URSI GASS 2020, Rome, Italy, 29 August - 5 September 2020

Radio-Astronomical Imaging with WSClean and Image-Domain Gridding

Veenboer B., Van der Tol S., Offringa A. R., Romein J. W., and Dijkema T. J.
ASTRON (Netherlands Institute for Radio Astronomy)

Abstract

Modern radio observatories provide astronomers with large data sets. To turn these data sets into science, astronomers require imaging techniques that provide state-of-the-art image quality (e.g. resolution and dynamic range) at high performance (e.g. modest runtime). We integrated an efficient GPU-accelerated implementation of the Image-Domain Gridding (IDG) algorithm into the widely used WSClean imager to fulfill these needs. WSClean features several novel deconvolution techniques, such as auto-masked multi-scale, and parallel cleaning, while IDG is capable of correcting full-polarization direction-dependent effects (DDEs) during imaging. In this paper we give an overview of the current state of our work: we introduce IDG, we present performance of IDG on contemporary accelerator hardware, we show the impact of DDE correction on image quality and we discuss its suitability for the demanding EoR science case as an example. All in all, WSClean+IDG is an excellent imager for current and future radio-telescopes.

1 Introduction

The LOw-Frequency ARray (LOFAR) [6] radio telescope is currently the most sensitive low-frequency radio telescope and comprises tens of receivers (each consisting of many antennas) centered around The Netherlands with remote receivers all over Europe, with a maximum baseline length of about 2000 km. The future Square-Kilometre Array (SKA) will be even larger (more receivers) and will be built as two large arrays, one in South Africa (SKA1-Low) and one in Western Australia (SKA1-Mid).

While the design and construction of such large instruments are very challenging in itself, another challenge awaits when the instrument is put into use. A radio telescope such as the SKA produces enormous amounts of data (at a rate multiple times the global internet traffic) and it is impossible to store all this data. Efficient processing techniques are therefore required to process all measurements produced by the instrument in (near) real-time and create data products for use by astronomers around the world.

However, even for existing telescopes such as LOFAR, data processing remains challenging – in particular, when so-called *direction-dependent effects* (DDEs) that vary over the field of view have to be taken into account [7]. Examples of DDEs are the ionosphere and the primary beam of the instrument. Only when the DDEs are corrected for, the high dynamic ranges made possible by the high sensitivity and large fields of view of the new generation of radio interferometers can be achieved [8].

Motivated by the above challenges, we used hardware/software co-design to come up with the novel Image-Domain Gridding (IDG) algorithm that corrects for the DDEs, and runs highly efficient on accelerator hardware. IDG is now integrated into the widely used WSClean imaging application. This combination of WSClean+IDG is a good candidate for imaging data of current and future radio-telescopes.

We first briefly introduce Image-Domain Gridding in Section 2. Section 3 discusses implementations of this algorithm for CPUs, GPUs and FPGAs. Section 4 showcases the integration of IDG in WSClean and highlights the latest deconvolution features provided by WSClean. In Section 5 we discuss the use of WSClean+IDG for Epoch of Reinization (EoR) research. We can also use IDG for direction-dependent calibration, as we will show in Section 6.

2 IDG algorithm

The high cost of computing the convolution kernels in W-projection and AW-projection [8] was the main motivation for the development of the *Image-Domain Gridding* algorithm (IDG) [2]. IDG effectively performs the same operation as classical gridding and degridding with AW-projection, except that it is designed to circumvent the computation of convolution kernels altogether.

At the center of the algorithm are so-called *subgrids*, which represent low-resolution versions of the sky brightness for a small subset of visibilities, see Fig. 1. These subgrids are placed onto the grid in such a way that they cover a subset of visibilities, together with their associated W-term and A-term convolution kernels, i.e., a subgrid covers all grid points that are affected by the convolved visibility.

