GPU acceleration of the SAGECal calibration package for the SKA

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

ADASS, Groningen, 10/09/2019

An image from uncalibrated visibilities

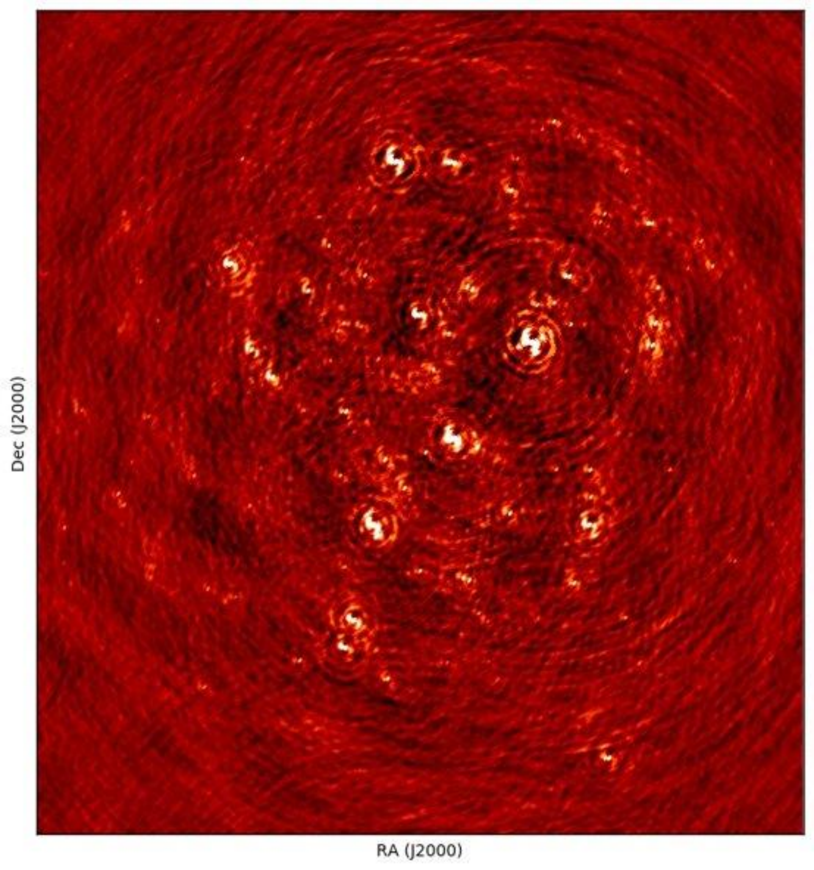

