University of Groningen

cuFFS

Sridhar, S. S.; Heald, G.; van der Hulst, J. M.

Published in:
Astronomy and Computing

DOI:
[10.1016/j.ascom.2018.10.002](https://doi.org/10.1016/j.ascom.2018.10.002)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2018

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Sridhar, S. S., Heald, G., & van der Hulst, J. M. (2018). cuFFS: A GPU-accelerated code for Fast Faraday rotation measure Synthesis. *Astronomy and Computing*, 25, 205-212.
<https://doi.org/10.1016/j.ascom.2018.10.002>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Full length article

cuFFS: A GPU-accelerated code for Fast Faraday rotation measure Synthesis

S.S. Sridhar^{a,b,c,*}, G. Heald^c, J.M. van der Hulst^a^a Kapteyn Astronomical Institute, University of Groningen, Postbus 800, 9700AV Groningen, The Netherlands^b ASTRON, The Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands^c CSIRO Astronomy and Space Science, 26 Dick Perry Ave, Kensington, WA 6151, Australia

ARTICLE INFO

Article history:

Received 18 April 2018

Accepted 6 October 2018

Available online 23 October 2018

Keywords:

GPGPU

Methods: data analysis

Techniques: image processing

Techniques: polarimetric

Computing methodologies: graphics processors

ABSTRACT

Rotation measure (RM) synthesis is a widely used polarization processing algorithm for reconstructing polarized structures along the line of sight. Performing RM synthesis on large datasets produced by telescopes like LOFAR can be computationally intensive as the computational cost is proportional to the product of the number of input frequency channels, the number of output Faraday depth values to be evaluated and the number of lines of sight present in the data cube. The required computational cost is likely to get worse due to the planned large area sky surveys with telescopes like the Low Frequency Array (LOFAR), the Murchison Widefield Array (MWA), and eventually the Square Kilometre Array (SKA). The massively parallel General Purpose Graphical Processing Units (GPGPUs) can be used to execute some of the computationally intensive astronomical image processing algorithms including RM synthesis. In this paper, we present a GPU-accelerated code, called cuFFS or CUDA-accelerated Fast Faraday Synthesis, to perform Faraday rotation measure synthesis. Compared to a fast single-threaded and vectorized CPU implementation, depending on the structure and format of the data cubes, our code achieves an increase in speed of up to two orders of magnitude. During testing, we noticed that the disk I/O when using the Flexible Image Transport System (FITS) data format is a major bottleneck and to reduce the time spent on disk I/O, our code supports the faster HDFITS format in addition to the standard FITS format. The code is written in C with GPU-acceleration achieved using Nvidia's CUDA parallel computing platform. The code is available at <https://github.com/sarrvesh/cuFFS>.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Until the early 2000s, improvements in the execution speeds of scientific programs were largely due to the increase in CPU clock rates. This however changed in last decade when chip manufacturers realized that it was hard to keep increasing the clock rate as it led to overheating of the chips due to thermal losses (see for example Patterson and Hennessy, 2014).

Processor manufacturers work around the “thermal wall” by implementing multiple processing units (or cores) on a single chip leading to the design of multicore processors. While each core on the chip is not faster than its predecessor, multicore processors achieve higher compute throughput using parallelization. GPUs extend this strategy of implementing multiple cores on a single chip to an extreme limit. In comparison to multicore CPUs, GPUs achieve high compute throughput by maximizing the number of processing units (or cores) on a chip. For example, NVIDIA's Tesla

P100 card has 56 streaming microprocessors each with 64 processing units resulting in a total of 3584 single precision (or 1792 double precision) floating point compute cores on a single chip. An added advantage of using GPUs over CPU compute clusters is that GPUs consume less energy measured in Watts per FLOP¹ (Trancoso and Charalambous, 2005; Michalakes and Vachharajani, 2008).

While GPUs were originally designed to be used as graphics coprocessors² to handle higher resolution and display rates demanded by the gaming industry, the massive compute capability of GPUs has been exploited leading to the creation of a new field called GPGPU (General Purpose GPU) computing (Owens et al., 2007). While the initial scientific programs that were run on GPUs had to be ported to the GPU native programming languages, software solutions made available in the last decade have made it far easier to develop GPU codes (Nvidia, 2017). A review of GPU use in scientific computing can be found in Owens et al. (2007, 2008).

* Correspondence to: ASTRON, Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991PD, Dwingeloo, The Netherlands.

E-mail address: sarrvesh@astron.nl (S.S. Sridhar).

¹ Floating point operation.

² Coprocessors are secondary processors mounted on a device that complements the compute capabilities of the primary processor.

Graphical processing units have been used quite extensively in astronomy in areas of research ranging from real-time applications like adaptive optics (for example, see [Bouchez et al., 2012](#)) and correlators for radio interferometry ([Schaaf and Overeem, 2004](#)), easily parallelizable algorithms like N-body simulations ([Portegies Zwart et al., 2007](#); [Elsen et al., 2007](#)), to more complex algorithms in general relativistic magnetohydrodynamics ([Zink, 2011](#)). For a detailed discussion on the use of GPUs in astronomy, see [Fluke et al. \(2011\)](#) and [Fluke \(2012\)](#).