In IDG, gridding consists of four steps: (1) the visibilities are gridded onto subgrids by the *gridder kernel*; (2) A-term correction is applied to each pixel of the subgrid in image-space; (3) subgrids are Fourier transformed by applying an

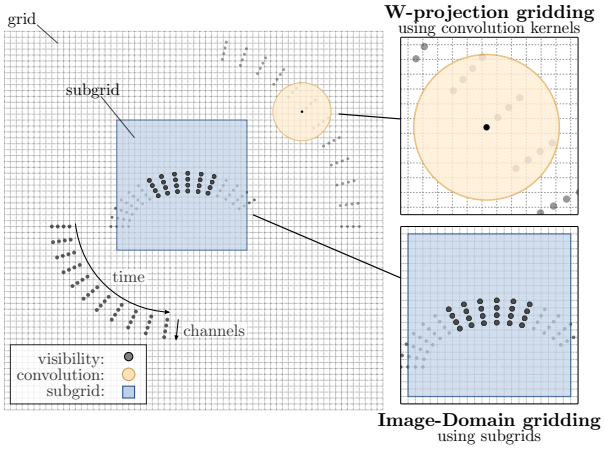

Figure 1. In traditional W-projection and AW-projection gridding, visibilities are gridded using convolutions in the frequency domain (top-right) as opposed to IDG, in which the W-term and A-term effects are corrected in the image domain (bottom right). For the latter, neighboring visibilities (indicated with thick dots) are gridded onto small 'subgrids'.

inverse 2D FFT; (4) the subgrids are added to the grid by the *adder kernel*. Degridding is similar to the gridding, but proceeds in reverse order: (1) subgrids are extracted from the grid by a *splitter kernel*; (2) subgrids are Fourier transformed by an FFT kernel; (3) A-term correction is applied; (4) the *degridder kernel* computes visibilities. A direct summation of visibilities onto the subgrid is computationally feasible as in practice the subgrid size is 4–6 orders of magnitude smaller then the size of the grid.

The number of operations in IDG scales linearly with the number of subgrids and more than quadratically with the size of the subgrid: $\mathcal{O}(N^2 \log N^2)$. Furthermore, depending on the number of visibilities per subgrid, gridding is either dominated by the operations in the gridder kernel or by the Fourier transformation of the subgrid and the adder kernel. The regular structure of the IDG algorithm allows for a highly-efficient implementation, as we will discuss in Section 3. Moreover, in IDG correction for direction-dependent effects take place outside of the critical path, and is therefore relatively cheap to perform.

3 IDG implementation

In [3] we demonstrate how we implemented and optimized IDG for CPUs and GPUs. To this end, we introduced an 'execution plan', which allows for a highly scalable implementation of various compute kernels. Furthermore, we implement the different steps of IDG gridding and degridding using various kernels where we apply optimizations such as vectorization, data reuse, prefetching, memory coalescing, latency hiding and loop transformations to maximize the number of operations performed per clock cycle. On CPUs, the achieved performance is significantly lower than the theoretical peak. This is caused by the evaluation of sine and cosine, which take about 80% of the total kernel runtime [3]. We found that GPUs from AMD and NVIDIA greatly benefit from hardware support for the evaluation of sines and cosines. Especially NVIDIA GPUs, with support for this operation in dedicated hardware, run highly efficient at about 90% of the theoretical peak [3].

Alternatives such as a GPU-accelerated implementation of W-projection perform less, at about 60% of the theoretical peak [12]. Thus while IDG additionally corrects for direction-dependent effects, it is even faster than a traditional W-projection gridder.

During the past years, we saw new FPGA technologies like hard Floating-Point Units, the ability to program them in a high-level language, and tight integration with CPU cores. In [4] we show that we can efficiently optimize for FPGA resource usage, but also that optimizing for a high clock speed is difficult. We compared a FPGA, GPU and CPU with almost identical theoretical peak performance and demonstrated that the FPGA and GPU outperform the CPU by about an order of magnitude, both in terms of throughput (the number of visibilities processed per second) as well as in terms of energy efficiency (see also Figure 2). In absolute terms, the GPU is the fastest and most energy-efficient device for IDG.

