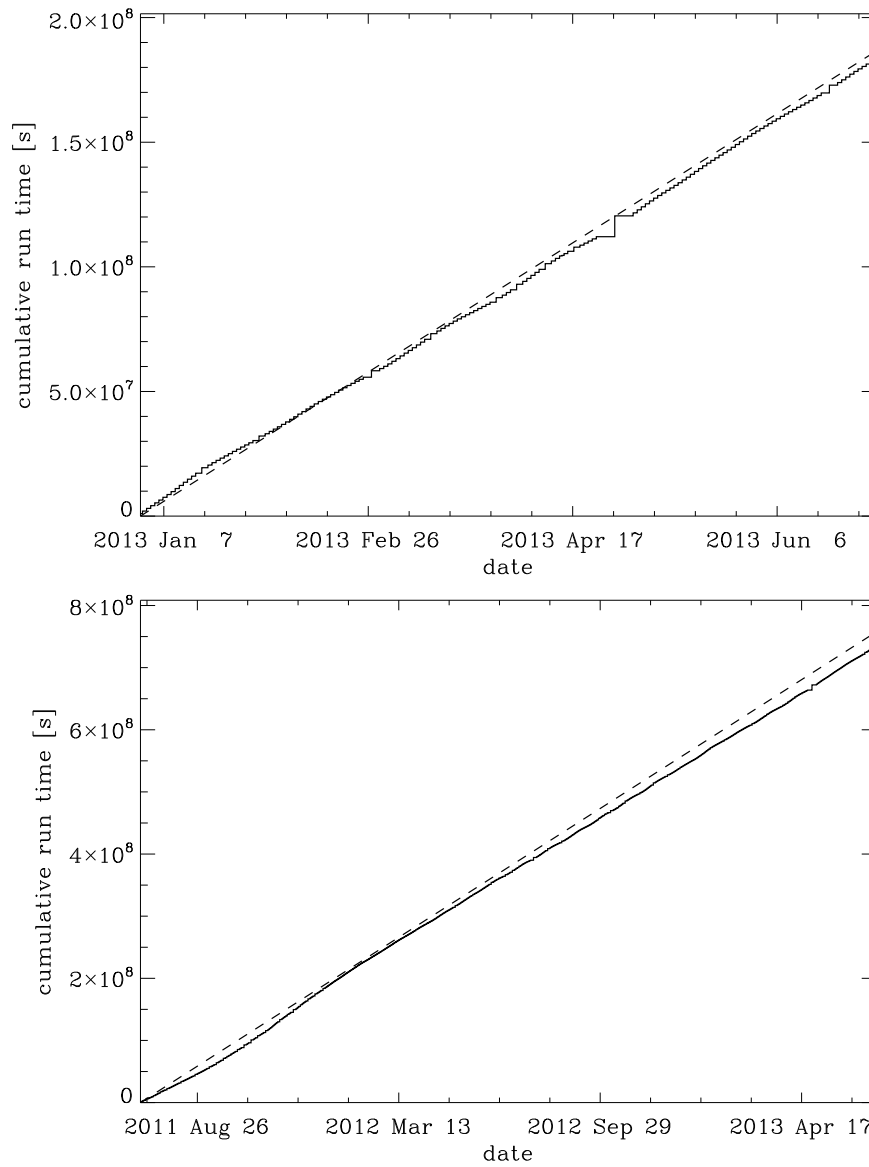


Figure 6. Cumulative run time at Stanford estimated from equation 1 (solid line) and the total amount of processing time available (dashed line). *Top:* for the first six months of 2013. *Bottom:* for a much longer interval.

Table 2. Fraction of pixels which consistently have the same disambiguation

date	AMBNEQ	AMBTFCR	CONF_DISAMBIG ≥ 60	CONF_DISAMBIG = 90
2011.02.01	100	0.96	0.929	0.981
	500	0.96	0.946	0.986
2012.06.01	100	0.96	0.931	0.951
	500	0.96	0.950	0.964
2012.01.01	100	0.96	0.941	0.980
	500	0.96	0.955	0.985

1. Other Measures of Performance

Todd asked about what is “good enough” in terms of reaching a minimum in energy. As an alternative to the energy, consider the number of pixels which have the same disambiguation for all the random number seeds run, as a measure of how believable the result is. Table 1 show the fraction of pixels for which the same solution is obtained for 14 different random number seeds for three dates and two annealing schedules, one with slightly shorter run time than the pipeline (AMBNEQ = 100) and one with moderately longer run time (AMBNEQ = 500). For pixels in which the annealed solution is returned (CONF_DISAMBIG = 90), at least 95% of the pixels have a consistent result for all the dates and annealing schedules.

Figures 7, 8 and 9 show the radial component of the magnetic field and contours of the confidence in the disambiguation for these three dates. In addition, the figures show the pixels for which no consistent solution is found with AMBNEQ = 100, AMBTFCR = 0.96 for small subareas. There are two scenarios in which inconsistent solutions are found: in quiet sun areas, a small patch of pixels above the noise threshold occasionally leads to a rectangle of inconsistent solutions, and in active region areas, there are occasionally small patches that show an inconsistent solution, much as in the HARP data (e.g., the checkerboard pattern). For 2012.06.01 (Fig. 9), which performs the worst by this measure, much of the strong field and corresponding areas of inconsistent results are very close to the limb. Aside from this tendency for worse performance very close to the rim, there is no obvious spatial pattern to the inconsistent pixels, and there are no large areas of contiguous inconsistent pixels.

