Selecting Ambiguity Resolution Parameters for the HMI Pipeline

1. NRT Data

According to T. Hoeksema, the target for disambiguating the Near Real Time (NRT) HARPs is 90% of the HARPs to be complete within one hour using a single 8-core node. Thus the objective here is to determine the cooling schedule for the annealing which results in the minimum energy ($E = \sum(|J_z| + |\boldsymbol{\nabla}\cdot\boldsymbol{B}|)$) within that time constraint.

First, determine the properties of the 90th percentile. Using one image per day at 00:00:00 TAI from 2012.07.24 to 2012.09.17 (all the days which are presently available) from the series hmi.Mharp_702s_nrt, the distributions of area are shown in Figure 1. The 90th percentile of the maximum area of an individual HARP occurs at 1.36×10^6 pixel, while the 90th percentile of the total area of all the HARPs at a given time occurs at 2.51×10^6 pixel. However, **the small sample size means that there is considerable uncertainty in these estimates**.

Consider two limiting factors: the time to run a single HARP, and the time to run all the HARPs. For the first factor, the run time for 1.36×10^6 pixel should be (approximately) 3600 s. To make a first estimate of the optimal parameters, consider HARP 1256, which is closest in area of the HARPs considered in detail (see Table 1), although still smaller than the 90th percentile. For this HARP, the final energy from the annealing algorithm, as a function of run time is shown in Figure 2.[1] By inspection, the minimum energy for a run time around 3600 s occurs for $10 \lesssim \mathrm{neq} \lesssim 100$.

Figure 3 shows the energy interpolated to a run time of 3600 s as a function of the cooling parameter neq. Each point is also labeled with the interpolated value of tfactr. For the specified run time, the minimum energy occurs at approximately neq = 20, with a corresponding tfactr ≈ 0.99.

Figure 4 shows a corresponding plot for HARP 1271, which is substantially smaller than HARP 1256. The time was selected such that the value of tfactr at neq = 20 is approximately the same as for HARP 1256. In this case, the minimum energy likely occurs for values of neq < 20, but the results for neq = 20 are still quite good.

The scaling of the run time with HARP area for this cooling schedule is shown in

---

[1] All the run times shown here include no smoothing, which is what I recommend for NRT HARPs.

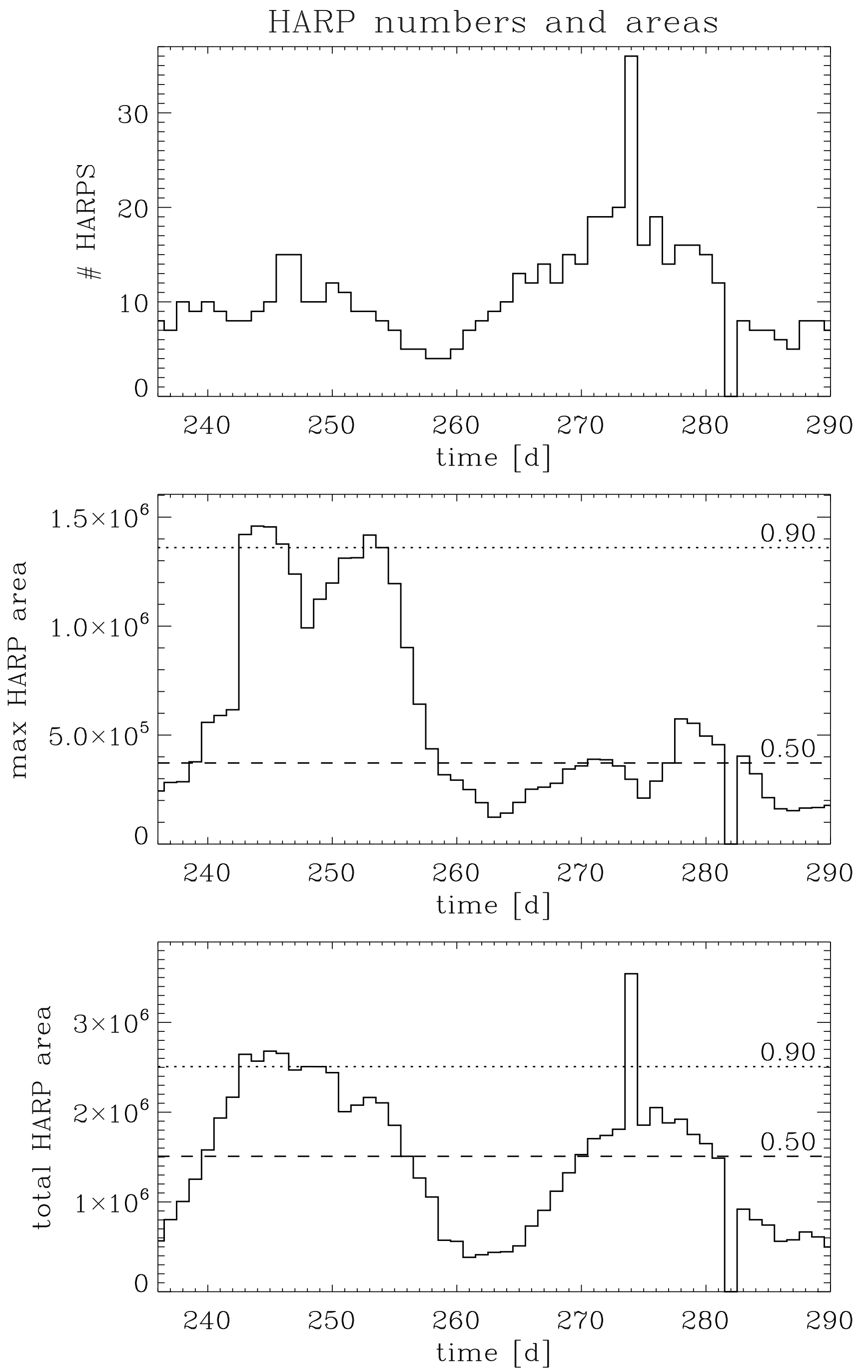

Fig. 1.— NRT HARP distributions. The middle panel shows the area in pixels of the largest HARP at that time. The bottom panel shows the sum of the areas of all the HARPs at that time. Time of 0 corresponds to 01 January 2012. The dotted horizontal lines show the 90th percentile; the dashed horizontal lines show the 50th percentile (median). The top panel shows the number of HARPs, although that is not used in this analysis.

Table 1. HARP area

HARP #	area (pixel)	time (NRT)	time (definitive)
1124	742000	1736.58	
1237	299719	583.293	
1241	62720	88.7842	
1246	15093	15.9886	
1248	21131	23.9715	
1249	194233	346.064	
1256	921393	2253.58	
1258	32844	40.7578	
1269	9450	9.10091	
1270	63187	86.9431	
1271	124092	201.827	
1272	21120	23.9565	
1273	5820	5.07851	
1274	21008	23.8037	
16	93240	143.079	

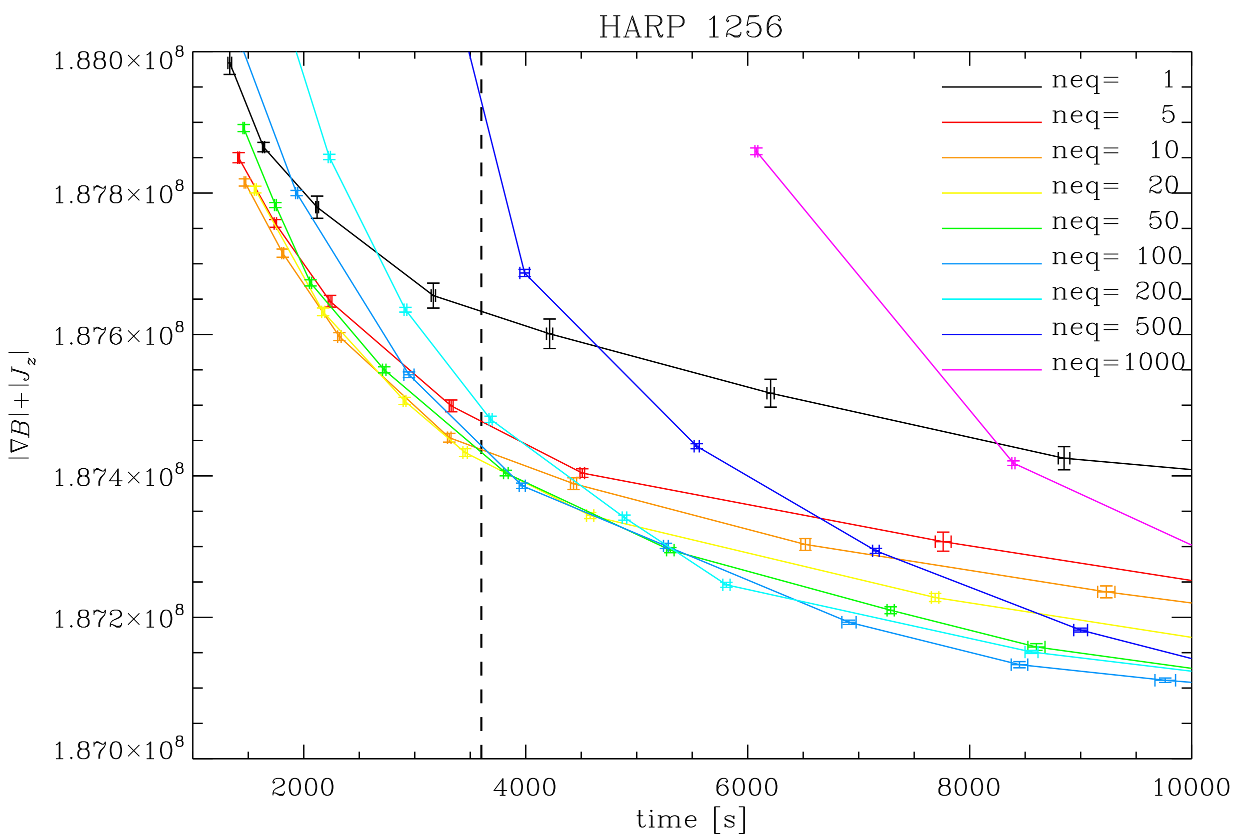

Fig. 2.— Annealing energy for HARP 1256 as a function of run time. Each color corresponds to a different value of the cooling parameter neq; each point within a given color corresponds to a different value of the cooling parameter tfactr. Each point is the mean over different random number seeds (typically 25), and the error bar is the standard error (standard deviation divided by square root of number of random number seeds). The vertical dashed line is at a run time of one hour.