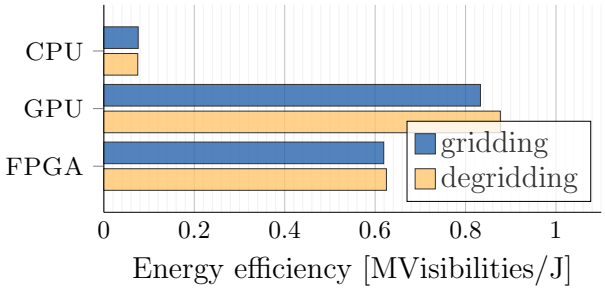

Figure 2. Image-Domain Gridding runs much more energy-efficient on accelerators (an Intel Arria 10 FPGA and an NVIDIA Maxwell GPU) than on a CPU (Intel Haswell). Both accelerators are much more energy-efficient then the CPU by processing about an order of magnitude more visibilities for every Joule consumed [4].

4 IDG in WSClean

WSClean is a widely used imager [9] that uses the w-stacking algorithm to implement inversion (gridding) and predict (degridding) and provides several novel deconvolution algorithms relevant for the SKA, including fast auto-masked multi-scale, and parallel cleaning. This makes WSClean more robust to calibration errors and an order of magnitude faster than alternative imagers.

We integrated IDG into WSClean such that features as data handling, deconvolution, etc. are maintained, while the existing inversion (gridding) and predict (degridding) functionality are provided by IDG. Together, WSClean and IDG

provide a unique mix of state-of-the-art imaging and deconvolution algorithms. This combination is now used in production by a variety of different radio observatories, such as LOFAR [6] and the MWA [10].

Figure 3 illustrates that the direction-dependent TEC and gain corrections provided by WSClean+IDG result in superior image quality compared to an image where these corrections are not applied. We are currently working towards creating an imaging pipeline (called 'Rapthor') for LOFAR data processing capable of direction-dependent calibration and imaging.

The source code for WSClean can be downloaded at: `https://sourceforge.net/projects/wsclean`. IDG is available at: `https://git.astron.nl/RD/idg`.

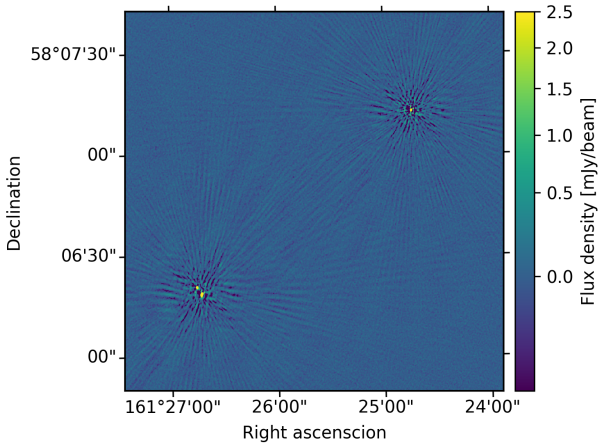

(a) WSClean with IDG beam corrections

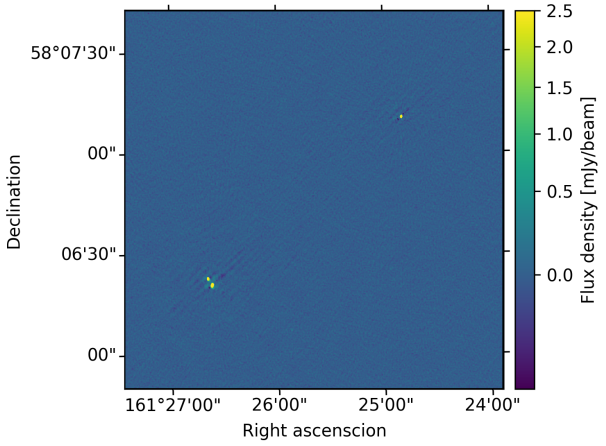

(b) WSClean with IDG beam, TEC and gain corrections

Figure 3. These images by Sweijen et. al. [13] are an extract of the 'Lockman Hole' field (at 46 Mhz, 0.4" resolution) and illustrate the benefits of using an imager capable of correcting for direction-dependent effects.