Within the field of radio astronomy, GPUs have largely been restricted to implementing fast real-time signal processing applications like correlators and real-time calibration pipelines (for example, see [Schaaf and Overeem, 2004](#); [Harris et al., 2008](#); [Ord et al., 2009](#); [Chennamangalam et al., 2014](#); [Price et al., 2016](#)), and pulsar de-dispersion pipelines ([Magro et al., 2011](#)). Recently, the Dynamic Radio Astronomy of Galactic Neutron stars and Extragalactic Transients (DRAGNET; [Bassa et al., 2017](#))³ project with the LOFAR radio telescope has started using GPUs to search for fast radio transients in real-time. GPUs have also been used to accelerate solving the Radio Interferometer Measurement Equation ([Hamaker et al., 1996](#)) to calibrate radio interferometry data ([Perkins et al., 2015](#)). Within the context of medical imaging, it has been demonstrated that the gridding procedure that needs to be carried out before applying two dimensional inverse Fourier transform to the data can be accelerated using GPUs ([Schomberg and Timmer, 1995](#); [Schwartz et al., 2006](#)). Furthermore, a GPU implementation of the fast Fourier transform library cuFFT is available as part of the official CUDA toolkit.

In this paper, we present a GPU-accelerated program – CUDA-accelerated Fast Faraday Synthesis (cuFFS) – to perform a commonly used technique in radio polarimetry called Faraday rotation measure (RM) synthesis. This paper is organized as follows: The theoretical background and computational complexity of the rotation measure synthesis method are presented in Section 2. Our GPU implementation of the RM Synthesis code is explained in Section 3 and benchmark tests are presented in Section 3.3. Finally, we present our conclusion and future prospects for improvements to the software package in Section 4.

2. Background

2.1. RM synthesis in theory

Linearly polarized synchrotron emission is an important observational probe to study astrophysical magnetic fields. Synchrotron emission is produced by cosmic ray electrons accelerating across magnetic field lines. The polarization angle of the synchrotron radiation depends on the orientation of the magnetic field in the plane of the sky. Thus resolved polarimetric observations allow us to probe and study the structure of astrophysical magnetic fields.⁴ However, in practice this is complicated owing to the fact that the observed polarization angle and polarization fraction⁵ are different from the intrinsic properties at the source due to propagation effects caused by Faraday rotation.

Faraday rotation is a phenomenon by which the plane of polarization of a linearly polarized electromagnetic wave gets rotated due to circular birefringence as it propagates through a magnetized medium like the interstellar medium or the ionosphere. The

observed polarization angle (χ) is related to the intrinsic polarization angle (χ_0) through a quantity called Faraday rotation measure (RM) as

$$\chi(\lambda^2) = \chi_0 + RM\lambda^2 \quad (1)$$

where RM is related to the strength of the magnetic field projected along the line of sight ($B_{\parallel} = \vec{B} \cdot \vec{l}$) as

$$RM = 0.81 \int_{\text{source}}^{\text{observer}} n_e \vec{B} \cdot d\vec{l}. \quad (2)$$

In the above equation, n_e is the number density of electrons measured in units of cm^{-3} , l is the pathlength in parsec, B is in μG , and ϕ is in rad m^{-2} .

The observed complex polarization vector $P(\lambda^2) = Q(\lambda^2) + iU(\lambda^2)$ is related to the observed polarization angle $\chi(\lambda^2)$, fractional polarization p , and total intensity I as

$$P(\lambda^2) = pIe^{2i\chi(\lambda^2)}. \quad (3)$$

Note that Eqs. (1) and (3) allow us to relate the observed polarization vector with the intrinsic polarization angle and the Faraday rotation measure. Replacing the Faraday rotation measure with a generalized Faraday depth ϕ , and since emission from multiple ϕ can contribute to the observed polarization vector $P(\lambda^2)$, following [Burn \(1966\)](#) we can write

$$P(\lambda^2) = \int_{-\infty}^{+\infty} pIe^{2i[\chi_0 + \phi\lambda^2]} d\phi = \int_{-\infty}^{+\infty} F(\phi)e^{2i\phi\lambda^2} d\phi. \quad (4)$$

where $F(\phi)$ is defined as the Faraday Dispersion Function $F(\phi) := pIe^{2i\chi_0}$. Inverting this Fourier transform-like equation,

$$F(\phi) = \int_{-\infty}^{+\infty} P(\lambda^2)e^{-2i\phi\lambda^2} d\lambda^2 \quad (5)$$

we can see that the Faraday Dispersion Function (FDF) represents the emitted polarized flux as a function of Faraday depth. Faraday depth is related to the strength of the magnetic field lines projected along the line of sight through the relation

$$\phi(r) = \int_r^{\text{observer}} n_e \vec{B} \cdot d\vec{l} \quad (6)$$

and is equal to RM only for the special case when there is a single Faraday rotating medium with no internal Faraday rotation between the source and the observer ([Vallee, 1980](#); [Brentjens and de Bruyn, 2005](#)). From Eqs. (4) and (5), it is easy to see that $P(\lambda^2)$ and $F(\phi)$ form a Fourier-like transform pair.