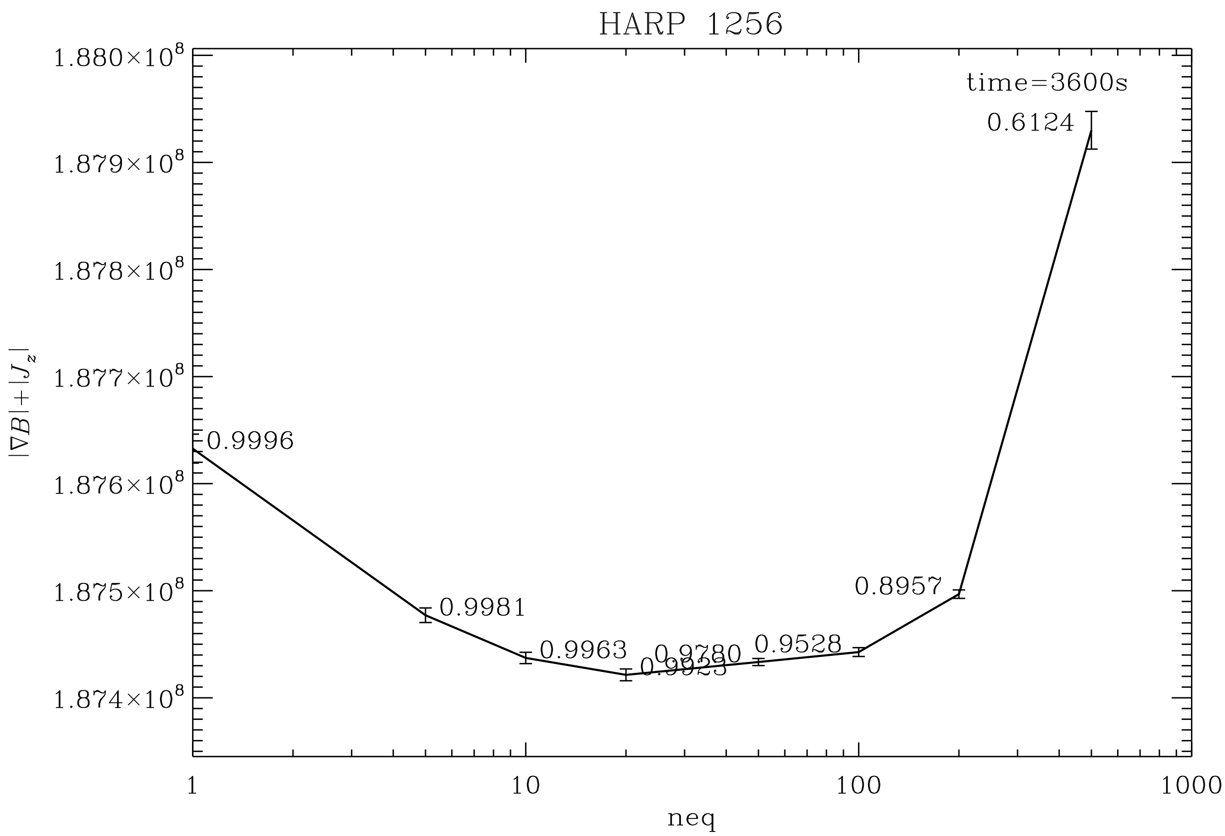

Fig. 3.— Annealing energy for HARP 1256 interpolated to a time of 3600 s. Each point is labeled with the interpolated value of tfactr. The interpolation is a simple linear scheme, and so has systematic biases, given the constant sign of the curvature of the curves shown in Fig. 2. Error bars are estimated by propagating the uncertainty from the different random number seeds through the interpolation, but do not account for the limitations of the linear interpolation.

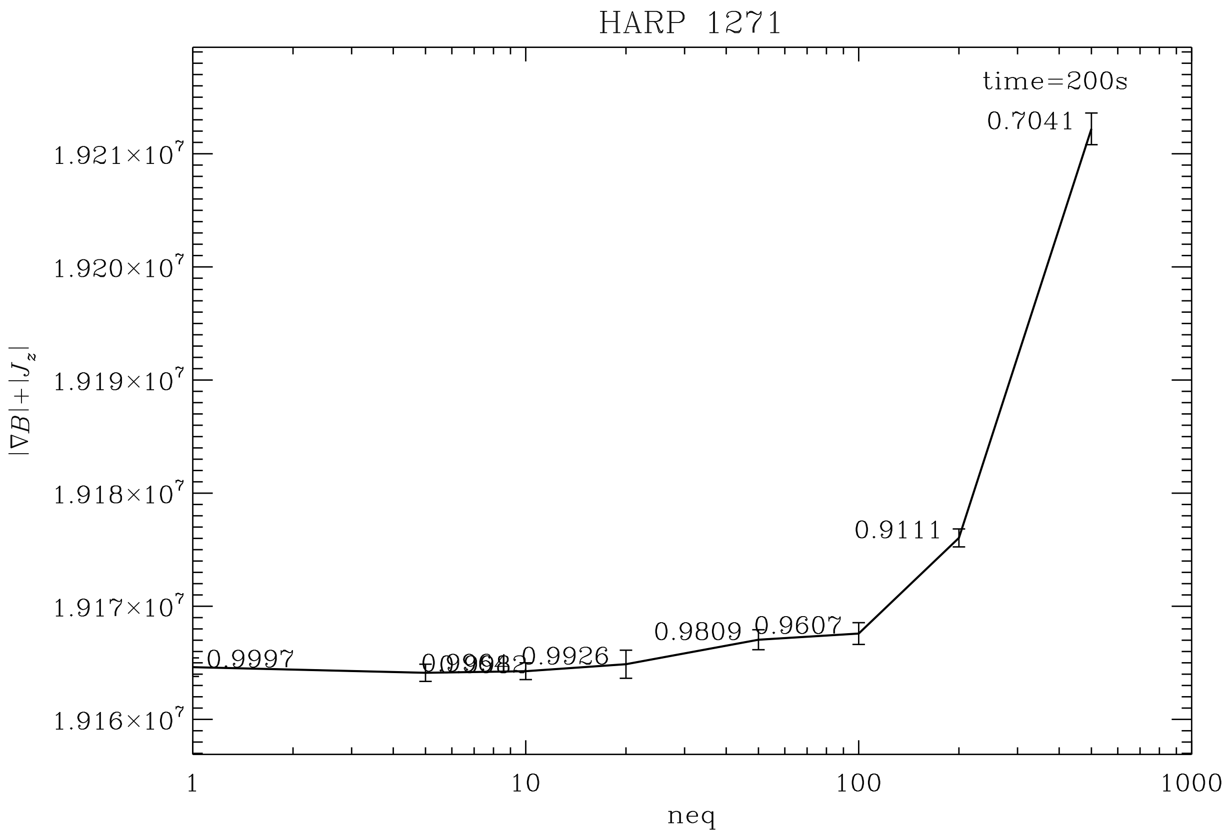

Fig. 4.— Annealing energy for HARP 1271 interpolated to a time of 200 s, in the same form as Fig. 3. The value of tfactr at neq = 20 is approximately the same as for HARP 1256.

Figure 5. It is approximately a power law, of the form

$$\text{time} \quad = \quad 1.50 \times 10^{-4} n_{\text{pix}}^{1.20} \tag{1}$$

from which we can estimate the run time for each of the HARPs in Table 1. However, **note that the 90th percentile HARP area lies outside the range considered, so the time estimate for this area is very sensitive to the value of the power law index, which is not particularly well constrained**. Using this power law, a HARP of area 1.36×10^6 pixel will run in approximately 3600 s.

The second limiting factor, the time to run all the HARPs, is likely automatically satisfied if the largest HARP is complete in one hour. That is, if we take the run time to be proportional to the total number of pixels in the HARPs (which is not a great approximation, given the scaling law, but should overestimate the total run time), and note that the 90th percentile of the total area occurs at less than twice the value of the 90th percentile of the area of the largest single HARP, then two cores should be sufficient to process all the HARPs 90% of the time, with typically one core processing the largest HARP, and the remaining core processing all the smaller HARPs. However, we can make a better estimation of the total run time by using the estimate of the run time given by equation 1 combined with the known NRT HARP areas. The results are shown in Figure 6. The 90th percentile in total run time is approximately 5200 s. For comparison, $8\,\text{CPUs} \times 720\,\text{s} = 5760\,\text{CPU} \cdot \text{s}$, so a single 8-core node will typically be able to keep up with processing the NRT HARPs 90% of the time, with the processing complete within one hour 90% of the time.

1.1. Recommendations

Two things would greatly help to constrain the best parameter choices:

- A much larger sample of NRT HARP sizes, to better determine the 90th percentile areas.
- A detailed analysis of a larger area HARP, so that the 90th percentile value does not lie outside the range of areas considered.

