

SAGECal Performance With Large Sky Models

Hanno Spreeuw, Sarod Yatawatta, Ben van Werkhoven, and Faruk Diblen

Abstract – As astronomical instruments become more sensitive, the requirements for calibration software become more stringent; without accurate calibration solutions, thermal noise levels in images will not be reached, and the scientific output of the instrument is degraded. Calibration requires bright sources with known properties, in particular with respect to their brightnesses as a function of frequency. However, for modern radio telescopes with a huge field of view, a single calibration source does not suffice; instead, a sky model with tens of thousands of sources is needed. In this work, we investigate the compute load for such complicated sky models, with up to 50,000 sources, for the SAGECal calibration package. We have chosen half of the sources in these models to be point sources and half of them extended, which we represent by Gaussian profiles.

1. Introduction

From its beginning, in the 1930s, radio astronomy progressed almost half a century with the use of simple calibration techniques, that is, by calibrating on very bright unresolved sources with known flux density. These “standard candles” were observed just before and/or after the observation of the target source. The gain solutions from the individual dishes as calculated from the calibrator observation were applied to the target observation before images were made. This obviously cannot account for any gain phase changes during the target observation, and it will be hard if not impossible to reach thermal noise levels in the images in this way. During the 1980s, a new calibration technique, known as self-calibration or hybrid mapping, found its way to aperture synthesis arrays [1], although it had been applied to very-long-baseline interferometry data earlier [2]. During a self-calibration loop, a sky model is produced from the observation itself; this is used for the next iteration to find more accurate gain solutions until these solutions converge and the root-mean-square noise in the image no longer decreases. It should be obvious that more complete representations of the radio sky, as probed by the interferometer, will lead to more accurate calibration solutions. The most complete representations, however, will consist of tens of thousands of extended sources and point sources. These large numbers of sources will lead to substantial

compute loads, first, when these complex sky models are Fourier transformed to the (u, v, w) spacings of the observation in order to find the Jones matrices that provide the optimal mapping from these “ground truth” visibilities to observed visibilities and, second, to compute the residuals after the Jones matrices have been derived. For this purpose, we will investigate the runtimes of the SAGECal calibration package [3, 4] for sky models of five different sizes to find a single set of calibration solutions; we have not run any self-calibration loops since the sky model is considered perfect. We also vary the number of stations in the observations that we calibrate. In this way, we derive the scaling of SAGECal performance for increasingly sensitive radio telescopes from the Low Frequency Array (LOFAR) [5] to the Square Kilometre Array (SKA) [6].

2. Method

We used the same set of five artificial observations as we did in earlier work [7]. These are mock observations with 64, 128, 256, 384, and 512 stations, more or less ranging from the number of stations in the International LOFAR Telescope [8, 9] to SKA1-low [10]. We ran SAGECal on a dual Intel Xeon E5-2660v3 (40 logical cores), with both an NVIDIA Titan X and a GeForce 1080 GPU. SAGECal was set to use 40 CPU threads and both of the GPUs for sufficient workloads without any preferred order for deploying the GPUs. In this work, we have set the number of clusters, that is, the number of calibration directions, to 1 for all runs. By keeping this dimension fixed, it becomes easier to focus on the effect of the complexity of the sky model on SAGECal runtimes. We chose half of the sources in each sky model to have Gaussian profiles and the other half to be point sources. The peak brightnesses, sky positions, and spectral indices are randomly distributed between some limits (likewise for the major axes and orientation of the Gaussians, with the size of the minor axis always equal to half the size of the major axis). As in previous work [7], we depict the times for finding the actual calibration solutions by running SAGECal with both the “-a 0” and the “-a 1” settings. Running with the latter setting will Fourier transform the sky model to visibility space, while running with the former also includes the time to derive calibration solutions. Hence, the assessment of SAGECal performance in this work is built on the assumption that the Fourier transform of the sky model can be kept in RAM, which may not be possible if the observation has many frequency channels. As in our earlier work, we simultaneously investigate the effect of an increasing number of antenna stations on the time to compute

Manuscript received 24 August 2020.

Hanno Spreeuw, Ben van Werkhoven, and Faruk Diblen are with Netherlands eScience Center, Science Park 140, Amsterdam, Netherlands; e-mail: h.spreeuw@esciencecenter.nl, b.vanwerkhoven@esciencecenter.nl, f.diblen@esciencecenter.nl.

Sarod Yatawatta is with ASTRON, Netherlands Institute for Radio Astronomy, Dwingeloo, Netherlands; e-mail: yatawatta@astron.nl.