Recovering $F(\phi)$ from the observed $P(\lambda^2)$ is however not straight-forward as we cannot measure $P(\lambda^2)$ for $\lambda^2 < 0$. Furthermore, $P(\lambda^2)$ is also not measured for all possible values of $\lambda^2 > 0$. Rotation measure synthesis is a technique that was formulated by [Brentjens and de Bruyn \(2005\)](#) which tries to recover $F(\phi)$ using a discrete sampling of $P(\lambda^2)$. In addition to reconstructing $F(\phi)$, the procedure also improves the sensitivity of the polarization measurement by adding up derotated polarization signal across the entire observed bandwidth. In the following section, we provide a brief overview of the procedure. For a detailed account on RM synthesis, we refer the reader to [Brentjens and de Bruyn \(2005\)](#) and [Heald \(2009\)](#).

The rotation measure synthesis technique works around the problem of incomplete λ^2 coverage by introducing a window function $W(\lambda^2)$ which is non-zero only for λ^2 values at which complex polarization $P(\lambda^2)$ has been measured. Thus, Eq. (4) can be modified as

$$\tilde{P}(\lambda^2) = W(\lambda^2)P(\lambda^2) = W(\lambda^2) \int_{-\infty}^{+\infty} F(\phi)e^{2i\phi\lambda^2} d\phi \quad (7)$$

³ www.astron.nl/dragnet/.

⁴ Polarized synchrotron emission probes only the ordered component of the magnetic field. Distinction between different components of astrophysical magnetic fields and the observational probes used to study them is beyond the scope of this paper. For a recent review on this topic, see [Fletcher and Klein \(2015\)](#) and [Beck \(2016\)](#).

⁵ Polarization fraction is defined as polarized intensity divided by total intensity of the synchrotron radiation.

Combining the above equation with Eq. (5) and applying the convolution theorem, the reconstructed Faraday dispersion function $\tilde{F}(\phi)$ is related to the “true” Faraday dispersion function as

$$\tilde{F}(\phi) = F(\phi) * R(\phi) = K \int_{-\infty}^{+\infty} \tilde{P}(\lambda^2) e^{-2i\phi\lambda^2} d\lambda^2 \quad (8)$$

where

$$K = \left(\int_{-\infty}^{+\infty} W(\lambda^2) d\lambda^2 \right)^{-1}. \quad (9)$$

From the above equation, it is easy to see that the recovered Faraday dispersion function ($F(\phi)$) is a convolution of the “true” Faraday dispersion function with smoothing kernel ($R(\phi)$) called the Rotation Measure Spread Function (RMSF). Conceptually, this is analogous to synthesis imaging where $\tilde{F}(\phi)$ is equivalent to the dirty image and $R(\phi)$ is equivalent to the dirty beam. The RMSF is related to the weight function $W(\lambda^2)$ as

$$R(\phi) = K \int_{-\infty}^{+\infty} W(\lambda^2) e^{-2i\phi\lambda^2} d\lambda^2. \quad (10)$$

2.2. Practical implementation of RM synthesis

As discussed in the previous section, the complex polarization vector is not measured for all values of λ^2 . Modern radio telescopes measure $P(\lambda^2)$ in narrow frequency channels spread across a wide bandwidth each with equal channelwidth, expressed in λ^2 space as $\delta\lambda^2$. In such a scenario, the integrals in Eqs. (8)–(10) can be replaced by sums over each frequency channel provided $\phi\delta\lambda^2 \ll 1$:

$$\tilde{F}(\phi) \approx K \sum_{i=1}^N \tilde{P}_i e^{-2i\phi(\lambda_i^2 - \lambda_0^2)} \quad (11)$$

$$R(\phi) \approx K \sum_{i=1}^N w_i e^{-2i\phi(\lambda_i^2 - \lambda_0^2)} \quad (12)$$

$$K = \left(\sum_{i=1}^N w_i \right)^{-1} \quad (13)$$

In the above equations, $\tilde{P}_i$ is the complex polarization vector measured at λ_i and $w_i = W(\lambda_i)$. Fig. 1 shows a RMSF calculated using Eq. (12) for a mock dataset consisting of 1000 frequency channels covering the 1–2 GHz frequency range (L-band). Notice in Eqs. (11) and (12), we have introduced an additional term λ_0^2 in the exponent which is set to the mean of λ_i^2 . Introducing this additional λ_0^2 term is useful as a cosmetic addition as it prevents rapid rotation of $\tilde{Q}(\phi)$ and $\tilde{U}(\phi)$ as can be seen in Fig. 1. This cosmetic addition allows one to make an accurate estimate of P_{peak} and thus make an accurate estimate of the polarization angle χ using the relation

$$\chi = 0.5 \tan^{-1} \left(\frac{U_{\phi\text{peak}}}{Q_{\phi\text{peak}}} \right) e^{-2i\phi\lambda_{\text{peak}}^2}. \quad (14)$$

Note that $Q_{\phi\text{peak}}$ and $U_{\phi\text{peak}}$ are the values of $\tilde{Q}(\phi)$ and $\tilde{U}(\phi)$ at Faraday depth ϕ_{peak} where $|\tilde{F}(\phi)|$ is maximum.

The aim of the RM synthesis software package is to compute Eqs. (11) and (13). Eq. (12) is needed only if one wishes to deconvolve $R(\phi)$ from the reconstructed Faraday dispersion function using a procedure called RM-CLEAN⁶ (Heald, 2009).

⁶ RM-CLEAN is not supported in the current version of cuFFS. We intend to implement RM-CLEAN and other advanced polarization processing techniques in the future.