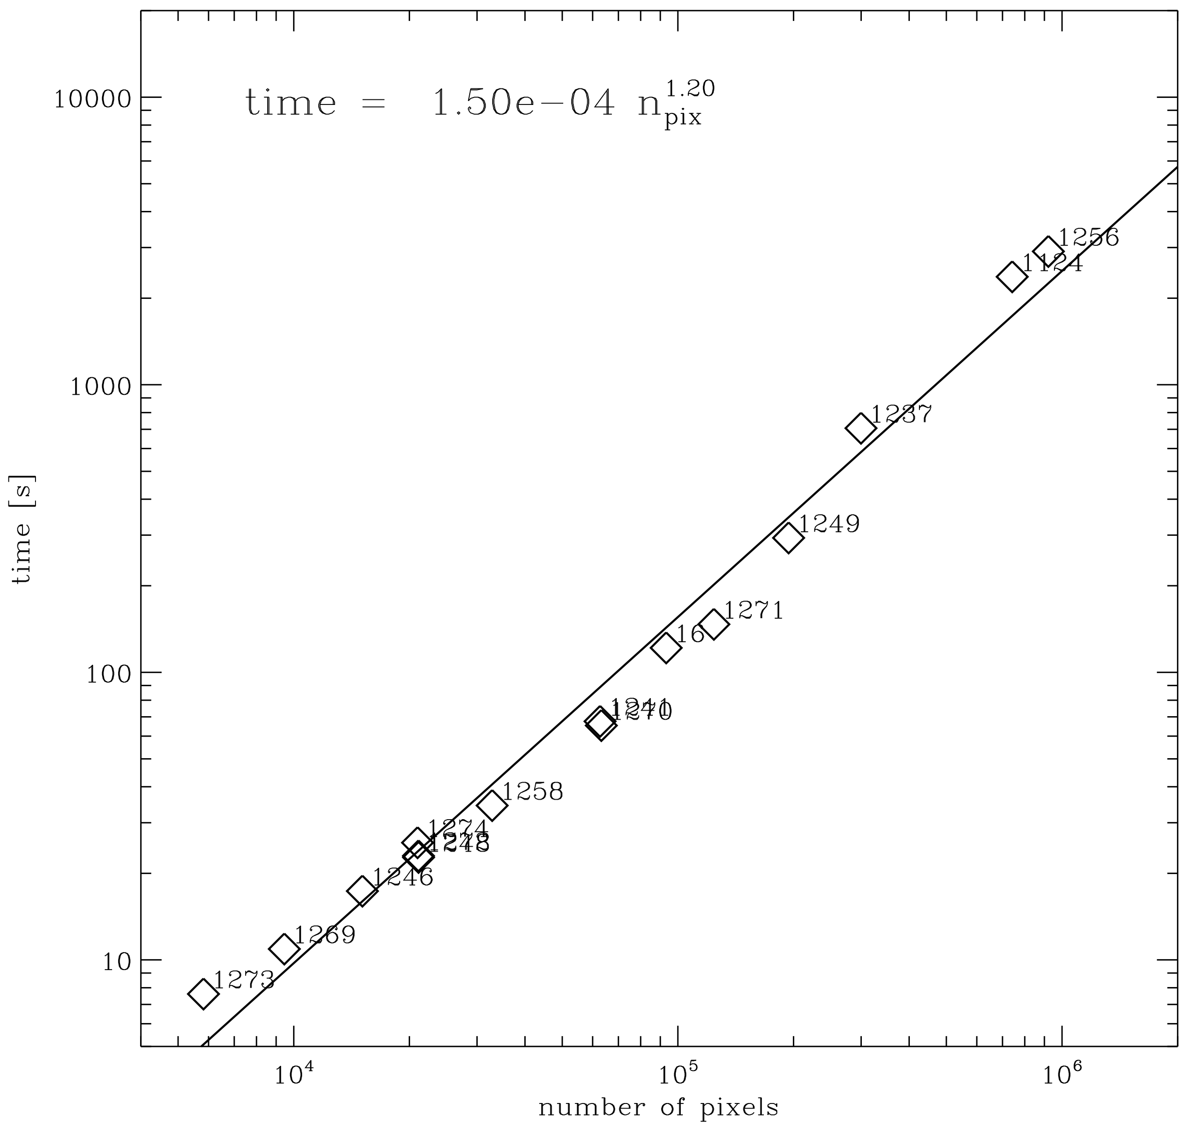

Fig. 5.— Scaling of run time with area of HARP in pixels for neq = 20, tfactr = 0.99.

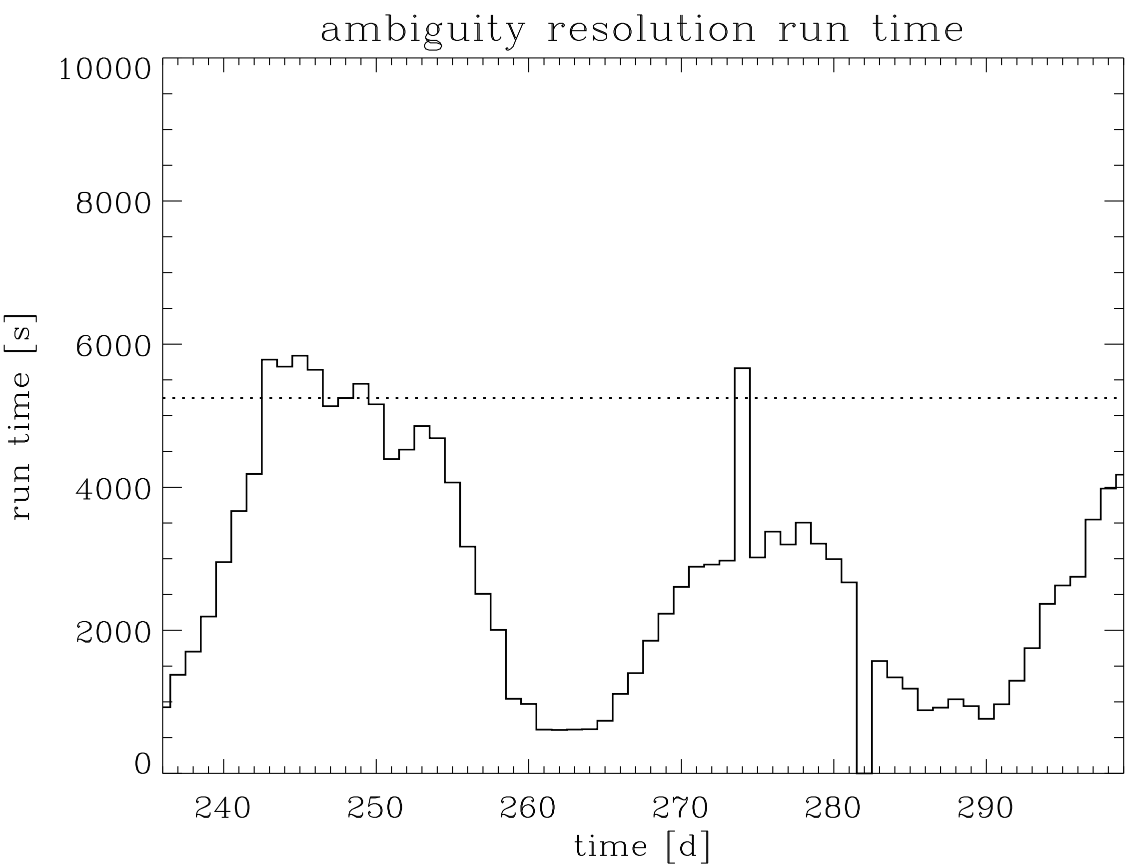

Fig. 6.— Distribution of run time for power law scaling of neq = 20, tfactr = 0.99. The dotted line shows the 90th percentile, which is at a run time of approximately 5200 s.

2. Other Measures of Performance

Todd asked about what is "good enough" in terms of reaching a minimum in energy. First, consider what the energy of the global minimum might be. Figures 7 and 8 show the energy as a function of run time for HARPs 1271 and 1237, including much longer run times than we have been considering for the pipeline.[2] In each case, the top plot shows a range of times that might realistically be considered for the pipeline. As run times are increased by orders of magnitude, the energy continues to decrease, showing that the global minimum in energy has not been reached. The behaviour is qualitatively similar for all the HARPs considered, although run times become prohibitively long for the largest HARPs.

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. The results of this for three HARPS are shown in Figure 9, and some examples for HARP 1237 are listed in Tables 2, and 3, for two different sets of annealing parameters corresponding to the parameters tentatively identified for NRT HARPs ($neq = 20$, $tfactr = 0.99$), and the parameters which produced the lowest energy ($neq = 100$, $tfactr = 0.992$). For this HARP, less than 1% of the pixels with field strengths above 300 G do not have the same ambiguity resolution for all 50 random number seeds considered. For field strengths above 150 G, the result is much worse: nearly half of the pixels with strong fields have no consistent answer, with an improvement from 48% to 45% when going from one set of annealing parameters to the other. This threshold of 150 G is important, because this is approximately the value of the noise estimated from Yang's fits to the field strength (100 G) plus the constant (50 G) Xudong has been using to generate the noise masks for the disambiguation.

Figure 10 shows the fraction of all pixels (i.e., with field strength threshold of 0.) for which no consistent answer is found as a function of the annealing energy for HARP 1271, and a wide range of annealing parameters. The fraction appears to decreases monotonically as the energy decreases, making the energy a reasonable proxy for the fraction. As the energy approaches a minimum, the fraction appears to decrease rapidly, suggesting that one might be able to estimate the minimum energy from this type of plot.

[2]Time estimates for these plots are imprecise because an assortment of different processors were used, and only 10 random number seeds were run for each cooling schedule.