An image from calibrated visibilities

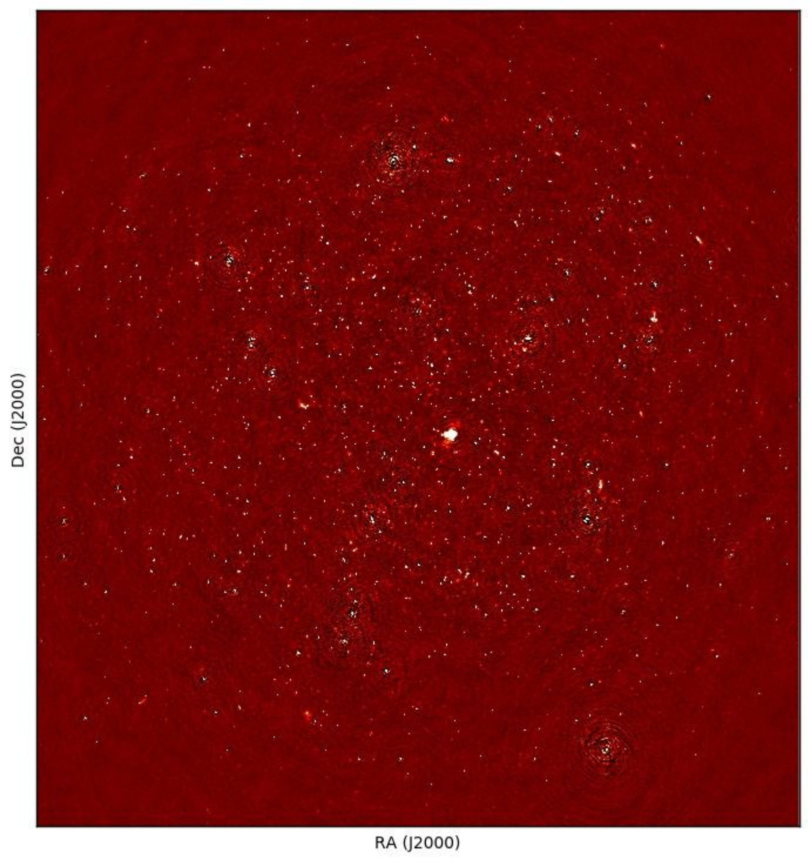

The International LOFAR Telescope (ILT)

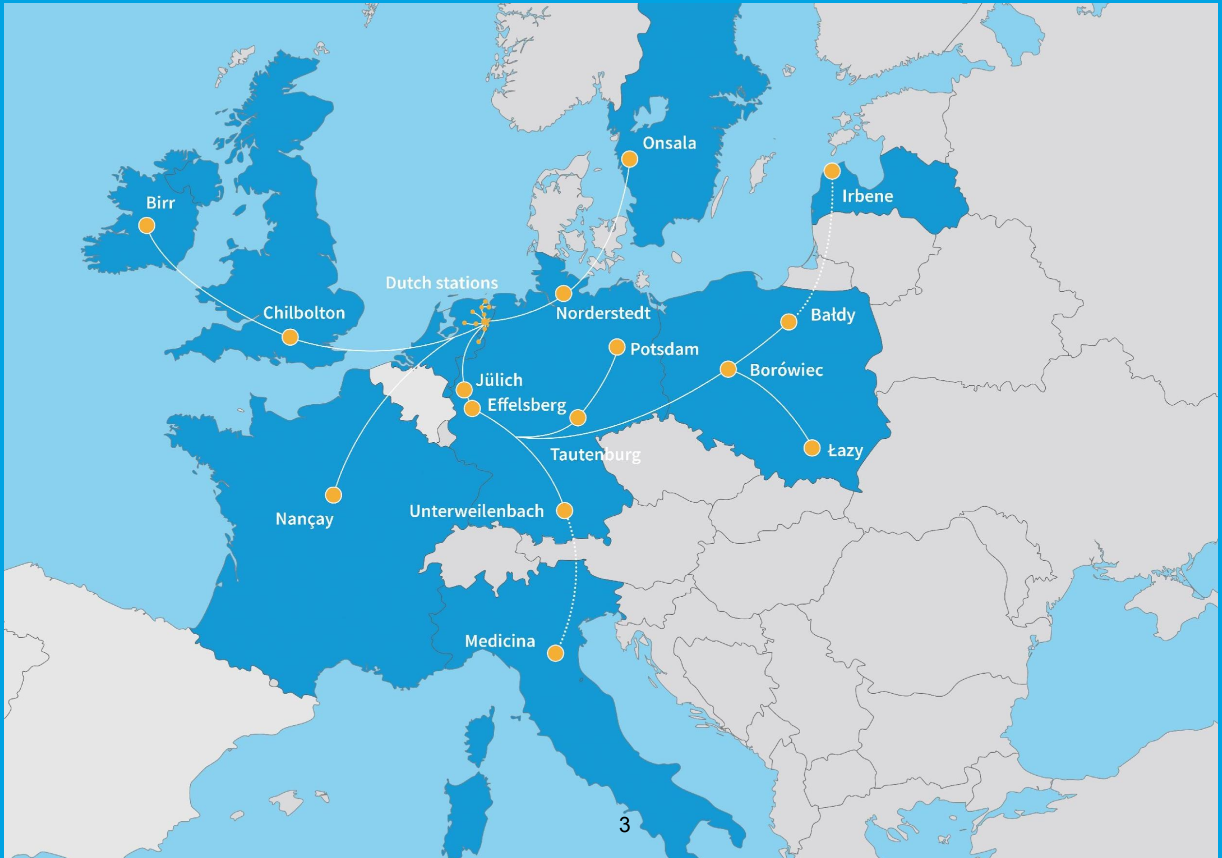

The Square Kilometre Array (SKA)