Since $\tilde{F}(\phi)$ and $\tilde{P}(\lambda^2)$ in Eq. (11) are complex vectors, the equation can be written in terms of its Stokes Q and Stokes U components as

$$\tilde{Q}(\phi) = K \sum_{i=1}^N Q_{\lambda i} \cos 2\phi(\lambda_i^2 - \lambda_0^2) + U_{\lambda i} \sin 2\phi(\lambda_i^2 - \lambda_0^2) \quad (15)$$

$$\tilde{U}(\phi) = K \sum_{i=1}^N U_{\lambda i} \cos 2\phi(\lambda_i^2 - \lambda_0^2) - Q_{\lambda i} \sin 2\phi(\lambda_i^2 - \lambda_0^2). \quad (16)$$

$$\tilde{F}(\phi) = \sqrt{\tilde{Q}^2(\phi) + \tilde{U}^2(\phi)} \quad (17)$$

The above two equations compute the recovered Faraday dispersion function for a given value of ϕ . In practice, the above equations need to be computed for a wide range of values of ϕ to get a sense for the variation of the intrinsic polarization vector as a function of Faraday depth. Thus the above equations become

$$\begin{aligned} \tilde{Q}(\phi_j) &= K \sum_{i=1}^N Q_{\lambda i} \cos 2\phi_j(\lambda_i^2 - \lambda_0^2) + U_{\lambda i} \sin 2\phi_j(\lambda_i^2 - \lambda_0^2); \\ \forall \phi_j &\in [\phi_{\min}, \phi_{\max}] \end{aligned} \quad (18)$$

$$\begin{aligned} \tilde{U}(\phi_j) &= K \sum_{i=1}^N U_{\lambda i} \cos 2\phi_j(\lambda_i^2 - \lambda_0^2) - Q_{\lambda i} \sin 2\phi_j(\lambda_i^2 - \lambda_0^2); \\ \forall \phi_j &\in [\phi_{\min}, \phi_{\max}]. \end{aligned} \quad (19)$$

These two equations need to be computed for each line of sight in a radio image to construct a 3D RM-cube that contains the Faraday dispersion function for all lines of sight.

The use of the window function to alleviate the problem related to incomplete λ^2 coverage implies that the choice of the window function determines the observation's ability to resolve structures along the Faraday depth axis and the range of Faraday depth's for which the sensitivity is greater than 50%. If $\delta\lambda^2$ and $\Delta\lambda^2$ are the channel width and the distribution of λ^2 achieved by the observation, the resolution ($\delta\phi$) and the maximum Faraday depth for which the sensitivity is greater than 50% ($\|\phi_{\max}\|$) are

$$\delta\phi \approx \frac{2\sqrt{3}}{\Delta\lambda^2} \quad (20)$$

$$\|\phi_{\max}\| \approx \frac{\sqrt{3}}{\delta\lambda^2}. \quad (21)$$

2.3. RM synthesis: Computational costs

From Eqs. (15) and (16), it is easy to see that K , 2ϕ , and $\lambda_i^2 - \lambda_0^2$ need to be computed only once for the entire dataset and have negligible compute cost. The compute cost to calculate $\tilde{Q}(\phi)$, $\tilde{U}(\phi)$, and $\tilde{F}(\phi)$ for one value of ϕ using inputs at one λ^2 value is 13 FLOP (7 multiplications, 2 trigonometric functions, 2 additions, 1 subtraction and 1 square root operation).

Thus to calculate $\tilde{Q}(\phi)$, $\tilde{U}(\phi)$, and $\tilde{F}(\phi)$ for one Faraday depth value using N_{chan} frequency channels is

- $(13 + 2) \times N_{\text{chan}}$ FLOP to compute Eqs. (15)–(17) for a given ϕ using all λ_i^2 .
- $2(N_{\text{chan}} - 1)$ additions for summations over all i .
- 2 multiplications to scale the values by K .

The above operations need to be repeated N_{ϕ} times to create RM-cubes of $\tilde{Q}(\phi)$, $\tilde{U}(\phi)$, and $\tilde{F}(\phi)$ with N_{ϕ} Faraday depth planes using N_{chan} input frequency channels. This amounts to a total computation cost of $15N_{\text{chan}}N_{\phi}$ FLOP for a single line of sight. Since each line of sight in the input datacubes can be treated independently, we can scale the computation cost further by the number of spatial

Fig. 1. Rotation Measure Spread Function (RMSF) for a mock dataset consisting of 1000 frequency channels covering the 1–2 GHz frequency range with $\lambda_0^2 = 0$ (left) and $\lambda_0^2 \neq 0$ (right). Notice that introducing the λ_0^2 term reduces the rapid rotation of $\tilde{Q}(\phi)$ and $\tilde{U}(\phi)$.

pixels in the input datacubes making the final computational cost to be $\sim 15N_{\text{chan}}N_{\phi}N_{\text{los}}$.

For a typical polarimetry data set from the Westerbork Synthesis Radio Telescope, for example from the Westerbork Spitzer Nearby Galaxy Survey (SINGS) (Braun et al., 2007; Heald et al., 2009), which has 512^2 spatial pixels, 803 frequency channels and 401 Faraday depth planes, the computational cost required to carry out RM synthesis is about 1.27 TFLOP.⁷

To make this estimate, we assumed that each arithmetic and trigonometric functions amount to a single floating point operation. In practice, this is however not true. For example, standard libraries evaluate trigonometric functions as Taylor series expansions and the number of Taylor terms used in the expansion can be implementation-dependent. Furthermore, the exact number of instructions needed for different mathematical operations available in the standard C libraries can vary depending on the implementation, system hardware and compiler optimization options. Thus the above compute cost should be treated only as a lower limit and not as a precise value.