5 IDG for EoR

Experiments that try to observe the 21-cm redshifted signals from the epoch of reionisation (EoR) using interferometric low-frequency instruments have stringent requirements on the processing accuracy. In 21-cm EoR power spectrum experiments, foregrounds are distinguished from the 21-cm signals by their spectral smoothness. Therefore, the spectral accuracy is a particularly important aspect of a gridder. Offringa et al. [5] demonstrate that traditional algorithms, with standard settings, are not accurate enough for 21-cm signal extraction. Of the various methods, IDG shows the highest accuracy with the lowest imaging time. Therefore, Image-Domain Gridding is overall the most suitable algorithm for 21-cm EoR power spectrum experiments, including for future analyses of data from the Square Kilometre Array (SKA).

6 IDG for direction-dependent calibration

Many DD-calibration algorithms split the sky in facets and solves for a gain in each facet, e.g. [14]. Van der Tol et. al. [11] propose a new method in which a smooth screen is directly fitted to the data. This is advantageous in combination with gridders that apply smooth screens, as it mitigates an extra step in which a screen is fitted to the solutions after solving for facet-bases gains, which can introduce additional errors. Moreover, because a spatial smoothness constraint is enforced within the solver, a higher S/N ratio of the solutions can be achieved. Van der Tol demonstrates that such an algorithm can be efficiently implemented using a degridder, and continues to show that a highly efficiently implementation of this algorithm can be implemented using IDG as the degridder.

7 Conclusion

We created an imager (WSClean+IDG) that combines state of the art in deconvolution techniques with an efficient GPU-accelerated imager that corrects for polarized direction-dependent effects during imaging. This imager offers unprecedented imaging capabilities for current and future radio telescopes.

8 Acknowledgements

This work is funded by the Netherlands eScience Center (NLeSC), under grant no 027.016.G07 (Triple-A 2), the EU Horizon 2020 research and innovation programme under grant no 754304 (DEEP-EST) and by NWO (DAS-5). The authors would like to thank Frits Sweijen and Reinout vavn Weeren for commissioning WSClean+IDG and for providing the sky-images in this paper.

References

[1] The SKA Organisation, "Square Kilometre Array", https://www.skatelescope.org/

[2] Van der Tol, S. and Veenboer, B. and Offringa, André R, "Image Domain Gridding", A27, Astronomy & Astrophysics, 2018

[3] Veenboer, B, Petschow, M and Romein, J. W., Image-Domain Gridding on Graphics Processors, IPDPS, pp. 545–554, 2017

[4] Veenboer B. and Romein, J. W, Radio-Astronomical Imaging: FPGAs vs GPUs, Euro-Par, pp. 509–521, 2019

[5] Offringa, A. R., and others, Precision requirements for interferometric gridding in 21-cm power spectrum analysis, A12, Astronomy & Astrophysics, 2019

[6] van Haarlem, M. P. and others, LOFAR: The LOw-Frequency ARray, A2, Astronomy & Astrophysics, 2013

[7] van Weeren, R. J. and others, LOFAR facet calibration, The Astrophysical Journal Supplement Series, 2016

[8] Tasse, C. and others, Applying full polarization A-Projection to very wide field of view instruments: An imager for LOFAR, A105, Astronomy & Astrophysics, 2013

[9] WSClean: an implementation of a fast, generic wide-field imager for radio astronomy, Offringa, A. R., and others, Mon. Not. R. Astron. Soc, pp. 606–619, 2014

[10] Tingay, S. J. and others, The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies, Publ. Astron. Soc. Aust., e007, 2013

[11] Van der Tol, S. and others, Estimating continuous direction-dependent gain screens from radio interferometric visibilities and a large skymodel, ADASS, 2019

[12] Merry, B Faster GPU-based convolutional gridding via thread coarsening, Astronomy & Computing, pp. 140–145, 2016

[13] Sweijen, F., van Weeren, R.J., Morabito, L.K. and others, in prep.

[14] Kazemi, S. and others, Radio interferometric calibration using the SAGE algorithm, Mon. Not. R. Astron. Soc, pp. 1656–1666, 2011