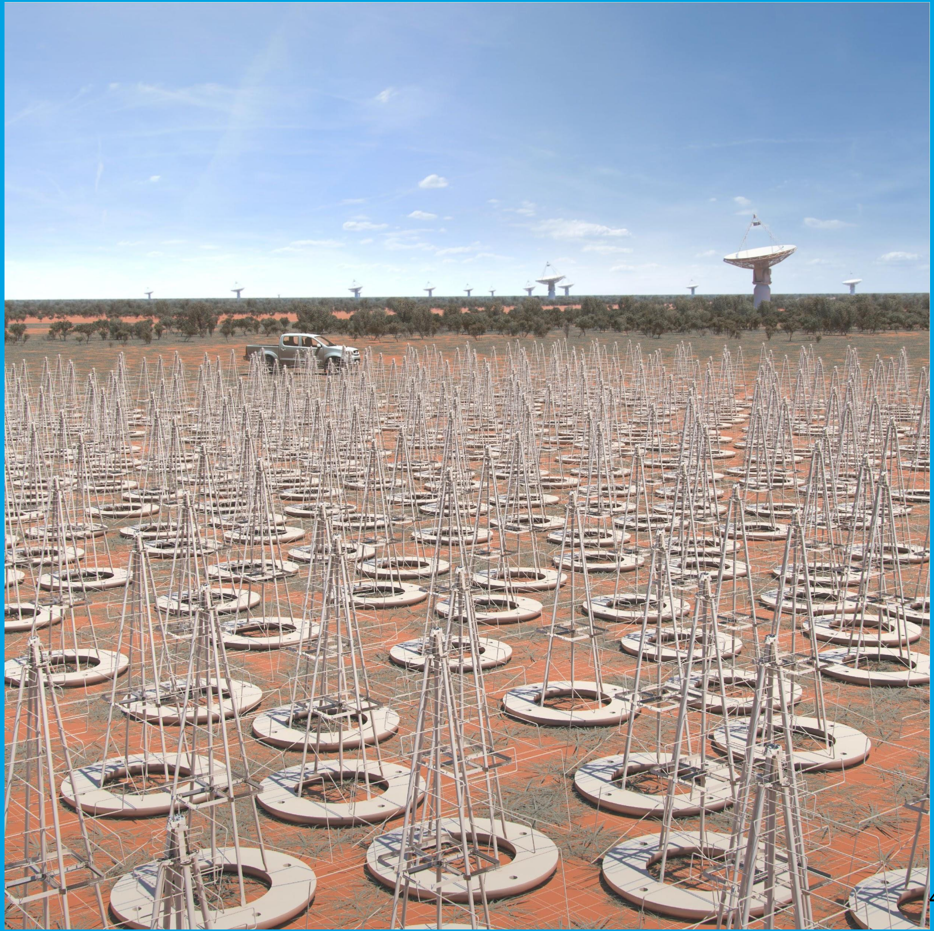

TECHNICAL INFORMATION

THE TELESCOPES

The Square Kilometre Array (SKA) is made up of arrays of antennas - SKA-mid observing mid to high frequencies and SKA-low observing low frequencies - to be spread over long distances. The SKA is to be constructed in two phases: Phase 1 (called SKA1) in South Africa and Australia; with Phase 2 (called SKA2) representing a significant increase in capabilities and expanding into other African countries, with the component in Australia also being expanded.