3. GPU implementation of RM Synthesis

cuFFS is written in C, and GPU acceleration is achieved using the CUDA Application Programming Interface (API). Our code is publicly available through Github.⁸ User input is provided to the program using a configuration file which is parsed using the structured configuration file parser libconfig.⁹ Sample input parset file is shown in Fig. 2.

Input and output data cubes can either be in FITS (Pence, 1999) or in HDFITS (Price et al., 2015) file formats. FITS files are read in and written out using the CFITSIO¹⁰ library while the HDFITS is accessed using the HDF5¹¹ library. Justification for supporting data formats in addition to the standard FITS format is provided in Section 3.2.

A fast CPU implementation of RM synthesis, like the one published in Brentjens (2007), computes the RM cubes by evaluating all Faraday depth values for each input channel. The main advantage of using such a strategy is that it minimizes disk I/O as each input channel is read only once and the output cubes are written to disk only once as long as the output cubes fit in the memory. If all output Faraday depth planes do not fit in memory,

the CPU program computes a subset of Faraday depth planes that fits in memory, writes the output planes to disk, and then proceeds to compute the next subset of Faraday depth planes. However, adapting a similar scheme is not very efficient on a GPU as such an implementation would involve numerous time-consuming data transfers between the *host* (CPU) and the *device* (GPU) especially when the output cubes demand larger memory than what is available on the GPU.

Since data needs to be moved between the CPU and the GPU for processing, computational efficiency can be maximized by minimizing the transfer-to-compute time ratio. This can be achieved by ensuring that a given piece of data is moved between the CPU and the GPU only once. Our GPU implementation ensures this by processing each independent line of sight separately. For a given line of sight, cuFFS carries out the following steps:

1. Transfer $\tilde{Q}(\lambda^2)$ and $\tilde{U}(\lambda^2)$ to the GPU.
2. Each thread on the GPU computes $\tilde{Q}(\phi)$, $\tilde{U}(\phi)$, and $\tilde{F}(\phi)$ for a given Faraday depth ϕ_z .
3. Transfer the computed $\tilde{Q}(\phi)$, $\tilde{U}(\phi)$, and $\tilde{F}(\phi)$ to the CPU.

By repeating these steps for each line of sight in the input datacubes, the program constructs the final RM cubes.

Modern GPUs come with substantial amount of main memory which implies that multiple lines of sight can be processed at the same time. Additionally, transferring a larger chunk of data between host memory and device memory is more efficient than issuing multiple data transfers. As a result, in practice, cuFFS processes multiple lines of sight at once. The program internally estimates the number of lines of sight that can be processed at once based on the amount of memory and the number of processors available on the GPU.

3.1. Science verification

To verify the output produced by cuFFS, we compared the output ($\tilde{F}(\phi)$) produced by cuFFS with a CPU implementation of RM synthesis from Brentjens (2007). The CPU and the GPU codes were used to process a data cube from the Westerbork Synthesis Radio Telescope Spitzer Infrared Nearby Galaxies Survey (WSRT-SINGS; Heald et al., 2009). We compared the output data cubes ($\tilde{F}(\phi)$) on a pixel-by-pixel basis and found that the mean of the ratio between the outputs produced by the GPU and the CPU to be 1.0 with a standard deviation of 1.5×10^{-6} implying that the output produced by cuFFS matches the output produced by the CPU code very well.

Fig. 3 shows the ratio of the pixels values produced by the GPU and the CPU code ($\tilde{F}_{\text{GPU}}(\phi)/\tilde{F}_{\text{CPU}}(\phi)$) plotted as a function of the output pixel values in $\tilde{F}_{\text{GPU}}(\phi)$. As can be seen from Fig. 3, the pixel values in the two output cubes match very well.

⁷ 1 TFLOP = 1 trillion floating-point operations.

⁸ <https://github.com/sarrvesh/cuFFS>.

⁹ <https://github.com/hyperrealm/libconfig>.

¹⁰ <https://heasarc.gsfc.nasa.gov/fitsio/>.

¹¹ <https://support.hdfgroup.org/HDF5/>.

```

// How many GPUs to use?
// In version 1.0, this is always 1.
nGPU = 1;

// What is the input format?
// Can be "FITS" or "HDFITS"
fileFormat = "FITS";

// Define how fits files are stored on disk
qCubeName = "cube-Q.rot.fits";
uCubeName = "cube-U.rot.fits";
freqFileName = "frequencies-Q.txt";

// Define the faraday depth axis.
// All units in rad/m/m
plotRMSF = False;
phiMin = -500.;
nPhi = 401; // integer
dPhi = 2.5;

// Prefix for output filenames
outPrefix = "trial1_";

```

Fig. 2. Sample input parset to cuFFS.

Fig. 3. **Left:** Comparison between the output pixel values produced by cuFFS and a CPU implementation by [Brentjens \(2007\)](#). **Right:** Zoomed-in version of the figure shown in the left panel to better display the scatter in pixel values computed by the GPU and the CPU codes. The scatter seen for the low output pixel values is probably due to numerical error.