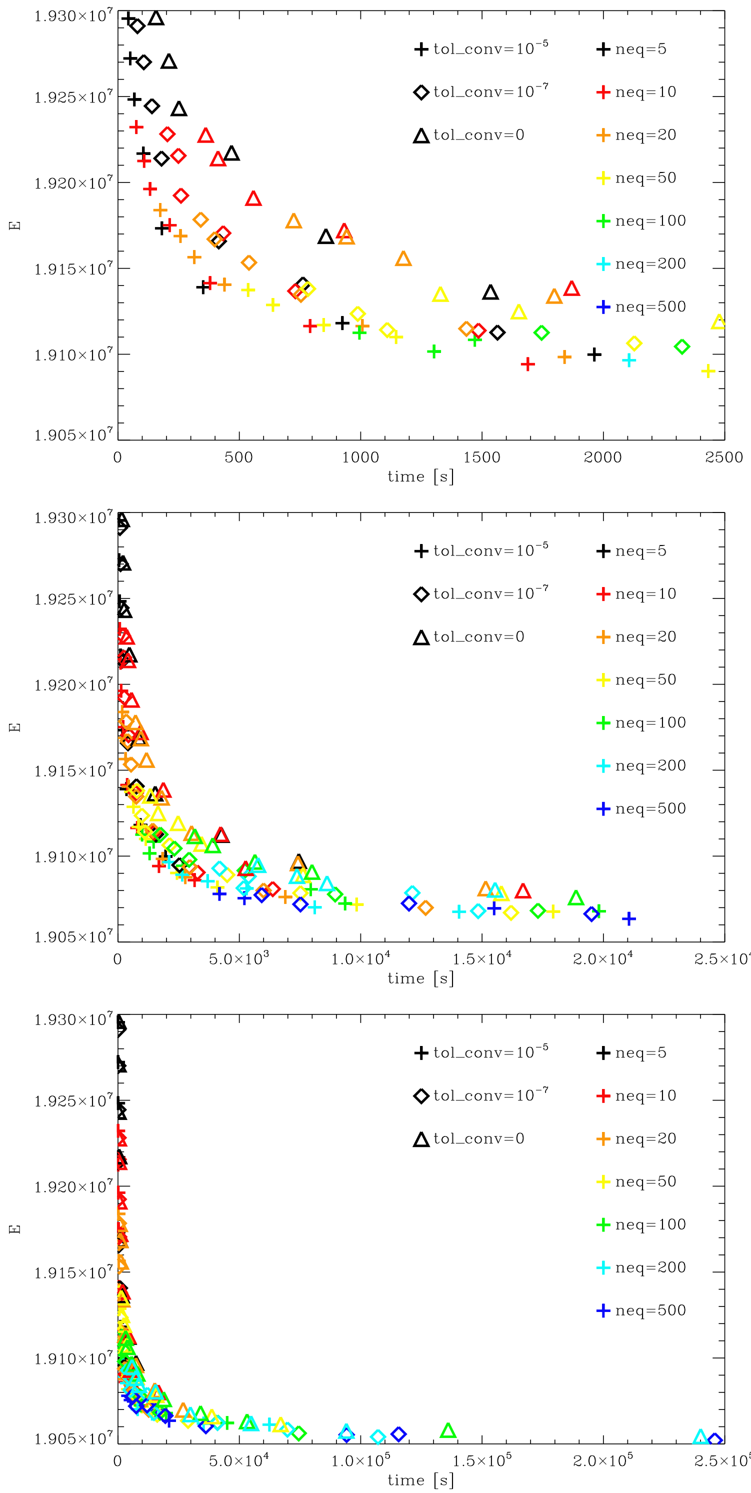

Fig. 7.— Annealing energy for HARP 1271, as a function of run time. The top plot shows a range of times that might realistically be considered for the pipeline (a time of ≈ 200 s for NRT and ≈ 500 s for definitive HARPs), and the x-axis range increases by a factor of ten for each subsequent row. The energy is still decreasing up to the longest times shown, suggesting that the global minium has still not been found.

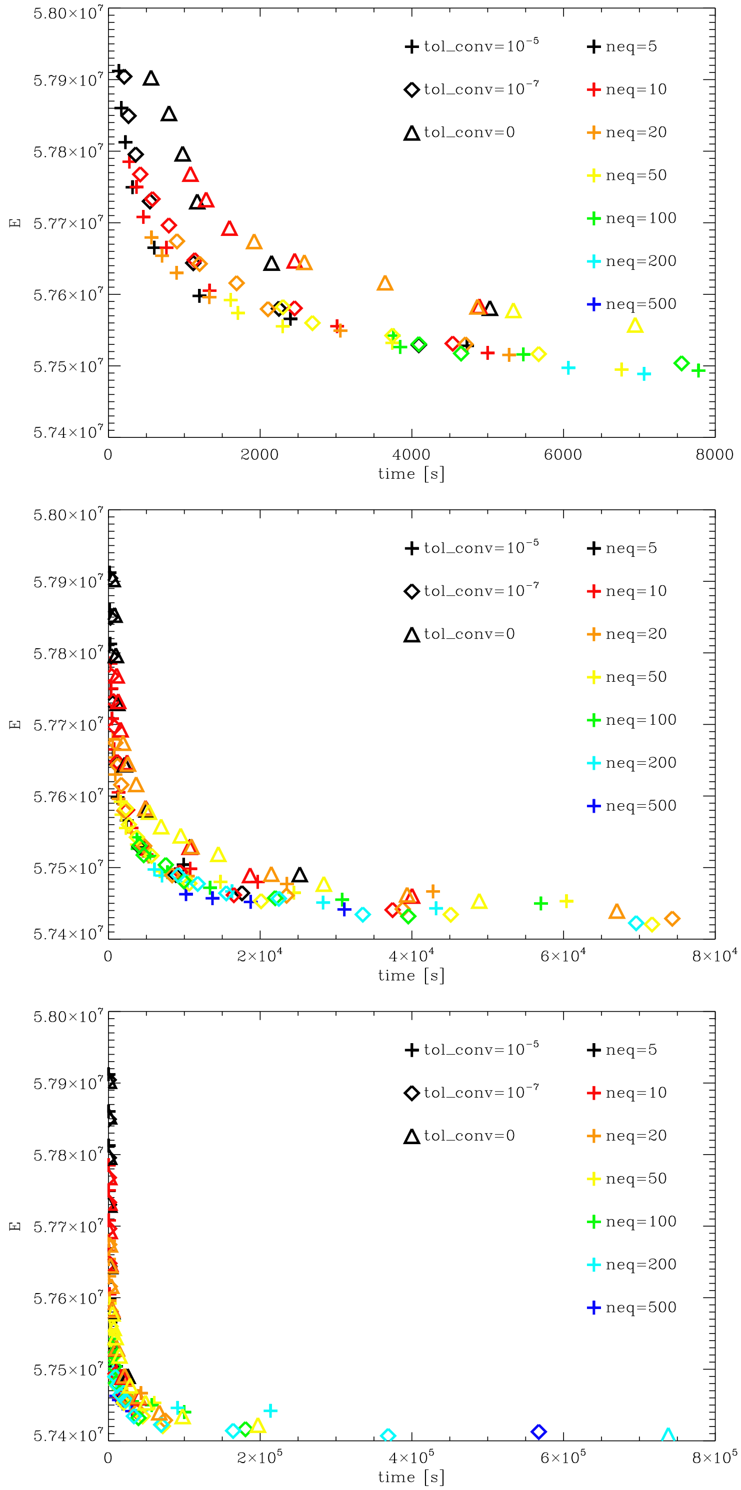

Fig. 8.— Same as Fig. 7 for HARP 1237, showing qualitatively similar behaviour.

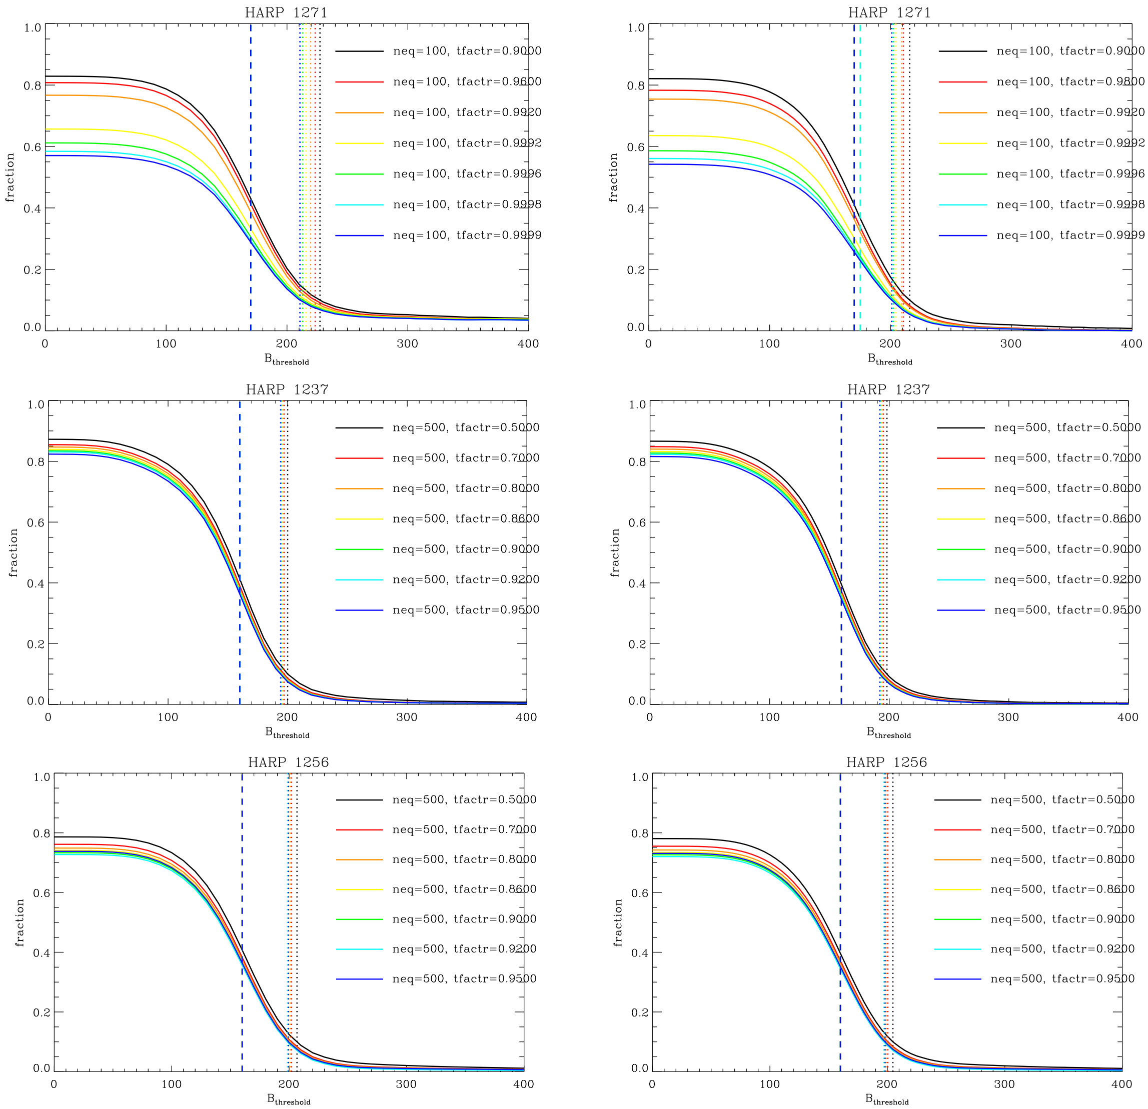

Fig. 9.— The fraction of pixels above a field strength threshold for which no consistent answer is found, as a function of the threshold, for HARPs 1271, 1237, and 1256 (HARP area increases from top to bottom), for a range of annealing parameters. Left: all pixels are included; right: a 10 pixel buffer is removed. The vertical dotted lines mark the field strength above which 10% of the pixels have no consistent answer. The vertical dashed lines mark the field strength at which the slope of the curve is a minimum. Interestingly, this occurs at approximately the same value for all the annealing parameters considered for a given HARP.