SKA1-mid

the SKA's mid-frequency instrument

Location: South Africa

Frequency range:
350 MHz to **15.3 GHz**
with a goal of 24 GHz

197 dishes
(including 64 MeerKAT dishes)

Maximum baseline:
150km

SKA1-low

the SKA's low-frequency instrument

Location: Australia

Frequency range:
50 MHz to **350 MHz**

~131,000
antennas spread between
512 stations

Maximum baseline:
~65km

Solving a scaling problem

- Number of SKA stations >> Number of LOFAR stations
- Data size ∞ (number of stations)^2
- Number of FLOPs should not ∞ data size
- Need large bandwidth and compute power

The SAGECal calibration package

- Space-Alternating Generalized Expectation-Maximization Calibration
- Taking the accuracy of calibration to the next level, e.g. for studying the HI signal from the Epoch of Reionization
- Over a decade of development (Yatawatta, Kazemi, Saroubi et al.). Recently by Yatawatta, Diblen, Van Werkhoven, Spreeuw
- Much faster than the usual Levenberg-Marquardt algorithm

The SAGECal calibration package (ctd.)

- Supports all source models: points, Gaussians, disks, rings, (widefield) shapelets (prolate spheroidal wave functions).
- Builds on casacore, but only for I/O with Measurement Sets
- CMake (CPU and GPU)
- Docker (CPU only)
- Singularity (GPU and CPU)

SAGECal on a many-core processor

- Dual Intel Xeon E5-2660 v3 @ 2.60GHz
- 40 logical cores, 20 physical cores
- Maximum power consumption per CPU 105W, so 210 W for the dual CPU
- Profiling with Intel VTune Amplifier

SAGECal threading efficiency

Effective CPU Utilization: 70.4% (28.179 out of 40 logical CPUs)

Effective CPU Utilization Histogram

This histogram displays a percentage of the wall time the specific number of CPUs were running simultaneously. Spin and Overhead time adds to the Idle CPU utilization value.