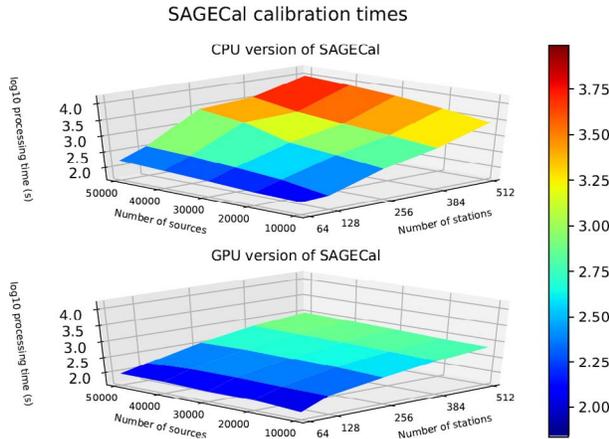


Figure 1. We measured the times SAGECal needs to calibrate five artificial data sets, having 64, 128, 256, 384, and 512 stations for five different sky models with varying numbers of sources but all in a single cluster (calibration direction). We have used a computer equipped with two CPUs (dual Xeon E5-2660v3, 40 logical cores, with number of CPU threads = 40 in all runs), equipped with an NVIDIA GeForce GTX 1080 and a Titan X GPU on ASTRON's DASS; see [13].

calibration solutions. Although SAGECal can process a single observation on multiple nodes, we do not consider that here. It was, however, discussed elsewhere [11, 12] in the context of consensus optimization.

3. Results

Figure 1 shows the results of our SAGECal runs for radio telescopes and sky models of five different sizes. It is clear from this figure that the GPU version of SAGECal scales better for increasing workloads. While the performance is similar for the lowest workload—99s for the GPU version of SAGECal versus 69s for its

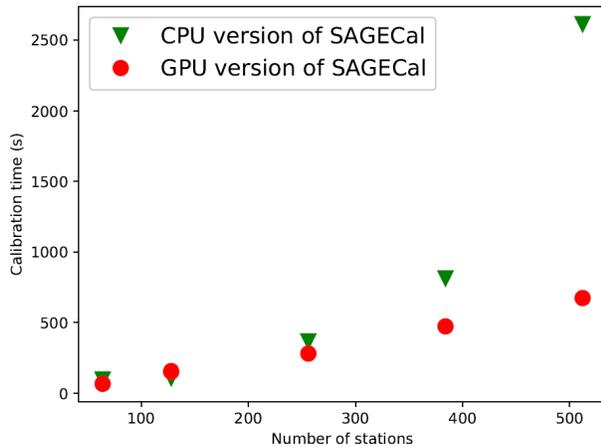


Figure 2. Slice from Figure 1 to show how the SAGECal runtimes are affected—on both CPU and GPU—by the size of the radio telescope—from LOFAR to SKA—when calibrating on our smallest sky model.

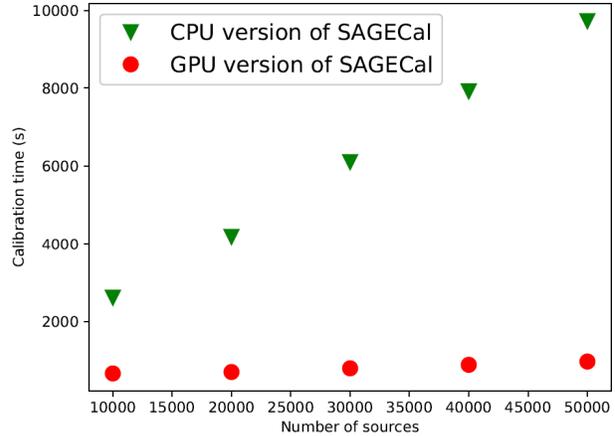


Figure 3. Slice from Figure 1 to show how the complexity of sky models affects the time SAGECal needs to derive calibration solutions for its CPU and GPU version.

CPU version—SAGECal on a GPU is 10 times faster than on a CPU for the heaviest workload, as we expected.

Figure 2 shows SAGECal scalability—from LOFAR to SKA—for our smallest sky model: $1e4$ sources. While the runtimes of the CPU version of SAGECal are comparable to the GPU version for 64 and 128 stations, it shows scaling issues for observations from even larger telescopes. The GPU version, on the other hand, shows close-to-linear scaling over the entire range of observations.

Figure 3 is another slice of the 3D plot from Figure 1 but in the orthogonal dimension across all sky models for 512 stations. It shows linear scaling for both devices, but the slope is much shallower for the GPU than for the CPU, resulting in a factor-10 difference in processing time for a sky model with $5e4$ sources.