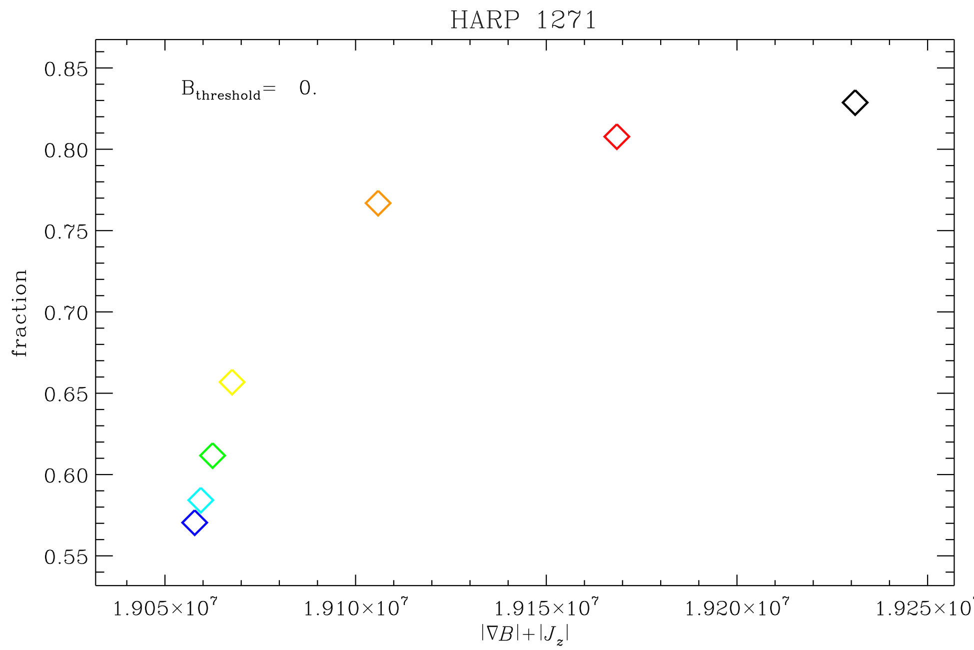

Fig. 10.— Fraction of pixels for which no consistent answer is found as a function of the annealing energy for HARP 1271, and a wide range of annealing parameters. Colors match the labels given in Fig. 9.

Table 2. Fraction "wrong" for HARP 1237, $neq = 20$, $tfactr = 0.99$.

threshold	# pixel	fraction of pixels	fraction above threshold
0.00000	254804	0.85014	0.85014
50.0000	230634	0.76950	0.84037
100.000	114694	0.38267	0.76278
150.000	21624	0.07214	0.48056
200.000	1380	0.00460	0.08217
250.000	175	0.00058	0.01512
300.000	67	0.00022	0.00752
350.000	40	0.00013	0.00573
400.000	26	0.00008	0.00466
450.000	18	0.00006	0.00398
500.000	8	0.00002	0.00221
550.000	4	0.00001	0.00135
600.000	2	0.00001	0.00081
650.000	0	0.00000	0.00000

Table 3. Fraction "wrong" for HARP 1237, $neq = 100$, $tfactr = 0.992$.

threshold	# pixel	fraction of pixels	fraction above threshold
0.00000	245573	0.819344	0.819344
50.0000	222046	0.740847	0.809085
100.000	109864	0.366557	0.730658
150.000	20434	0.0681772	0.454119
200.000	1233	0.00411385	0.0734191
250.000	156	0.000520488	0.0134843
300.000	68	0.000226879	0.00764217
350.000	43	0.000143468	0.00616311
400.000	28	9.34208e-05	0.00502332
450.000	19	6.33927e-05	0.00420168
500.000	9	3.00281e-05	0.00249377
550.000	5	1.66823e-05	0.00169779
600.000	2	6.67292e-06	0.000816660
650.000	0	0.00000	0.00000

3. Definitive Data

According to T. Hoeksema, the target for disambiguating the definitive HARPs is 80% of the HARPs, starting 400 days after the start of the mission, running on three 8-core nodes, which presumably amounts to $24\,\mathrm{CPUs} \times 720\,\mathrm{s} = 17280\,\mathrm{CPU} \cdot \mathrm{s}$. However, there is effectively no restriction on the longest run time, since these will already be delayed by several weeks, to allow a complete disk crossing to generate the HARP). Thus the objective here is to determine the cooling schedule for the annealing which results in the minimum energy ($E = \sum(|J_z| + |\boldsymbol{\nabla}\cdot\boldsymbol{B}|)$) within that time constraint.

First, determine the properties of the 80th percentile. Using one image per day at 00:00:00 TAI from 2011.06.01 to 2012.08.19 (from approximately 400 days after launch to approximately the latest date for which all the HARPs have been processed to date) from the series hmi.Mharp_702s, the distributions of area are shown in Figure 11. The 80th percentile of the maximum area of an individual HARP occurs at 0.96×10^6 pixel, while the 80th percentile of the total area of all the HARPs at a given time occurs at 2.53×10^6 pixel.

In this case, the only limiting factor is the time to run all the HARPs. If we take the run time to be proportional to the total number of pixels in the HARPs (which is not a great approximation, given the typical scaling laws, but will give a first estimate), then the time to run a 10^6 pixel HARP will be approximately $10^6\,\mathrm{pixel}/(2.5 \times 10^6\,\mathrm{pixel}/(17280\,\mathrm{s})) \approx 7{\times}10^3\,\mathrm{s}$. Again, consider HARP 1256 as the benchmark, and determine the optimal annealing parameters for this run time. Figure 12 shows the estimated energy at this run time as a function of neq. The minimum occurs near $neq = 100$, $tfactr = 0.98$. For the much smaller HARP 1271 (Figure 13), these values are also within a range which produces a good result.

The scaling of the run time with HARP area for this cooling schedule is shown in Figure 14. It is approximately a power law, of the form

$$\mathrm{time} \;=\; 4.06 \times 10^{-4} n_{\mathrm{pix}}^{1.20}. \qquad (2)$$

However, **note that there are some indications that this does deviate from a power law, with longer run times for large HARPs than would be predicted by the best fit power law**. This is potentially important, since it is the largest HARPs which dominate the processing time.

This scaling law can be used to estimate the run time for each individual HARP, based on its area, and combined with the HARP areas, to estimate the total run time at each time. The results of this, using the power law given in equation 2, are shown in Figure 15. The result is that the 80th percentile corresponds to approximately 13,500 s, while the target of 17000 s corresponds to the 88th percentile. This cooling schedule is slightly faster than

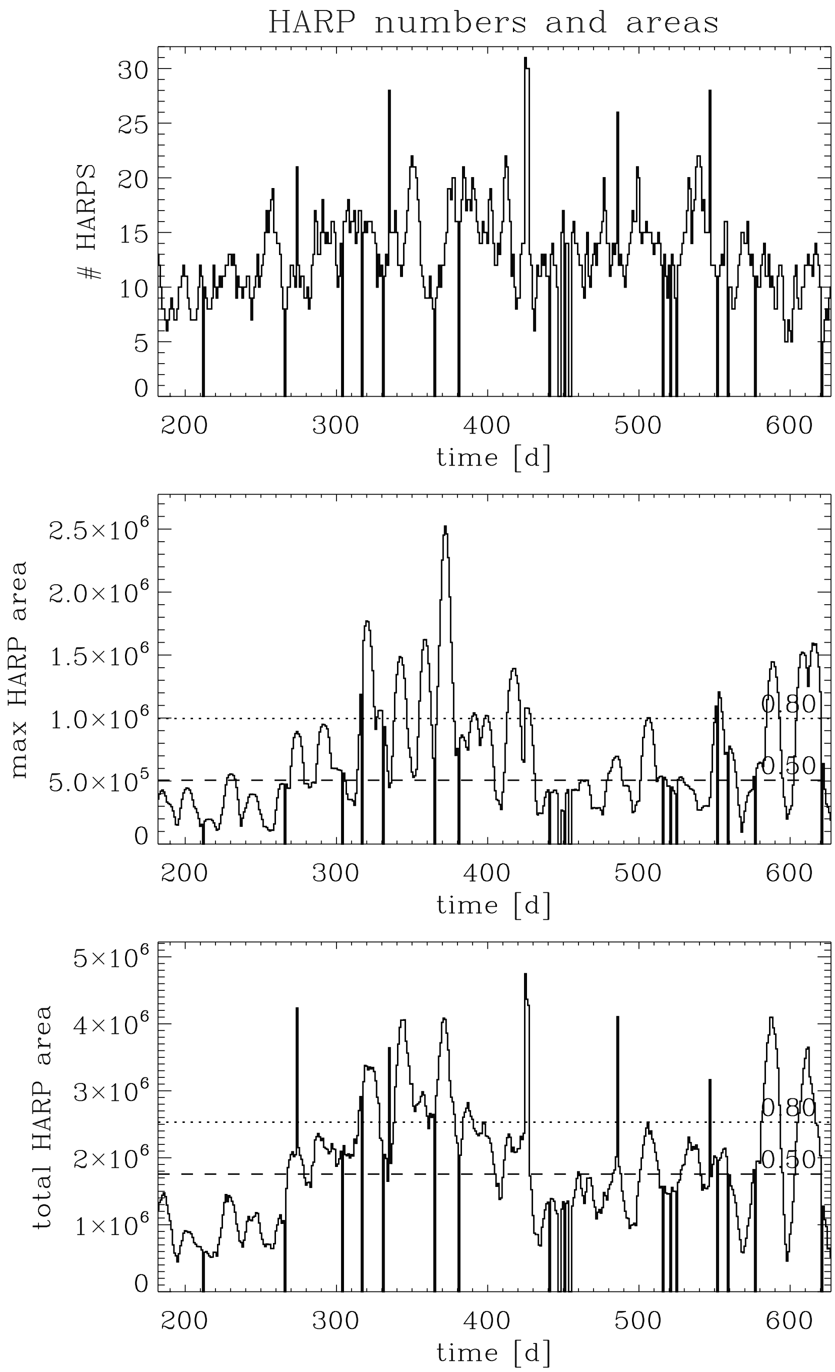

Fig. 11.— Definitive HARP distributions. The middle panel shows the area in pixels of the largest HARP at that time. The bottom panel shows the sum of the areas of all the HARPs at that time. Time of 0 corresponds to 01 January 2011. The dotted horizontal lines show the 80th percentile; the dashed horizontal lines show the 50th percentile (median). The top panel shows the number of HARPs, although that is not used in this analysis.