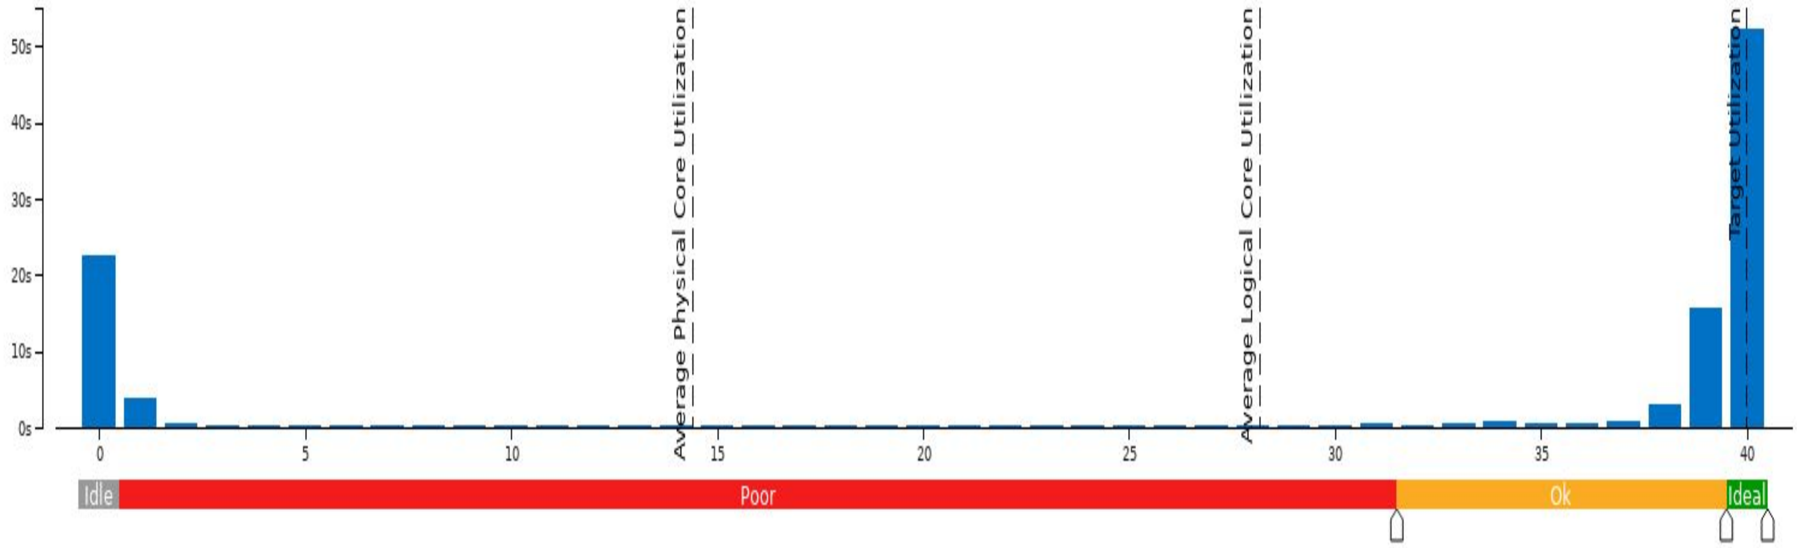

SAGECal on an NVIDIA Titan X (Pascal)

1.1. Kernel Performance Is Bound By Compute

For device "TITAN X (Pascal)" the kernel's memory utilization is significantly lower than its compute utilization. These utilization levels indicate that the performance of the kernel is most likely being limited by computation on the SMs.