2. Recommendations

- Compare run times at NWRA and Stanford for AMBNEQ = 100, AMBTFCR = 0.98 for dates examined here. **Done.** See Fig. 2.
- Decide how important it is to keep up with the 90th percentile. For example, is something like the 85th percentile acceptable as we head towards solar minimum?

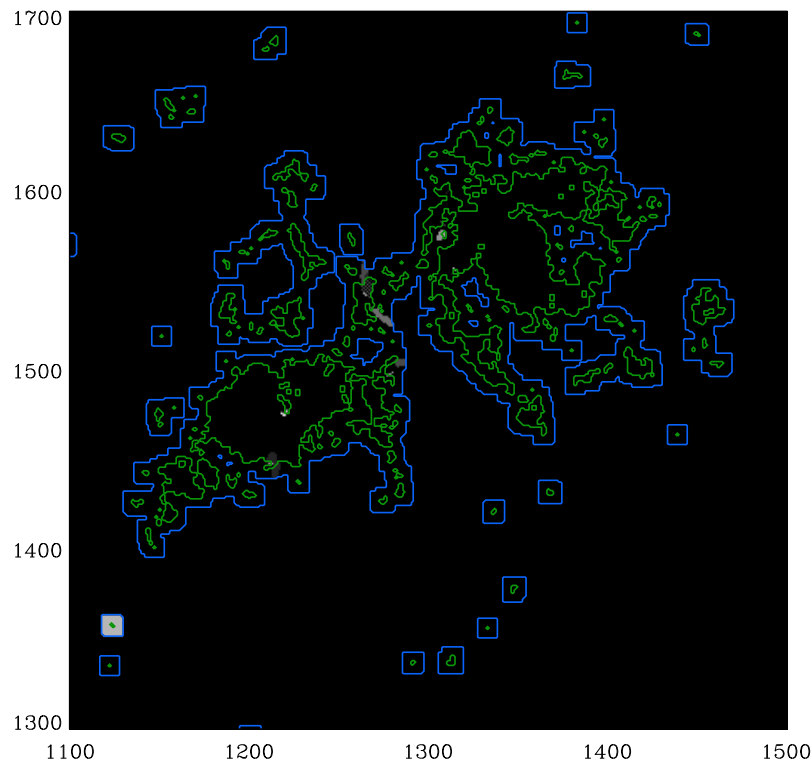
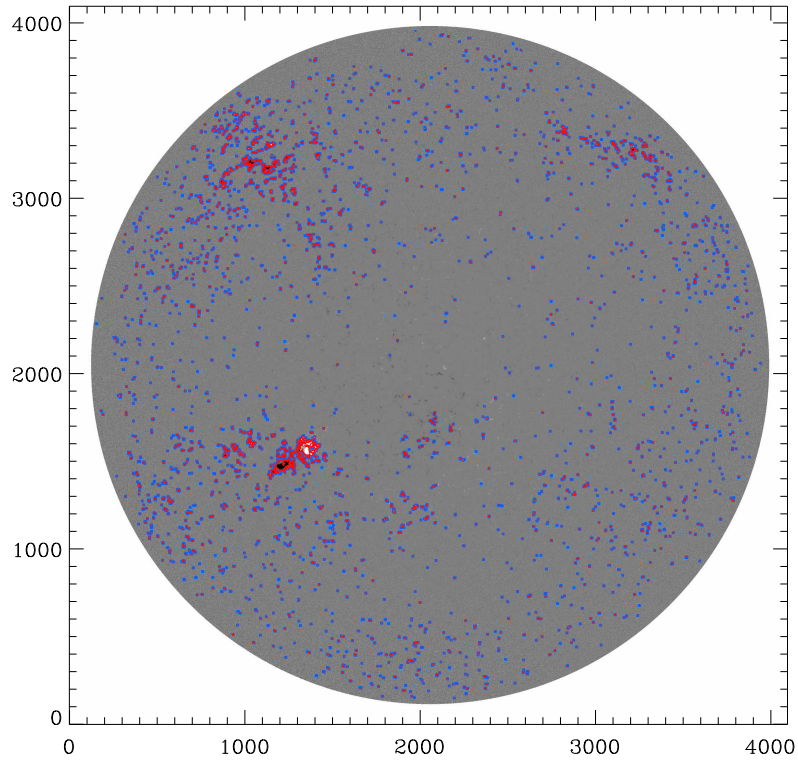


Figure 7. 2011.02.01 at 00:00:00 TAI. *Top:* radial component of the magnetic field, scaled to ± 500 G, with contours of CONF_DISAMBIG at values of 60 (blue) and 90 (red). *Bottom:* fraction of pixels which return the same solution between different random number seeds, with black being all seeds in agreement, and white being equal numbers of seeds with each solution.





Figure 8. 2012.01.01 at 00:00:00 TAI in the same format as Fig. 7.





Figure 9. 2012.06.01 at 00:00:00 TAI in the same format as Fig. 7.