3.2. A few comments on file formats

As discussed in the previous section, our code achieves parallelism by exploiting the fact that each line-of-sight and each Faraday depth output for a given line of sight in RM synthesis can be processed independently. However, during the initial development and testing phase, we quickly realized that reading and writing the FITS data cubes along the third data axis (or each line-of-sight) was the major bottleneck in our code. Even for small datacubes like the ones corresponding to WSRT observations, reading and writing datacubes along the third data axis was significantly longer than the amount of time required to carry out the actual RM synthesis computation. In addition to the bottleneck caused by the FITS I/O, it is becoming clear that the FITS data format is not well suited for large multi-dimensional datasets. For an excellent review on the limitations of the FITS data format in the era of big data (radio) astronomy, see [Thomas et al. \(2015\)](#).

One approach to alleviate the slow read/write speed associated with the FITS format would be to rotate (for example using the reorder task in *miriad* ([Sault et al., 1995](#))) the input datacubes

such that the frequency axis is the fastest-varying dimension. After running cuFFS, the output datacubes, which contain the Faraday depth axis as the fastest-varying dimension, can be derotated back to the original order.

Another approach to reduce the read/write time would be to adopt data formats other than the standard FITS file format. For a recent review on new data formats suited for storing astronomical datacubes, see [Shortridge \(2015\)](#). File formats like the Hierarchical Data Format version 5 (HDF5) are well suited for storing large datacubes and have a significantly lower read/write time compared to the FITS format especially while reading along the slowest-varying axis. Recently, [Price et al. \(2015\)](#) proposed a new format called HDFITS which merges the FITS data model with the HDF5 file format. Note that *data model* defines how a dataset is structured and *file format* specifies how the bits are stored on disk. A similar strategy has also been adopted by other data-intensive fields like climatology and oceanography where the NetCDF data format is built on top of the HDF5 file format. The main advantage of using HDFITS in our program is that the structure of the dataset stored in the HDF5 file is similar to the standard FITS format with the added advantage of higher read/write speed.

Table 1

Comparison of CPU and GPU execution times for different input and output data cubes.

Dataset	Number of			Execution time [hh:mm:ss]		
	Spatial pixels	Channels	ϕ planes	CPU ^a	GPU ^a	GPU ^b
WSRT	512 × 512	803	401	00:02:02	00:00:50 ^c + 00:00:08 ^d	00:00:25
LOFAR_LR	540 × 540	856	7000	00:39:30	00:04:37 ^c + 00:00:17 ^d	00:04:36
LOFAR_HR	4000 × 4000	856	7000	58:01:32	12:53:39 ^c + 00:13:47 ^d	02:28:02

Notes. CPU tests using a single threaded and vectorized code from (Brentjens, 2007) were carried out on a single node with Intel(R) Xeon(R) CPU E5-2690 processor at 2.6 GHz clock with 256 GB RAM. The CPU code was permitted to use a maximum of 200 GB RAM. GPU tests was carried out using an NVIDIA Tesla P100-SXM2-16 GB device.

^aFITS file format.

^bHDFITS format (the reported time includes time spent reordering the input datacube and the compute time).

^cTime spent reordering the input datacube to have the frequency axis as the first axis and reordering the output datacubes so that the Faraday depth axis is the third axis

^dCompute time.

To achieve higher read/write throughput and to be consistent with existing data format, the current version of cuFFS supports input and output datacubes in both FITS and HDFITS formats. While operating in the FITS mode, cuFFS assumes that the user has rotated the input cubes such that the frequency axis is the fast-varying dimension. Since HDFITS is efficient in reading along the frequency axis (see Table 1), no such prior rotation is required while using cuFFS in HDFITS mode. Note that regardless of the input format, the output Faraday depth cubes are ordered such that Faraday depth is the fast-varying dimension.

As mentioned before, numerous data formats have been proposed in the recent years for archiving large datasets. If desired, cuFFS can be modified to support additional data formats like N-dimensional Data Format (NDF; Jenness et al., 2015), Starlink Hierarchical Data System (HDS; Jenness, 2015) and Advanced Scientific Data Format (ASDF; Greenfield et al., 2015) with minimal change to the existing code.

3.3. Benchmarks and performance measurements

We compared the execution time of our GPU code with a single-threaded and vectorized CPU implementation of RM Synthesis from Brentjens (2007) using a combination of both real and mock datasets from current and upcoming telescope facilities. A brief description of each test dataset is presented below:

- **WSRT:** As a typical Westerbork polarimetry dataset, we used one of the datacubes produced as part of the Westerbork SINGS survey (WSRT-SINGS; Heald et al., 2009). After removing frequency channels corrupted by RFI, the final datacube consists of 803 frequency channels. The frequency setup used in the WSRT-SINGS survey results in a Faraday depth resolution of 144 rad m⁻². Based on these values, the input and the output datasets have dimensions 512 × 512 × 803 and 512 × 512 × 401 respectively.
- **LOFAR_LR:** Diffuse polarized emission from the Galactic foreground can be detected at 150 MHz in low resolution LOFAR maps with a resolution of a few arcmins (Jelić et al., 2015; Van Eck et al., 2017). The LOFAR low resolution dataset used for this test corresponds to a field around the nearby dwarf galaxy NGC 1569. The input datacubes are at a resolution of 250". The input Stokes Q and U cubes contain 856 frequency channels which are free of any radio frequency interference. With 48.8 kHz-wide channel maps covering the bandwidth ranging from 120.238 to 166.087 MHz, the output Faraday cubes will have a Faraday depth resolution of 1.17 rad m⁻² and sensitive to a maximum Faraday depth value of about 350 rad m⁻². Thus the dimensions of the input and the output datacubes are 540 × 540 × 856 and 540 × 540 × 7000.
- **LOFAR_HR:** The LOFAR high resolution dataset was produced from the same visibility data that was used to create the LOFAR_LR datacubes. The only difference between the LOFAR_HR and LOFAR_LR is that LOFAR_HR data cubes are at

a higher angular resolution of 26". The dimensions of the input and the output datacubes are 4000 × 4000 × 856 and 4000 × 4000 × 7000.