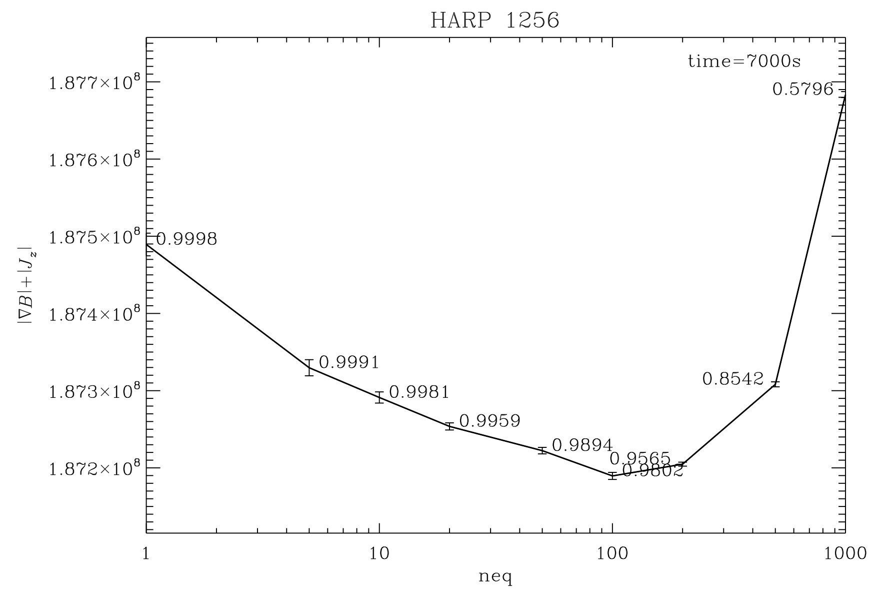

Fig. 12.— Annealing energy for HARP 1256 interpolated to a time of 7000 s, in the same format as Fig. 3.

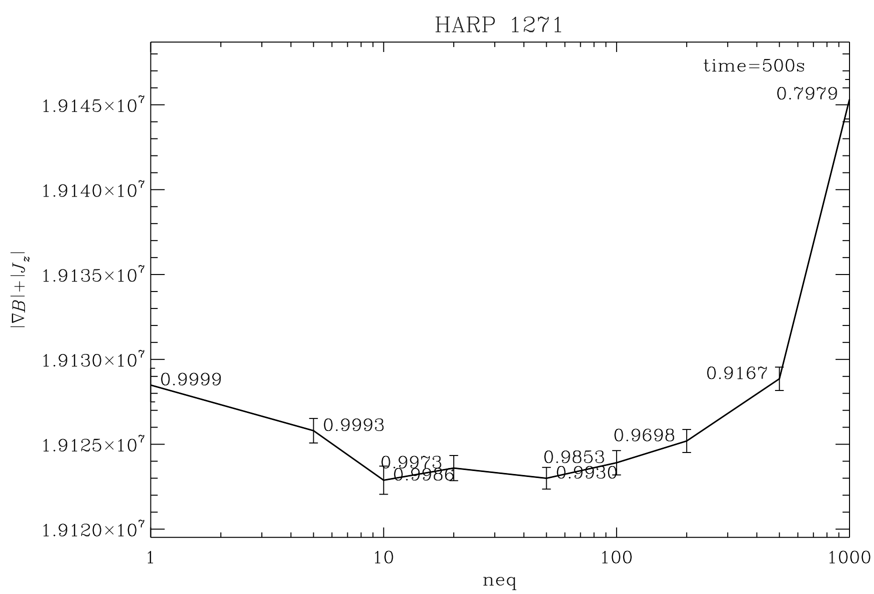

Fig. 13.— Annealing energy for HARP 1271 interpolated to a time of 500 s, in the same format as Fig. 3.

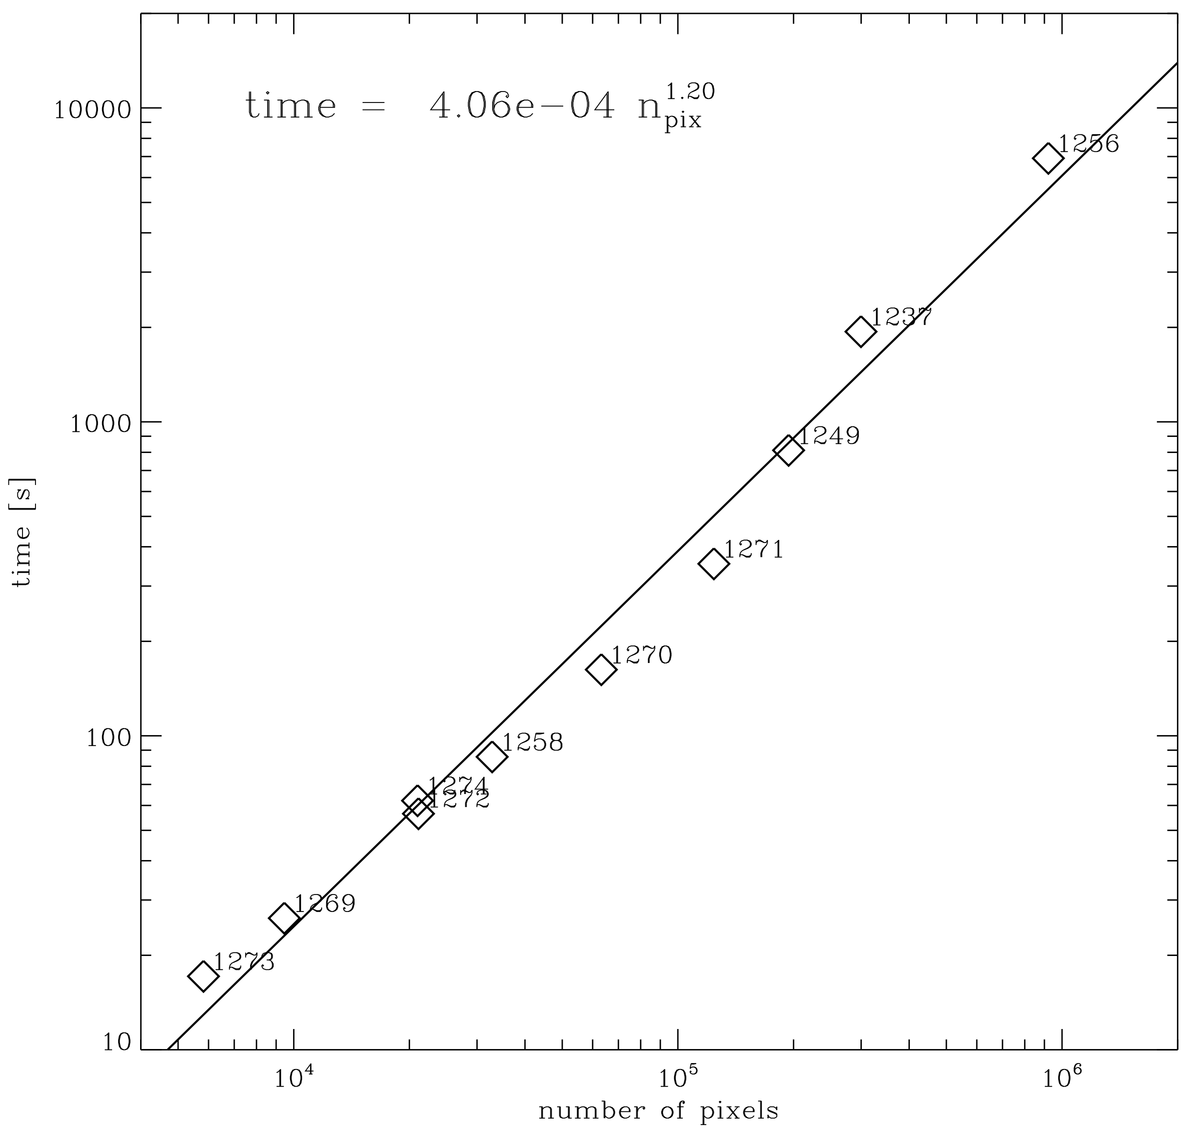

Fig. 14.— Scaling of run time with area of HARP in pixels for neq = 100, tfactr = 0.98. There is some indication that this deviates from a power law.

needed to achieve the desired run times, so consider a slightly longer run time of 8000 s for HARP 1256, as shown in Figure 16. In this case, the minimum in energy still occurs for $neq \approx 100$, but with a slightly larger $tfactr \approx 0.982$ than for a time of 7000 s. Thus, the value of neq for definitive HARPs can likely be set at 100, and small adjustments to the run time made by modifying tfactr only.