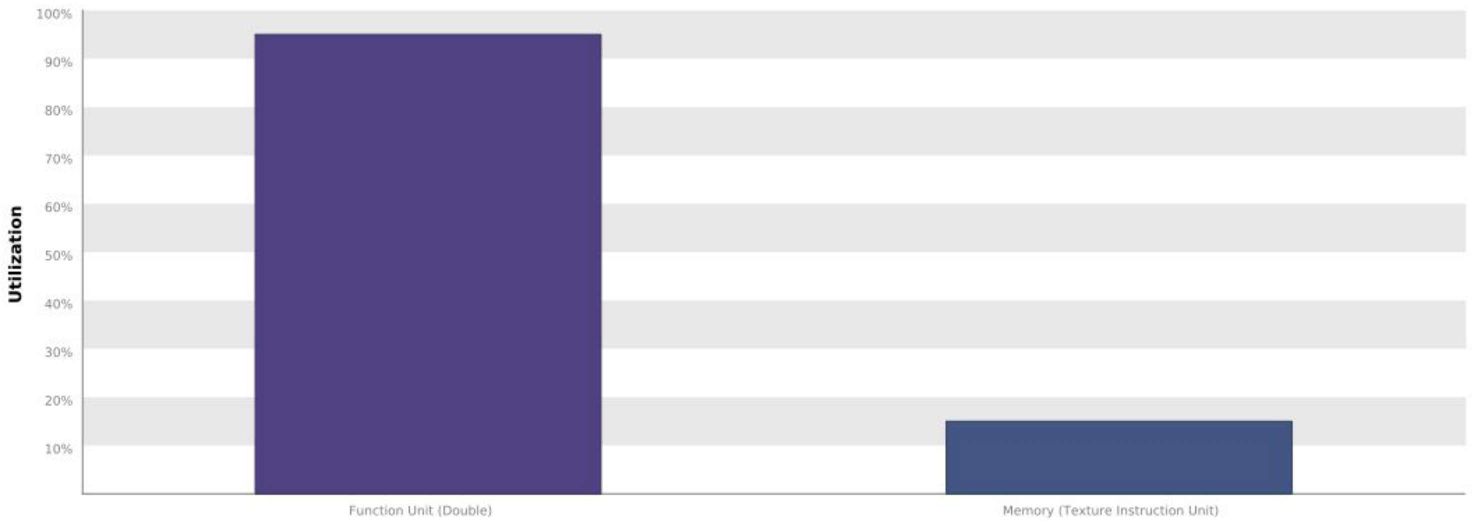

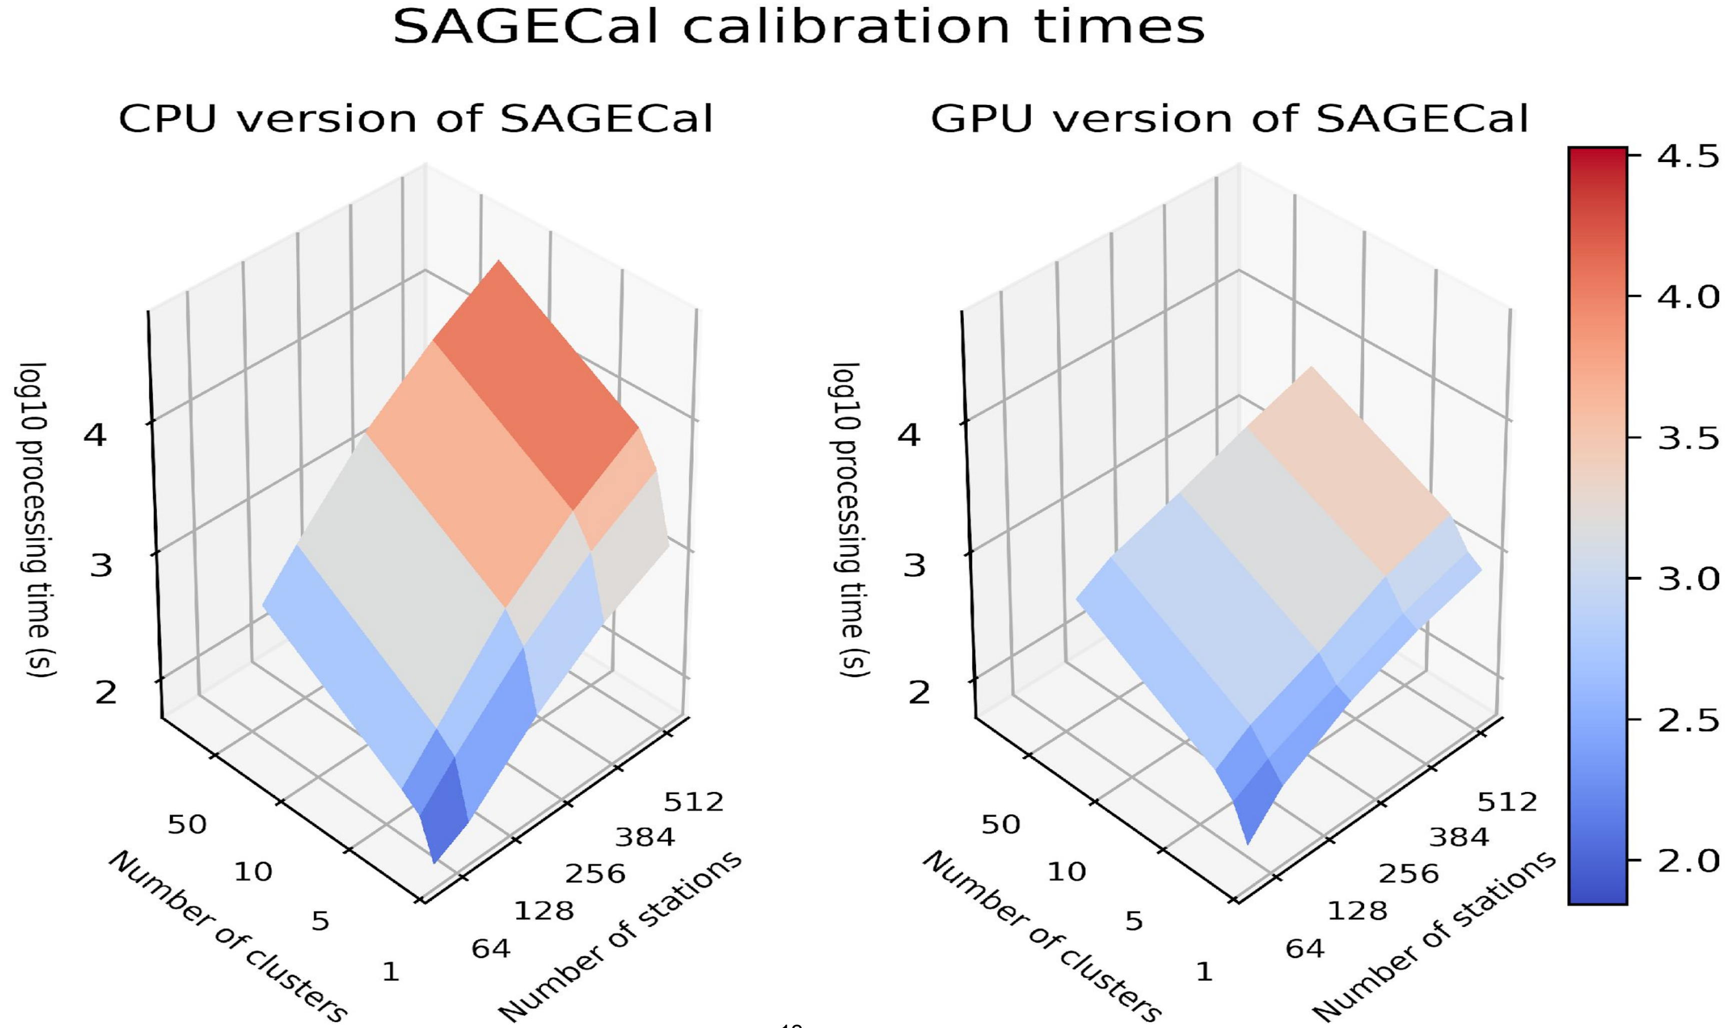

SAGECal calibration times
CPU version of SAGECal
GPU version of SAGECal
log10 processing time (s)
2
3
4
Number of clusters
1
5
10
50
Number of stations
64
128
256
384
512
2.0
2.5
3.0
3.5
4.0
4.5