All CPU and GPU benchmarks were carried out on CSIRO's Bracewell computer cluster located in Canberra, Australia. Each node on the cluster is powered by 28 Intel(R) Xeon(R) CPU E5-2690 v4 2.60 GHz processors with 256 GB RAM and is equipped with four NVidia Tesla P100-SXM2-16 GB GPUs. The CPU code was permitted to use up to 200 GB of RAM. The execution times for different datacubes using the CPU and the GPU implementation are listed in Table 1. Note that the execution times for the CPU code will increase almost linearly with decreasing RAM size.

From the execution times reported in Table 1, it can be seen quite clearly that our GPU implementation in both FITS and HDFITS mode is substantially faster than the execution times achieved using the CPU implementation even though the CPU implementation was allowed to use 15 times more main memory than our GPU program. Execution times reported in Table 1 while using cuFFS in FITS mode indicate that rotating and derotating the input and output datacubes still dominate the total execution time. If, however, the data rotation is done as part of the imaging pipeline that produces the Stokes Q and U cubes, then the GPU implementation is about two orders of magnitude faster than the CPU implementation for large datasets.

4. Conclusion and future outlook

In this paper, we have presented the GPU-accelerated CUDA program cuFFS, or CUDA-accelerated Fast Faraday Synthesis, to perform a widely used radio astronomy algorithm called rotation measure synthesis. The software package is made publicly available through Github. The software package is capable of handling input and output datacubes in either the FITS or the HDFITS format. If needed, additional data formats can also be supported with very little change to the existing code base. In addition to rotation measure synthesis, we plan to implement advanced polarization processing tools like RM-CLEAN (Heald, 2009) in the future version of our software package.

Comparing the execution time of our GPU implementation on NVidia Tesla P100-SXM2-16 GB GPU with a single-threaded, vectorized CPU implementation, we observe gains in execution times by at least a factor of 15 for smaller datasets from telescopes like WSRT and by at least by two orders of magnitude for larger datasets from telescopes like LOFAR, depending on the input and output data format. Further increase in speed can be achieved by using Fast Fourier Transforms (FFT) instead of Discrete Fourier Transforms (DFT) in Eqs. (15) and (16). CUDA accelerated Fast Fourier Transform can be accelerated using standard libraries like cuFFT.¹² We aim to implement these advanced features in future versions of cuFFS.

¹² <https://developer.nvidia.com/cufft>.

Finally, as discussed in Section 3.3 and Table 1, the main bottleneck in our implementation of RM synthesis is the FITS file access using CFITSIO. Even for a small WSRT dataset, rotating the cubes to speed up I/O can be an order of magnitude longer than the actual compute time. In light of the large data cubes that will be produced by the current and upcoming radio telescopes, this I/O overload is certainly a cause for concern. To minimize the amount of time spent rearranging data, we make the following two recommendations to teams carrying out large polarization surveys:

- Adopt data formats like HDFITS which benefit from combining the FITS data model with the HDF5 file format. Accessing an array stored in HDFITS format along the slowest-varying dimension can be two orders of magnitude faster than FITS format.
- If surveys prefer to use the standard FITS format to store their data products, they should at least consider archiving the data products not intended for visualization (like $Q(\lambda)$, $U(\lambda)$, $Q(\phi)$, and $U(\phi)$ cubes) with frequency and Faraday depth as the first/fastest-varying axis.

Acknowledgments

We thank the anonymous referee for their helpful comments that improved the manuscript. SSS acknowledges financial support from the Leids Kerkhoven-Bosscha Fonds (LKBF) for multiple visits to CSIRO in Perth, Australia where most of this work was carried out. SSS thanks Christopher Harris at the Pawsey Supercomputing Centre for useful discussions on CUDA programming. JMvdH acknowledges support from the European Research Council under the European Union's Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement no. 291531. We thank the IMT Scientific Computing group at CSIRO for their expertise and support of our work on the CSIRO Bracewell GPU cluster.