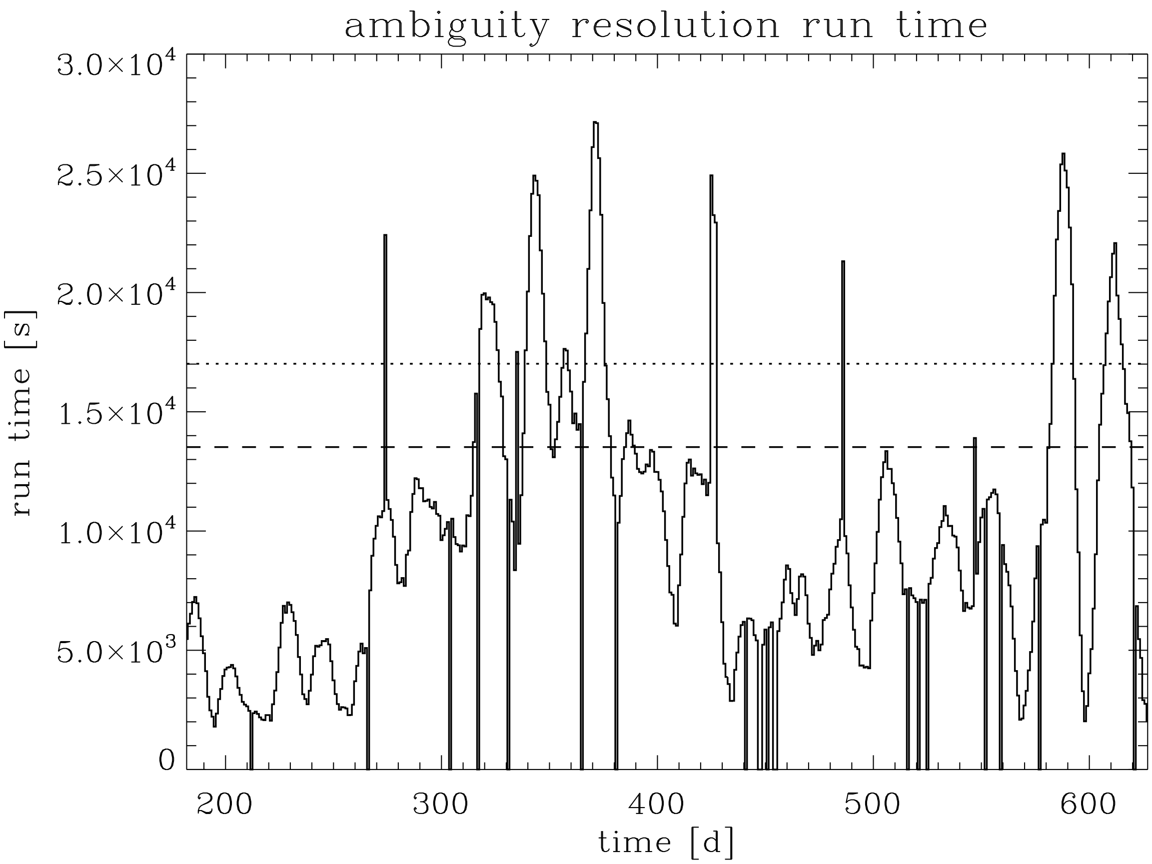

Fig. 15.— Distribution of run time for power law scaling of neq = 100, tfactr = 0.98. The dashed line shows the 80th percentile, while the dotted line shows the 88th percentile, which is at a run time of approximately 17000 s. Thus, this choice of annealing parameters results in shorter run time than targeted for the resources available.

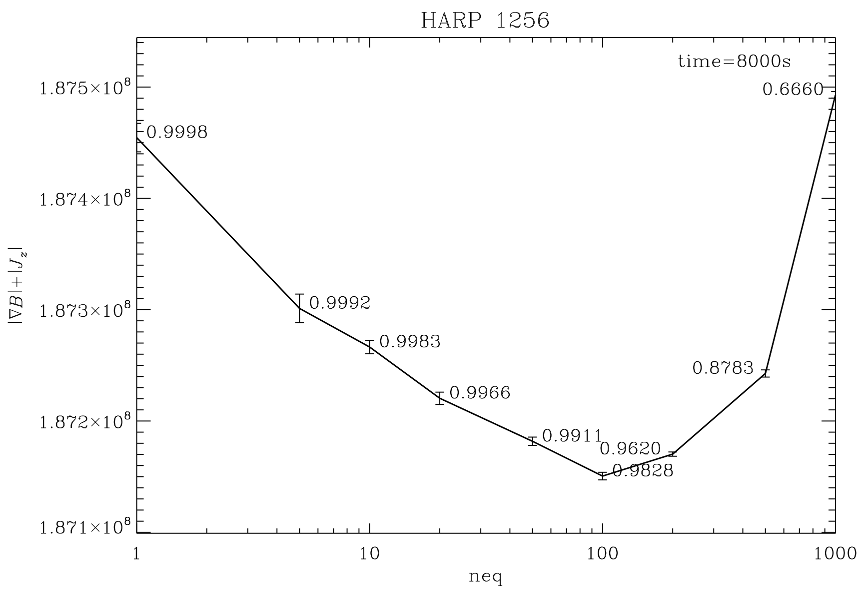

Fig. 16.— Annealing energy for HARP 1256 interpolated to a time of 8000 s, in the same format as Fig. 3. The minimum in energy still occurs for $neq \approx 100$, but with a slightly larger $tfactr \approx 0.982$ than for a time of 7000 s.

4. Setting the Threshold

In addition to the annealing parameters, also need to determine the transverse field strength above which the result of the annealing can be trusted, and below which the smoothing will be applied. Figure 9 suggests two ways to determine this.

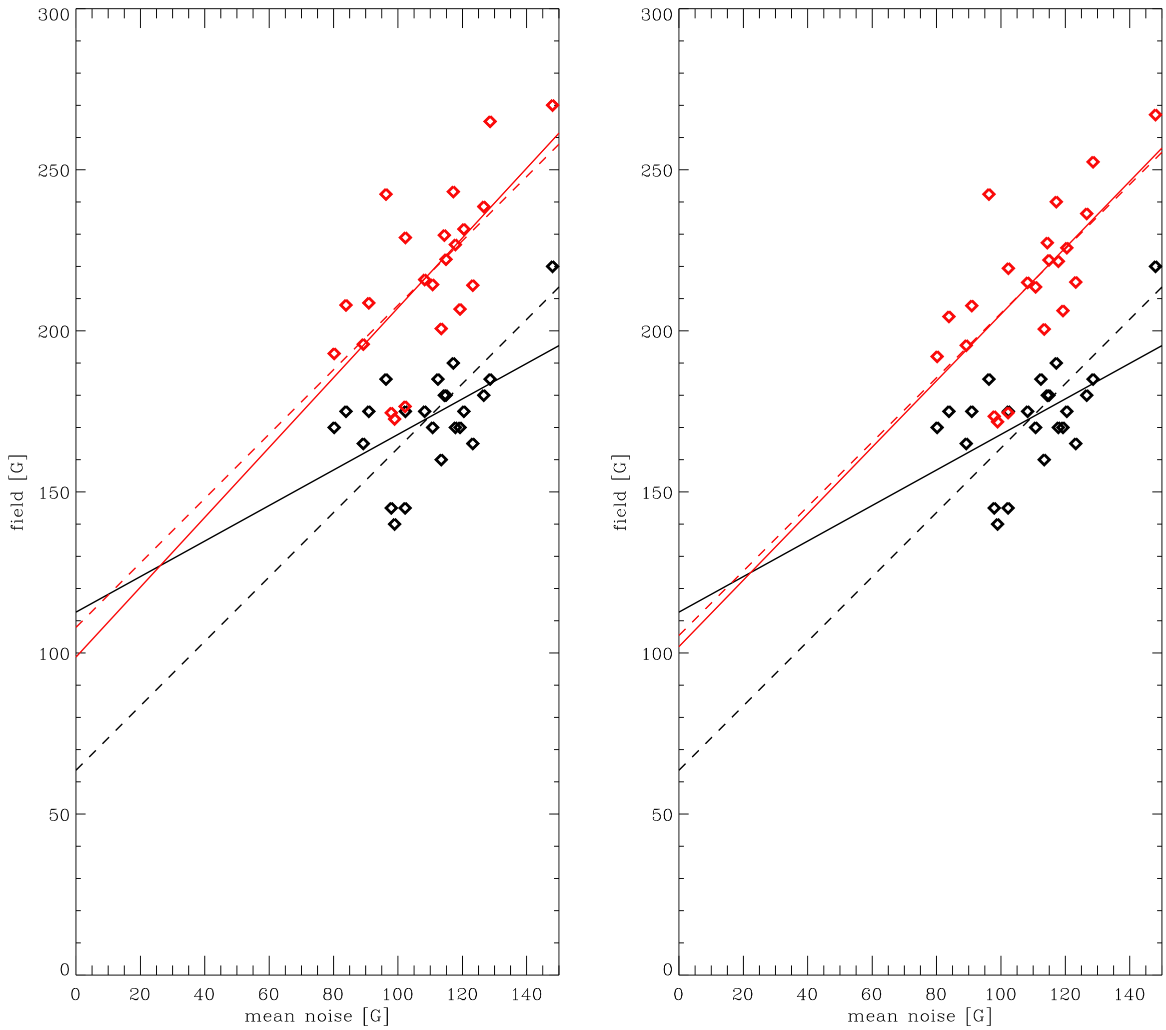

Fig. 17.— Threshold field strength as a function of the mean value of the noise mask for HARPs on 2011.12.05 at 00:00:00_TAI and 2011.12.31 at 19:00:00_TAI for NRT annealing parameters (left) and definitive annealing parameters (right). Black symbols are the value of the transverse field at which the fraction of "wrong" pixels is decreasing most rapidly (c.f. Fig. 9). Red are the value of the transverse field at which 10% of the pixels are "wrong". Solid lines are the regression fit; dashed lines are the best fit lines with a slope of 1.