SAGECal CPU and GPU scalability

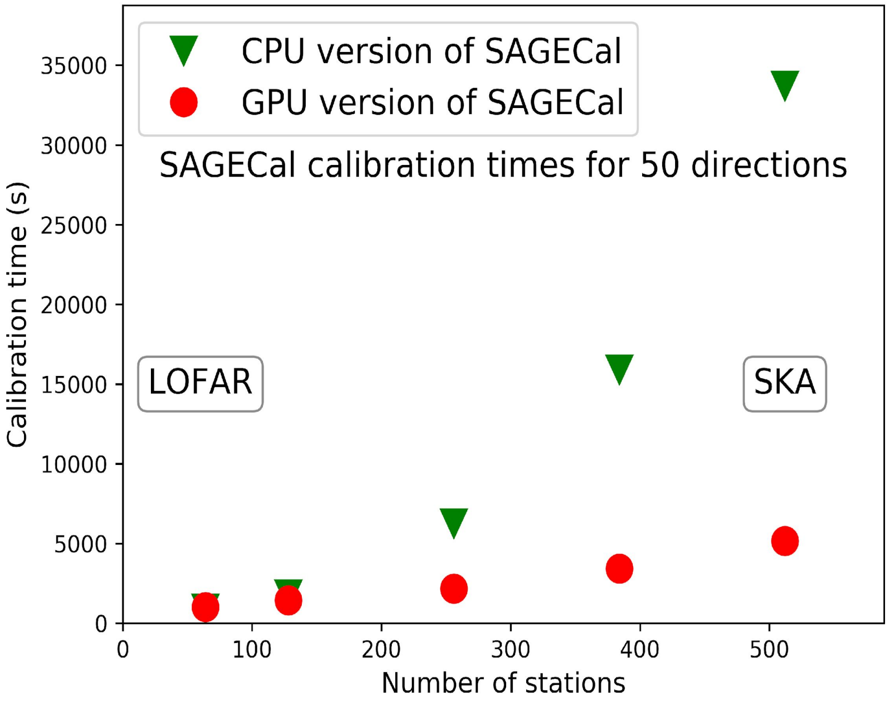

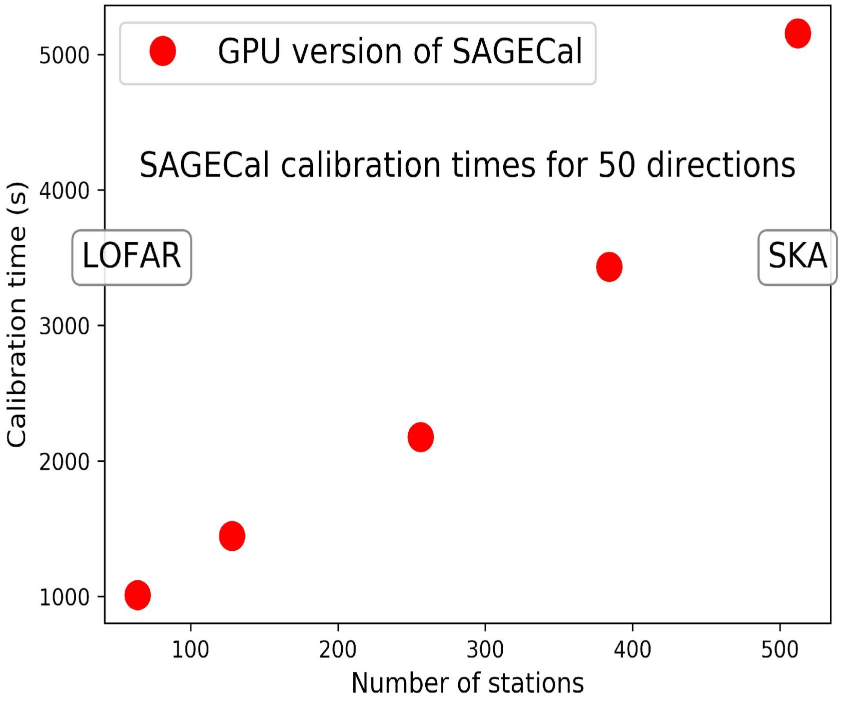

Sagecal for calibration in remote places, shortage of power

- NVIDIA Jetson Nano
- Power 5-10 W
- GPU: 128-core NVIDIA Maxwell™ architecture-based GPU
- CPU: Quad-core ARM® A57
- Memory: 4 GB 64-bit LPDDR4; 25.6 gigabytes/second

SAGECal on an NVIDIA Jetson Nano demo

Use SAGECal to calibrate this observation

- Short observation of 3C196
- 125 10s time samples, 61 stations.
- 153 MHz, just one 183 kHz channel
- Calibrate on SkyView image of 3C196 from TGSS (GMRT)

Compare SAGECal runtimes for a larger observation and sky model : many-core CPU vs. low power GPU

- Close to peak compute performance exploited on both devices
- SAGECal is highly optimized for both devices
- SAGECal runs 48% faster on a dual Intel Xeon E5-2660 v3 @ 2.60GHz than on a NVIDIA Jetson Nano
- But the Intel Xeon uses about 20 times more power

Conclusions

- SAGECal is highly optimized for many-core CPUs and for both many-core and low power GPUs.
- SAGECal performance scales better on a GPU than on a CPU
- SAGECal runtimes on a low power GPU are comparable to a many-core CPU.

Any questions?

- https://www.esciencecenter.nl/project/dirac by Sarod Yatawatta (PI)
- https://github.com/nlesc-dirac/sagecal
- https://github.com/nlesc-dirac/pytorch
- **Performance analysis of distributed radio interferometric calibration, IEEE SAM 2018, DOI 10.1109/SAM.2018.8448481**
- h.spreeuw@esciencecenter.nl

Let's stay in touch

+31 (0)20 460 4770

info@esciencecenter.nl

www.esciencecenter.nl

blog.esciencecenter.nl

eScienceCenter

eScienceCenter

linkd.in/1j2uS8S

netherlands eScience center