Figure 4 shows only the GPU values from Figure 3 to get a more accurate understanding of this SAGECal runtimes on a typical SKA1-low observation using large

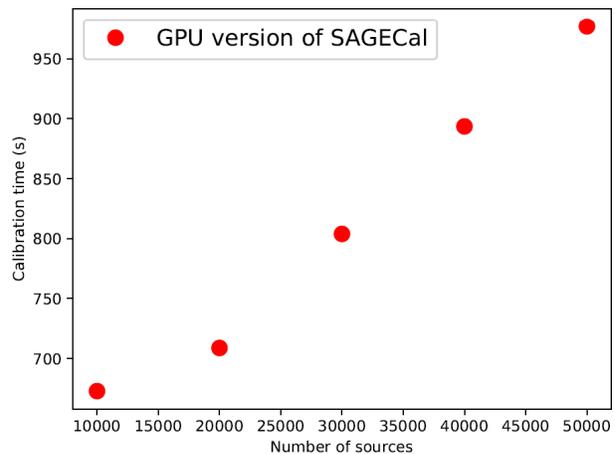


Figure 4. A zoom-in on the GPU values from Figure 3 for quantitative insight into the effect of larger sky models on SAGECal runtimes.

sky models. It is striking that, when comparing a relatively small sky model with 5000 extended sources and 5000 point sources to a sky model five times larger, the time to compute calibration solutions increases by only 45%. This is a promising result with regard to accurate calibration of large observations, such as the observations that are expected to be performed with SKA1-low.

4. Conclusions

In order to foresee possible compute bottlenecks in accurately calibrating data from increasingly large and sensitive radio telescopes, we have investigated SAGECal runtimes using mock observations. As expected from earlier work, we notice insufficient linear scaling from a typical LOFAR observation to a SKA1-low observation even when calibrating on a relatively small sky model—consisting of $5e3$ point sources and $5e3$ extended sources—on a CPU with 40 logical cores. The GPU version does show linear scaling across all telescope sizes. With regard to SAGECal runtimes in the orthogonal dimension—from sky models with $1e4$ to $5e4$ sources for an SKA1-low observation—we see linear scaling on both the CPU and the GPU with a much shallower slope for the latter. This is reassuring considering the most accurate and therefore the most computing-intensive and data-intensive calibration ever that is coming up with the start of the early science observations of the SKA.

5. References

1. A. C. S. Readhead and P. N. Wilkinson, “The Mapping of Compact Radio Sources From VLBI Data,” *The Astrophysical Journal*, **223**, July 1978, pp. 25-36.
2. T. J. Pearson and A. C. S. Readhead, “Image Formation by Self-Calibration in Radio Astronomy,” *Annual Review of Astronomy and Astrophysics*, **22**, January 1984, pp. 97-130.
3. S. Yatawatta, S. Kazemi, and S. Zaroubi, “GPU Accelerated Nonlinear Optimization in Radio Interferometric Calibration,” *Innovative Parallel Computing—Foundations & Applications of GPU, Manycore, and Heterogeneous Systems (INPAR 2012)*, San Jose, CA, USA, May 2012, pp. 1-6, doi:10.1109/InPar.2012.6339610.
4. S. Yatawatta, F. Diblen, H. Spreeuw, B. van Werkhoven, *SAGECal*, version 1.0, <https://github.com/nlesc-dirac/sagecal>, 2017.
5. <https://www.lofar.org>.
6. “SKA: Square Kilometre Array,” <https://www.skatelescope.org>.
7. H. Spreeuw, B. van Werkhoven, S. Yatawatta, and F. Diblen, “GPU Acceleration of the SAGECal Calibration Package for the SKA,” arXiv e-prints, October 2019, arXiv:1910.13908.
8. M. P. van Haarlem, M. W. Wise, A. W. Gunst, G. Heald, J. P. McKean, et al., “LOFAR: The LOW-Frequency ARray,” *Astronomy & Astrophysics*, **556**, August 2013, p. A2.
9. “LOFAR Science 2019,” https://www.astron.nl/lofarscience2019/Documents/Monday/LUM_Pizzo.pdf
10. “The Square Kilometre Array Fact Sheet for Industry,” <https://www.skatelescope.org/wpcontent/uploads/2018/08/16231-factsheet-telescopesv71.pdf>.
11. S. Yatawatta, F. Diblen, H. Spreeuw, and L. V. E. Koopmans, “Data Multiplexing in Radio Interferometric Calibration,” *Monthly Notices of the Royal Astronomical Society*, **475**, 1, December 2017, pp. 708-715, <https://doi.org/10.1093/mnras/stx3130>.
12. S. Yatawatta, “Distributed Radio Interferometric Calibration,” *Monthly Notices of the Royal Astronomical Society*, **449**, 4, April 2015, pp. 4506-4514, 04 2015, <https://doi.org/10.1093/mnras/stv596>.
13. “Distributed ASCI Supercomputer,” <https://www.cs.vu.nl/das5/ASTRON.shtml>.