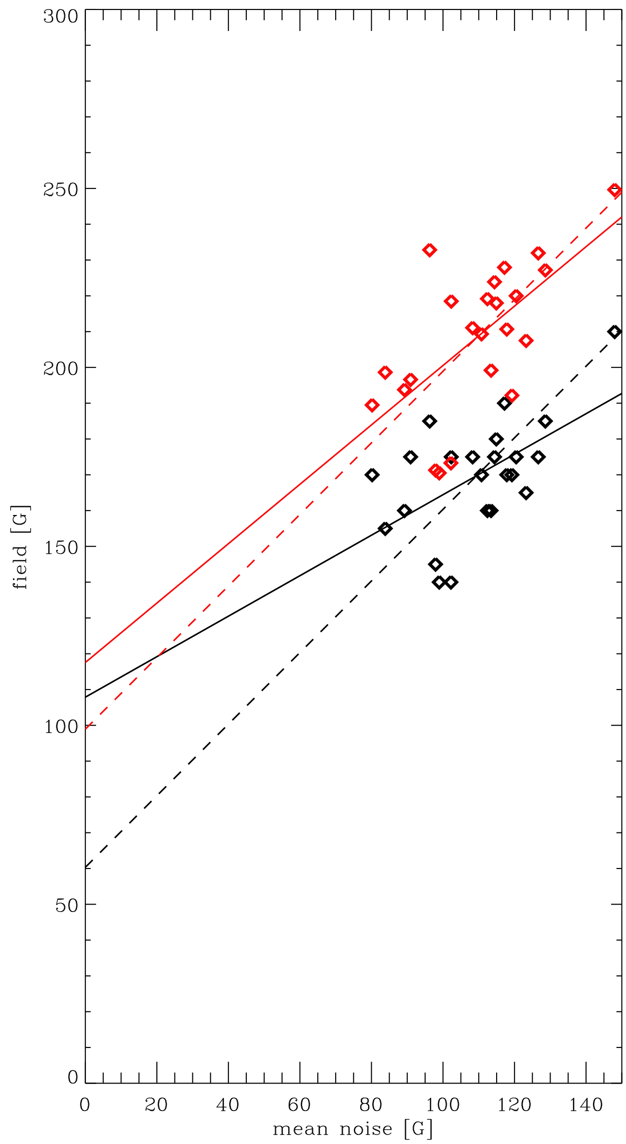

Fig. 18.— Same as Fig. 17, right but with the fraction wrong calculated excluding a buffer of 10 pixels.

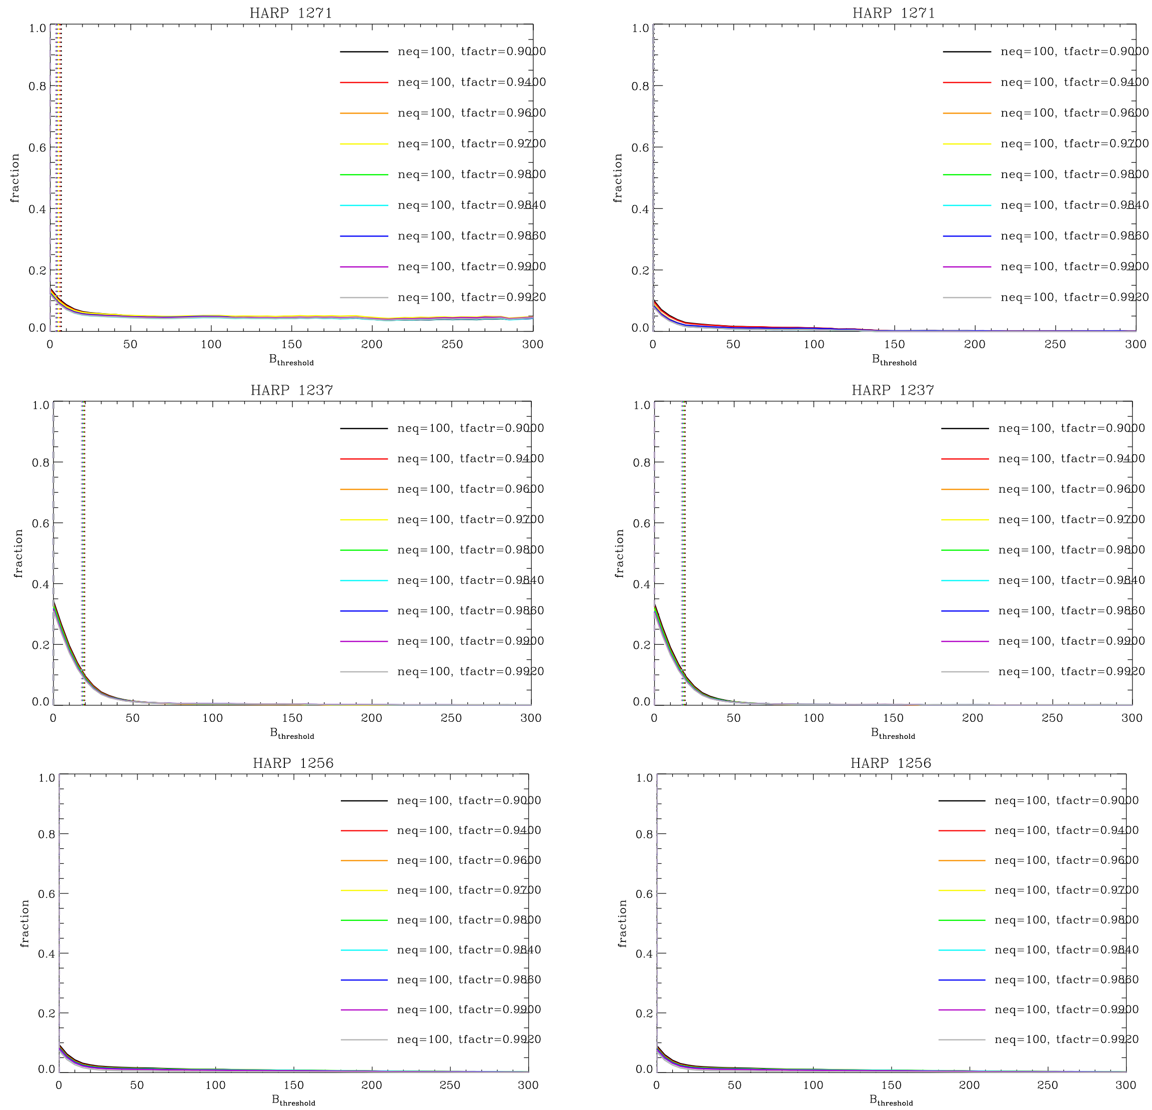

Fig. 19.— The fraction of pixels above a field strength threshold for which no consistent answer is found, as a function of the threshold above the noise mask, eroded by one pixel, for HARPs 1271, 1237, and 1256 (HARP area increases from top to bottom), for a range of annealing parameters. The vertical dotted lines mark the field strength above which 10% of the pixels have no consistent answer. Left: all pixels are included; right: a 10 pixel buffer is removed.

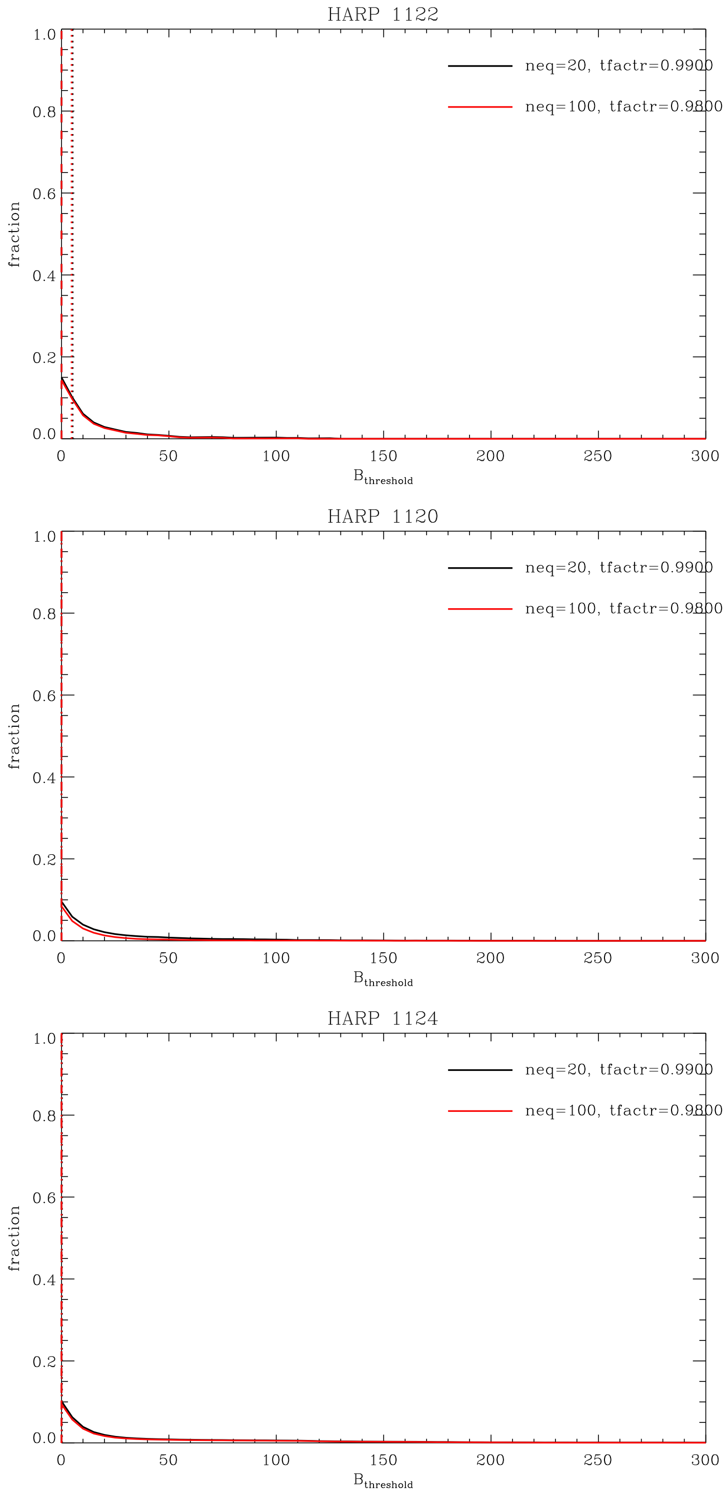

Fig. 20.— The fraction of pixels above a field strength threshold for which no consistent answer is found, as a function of the threshold, for HARPs 1122, 1120, and 1124 (HARP area increases from top to bottom), for a range of annealing parameters. The vertical dotted lines mark the field strength above which 10% of the pixels have no consistent answer.