References

- Bassa, C.G., Pleunis, Z., Hessels, J.W.T., 2017. Enabling pulsar and fast transient searches using coherent dedispersion. *Astron. Comput.* 18, 40–46. doi:10.1016/j.ascom.2017.01.004, arXiv:1607.00909.
- Beck, R., 2016. Magnetic fields in spiral galaxies. *Astron. Astrophys. Rev.* 24, 4. doi:10.1007/s00159-015-0084-4, arXiv:1509.04522.
- Bouchez, A.H., Acton, D.S., Basic, R., Conard, R., Espelanda, B., Esposito, S., Filgueira, J., Gallien, D., McLeod, B.A., Pinae, E., et al., 2012. The giant magellan telescope adaptive optics program. *Proc. SPIE* 9148, 91480W–1.
- Braun, R., Oosterloo, T.A., Morganti, R., Klein, U., Beck, R., 2007. The Westerbork SINGS survey. I. Overview and image atlas. *Astron. Astrophys.* 461, 455–470. doi:10.1051/0004-6361/20066092, arXiv:astro-ph/0610114.
- Brentjens, M.A., 2007. Radio polarimetry in 2.5D (Ph.D. thesis), University of Groningen.
- Brentjens, M.A., de Bruyn, A.G., 2005. Faraday rotation measure synthesis. *Astron. Astrophys.* 441, 1217–1228. doi:10.1051/0004-6361:20052990, arXiv:astro-ph/0507349.
- Burn, B.J., 1966. On the depolarization of discrete radio sources by Faraday dispersion. *Mon. Not. R. Astron. Soc.* 133, 67.
- Chennamangalam, J., Scott, S., Jones, G., Chen, H., Ford, J., Kepley, A., Lorimer, D.R., Nie, J., Prestage, R., Rosh, D.A., Wagner, M., Werthimer, D., 2014. A GPU-Based Wide-Band Radio Spectrometer. *Publ. Astron. Soc. Aust.* 31, e048. doi:10.1017/pasa.2014.43, arXiv:1411.0436.
- Elsen, E., Vishal, V., Houston, M., Pande, V., Hanrahan, P., Darve, E., 2007. N-body simulations on GPUs. *ArXiv e-prints* arXiv:0706.3060.
- Fletcher, A., Klein, U., 2015. *Galactic and Intergalactic Magnetic Fields*. Heidelberg: Springer.
- Fluke, C.J., 2012. Accelerating the rate of astronomical discovery with GPU-powered clusters. In: Ballester, P., Egret, D., Lorente, N.P.F. (Eds.), *Astronomical Data Analysis Software and Systems XXI*. In: *Astronomical Society of the Pacific Conference Series*, vol. 461, p. 3, arXiv:1111.5081.
- Fluke, C.J., Barnes, D.G., Barsdell, B.R., Hassan, A.H., 2011. Astrophysical supercomputing with GPUs: critical decisions for early adopters. *Publ. Astron. Soc. Aust.* 28, 15–27. doi:10.1071/AS10019, arXiv:1008.4623.
- Greenfield, P., Droettboom, M., Bray, E., 2015. ASDF: A new data format for astronomy. *Astron. Comput.* 12, 240–251. doi:10.1016/j.ascom.2015.06.004.
- Hamaker, J.P., Bregman, J.D., Sault, R.J., 1996. *Understanding radio polarimetry. I. Mathematical foundations*. *Astron. Astrophys. Suppl. Ser.* 117, 137–147.
- Harris, C., Haines, K., Staveley-Smith, L., 2008. GPU accelerated radio astronomy signal convolution. *Exp. Astron.* 22, 129–141. doi:10.1007/s10686-008-9114-9.
- Heald, G., 2009. The Faraday rotation measure synthesis technique. In: Strassmeier, K.G., Kosovichev, A.G., Beckman, J.E. (Eds.), *IAU Symposium*. In: *IAU Symposium*, Vol. 259, pp. 591–602. doi:10.1017/S1743921309031421.
- Heald, G., Braun, R., Edmonds, R., 2009. The Westerbork SINGS survey. II Polarization, Faraday rotation, and magnetic fields. *Astron. Astrophys.* 503, 409–435. doi:10.1051/0004-6361/200912240, arXiv:0905.3995.
- Jelić, V., de Bruyn, A.G., Pandey, V.N., Mevius, M., Haverkorn, M., Brentjens, M.A., Koopmans, L.V.E., Zaroubi, S., Abdalla, F.B., Asad, K.M.B., Bus, S., Chapman, E., Ciardi, B., Fernandez, E.R., Ghosh, A., Harker, G., Iliev, I.T., Jensen, H., Kazemi, S., Mellema, G., Offringa, A.R., Patil, A.H., Vedantham, H.K., Yatawatta, S., 2015. Linear polarization structures in LOFAR observations of the interstellar medium in the 3C 196 field. *Astron. Astrophys.* 583, A137. doi:10.1051/0004-6361/201526638, arXiv:1508.06650.
- Jenness, T., 2015. Reimplementing the Hierarchical Data System using HDF5. *Astron. Comput.* 12, 221–228. doi:10.1016/j.ascom.2015.02.003, arXiv:1502.04029.
- Jenness, T., Berry, D.S., Currie, M.J., Draper, P.W., Economou, F., Gray, N., McIlwrath, B., Shortridge, K., Taylor, M.B., Wallace, P.T., Warren-Smith, R.F., 2015. Learning from 25 years of the extensible N-dimensional data format. *Astron. Comput.* 12, 146–161. doi:10.1016/j.ascom.2014.11.001, arXiv:1410.7513.
- Magro, A., Karastergiou, A., Salvini, S., Mort, B., Dulwich, F., Zarb Adami, K., 2011. Real-time, fast radio transient searches with GPU de-dispersion. *Mon. Not. R. Astron. Soc.* 417, 2642–2650. doi:10.1111/j.1365-2966.2011.19426.x, arXiv:1107.2516.
- Michalakes, J., Vachharajani, M., 2008. GPU acceleration of numerical weather prediction. *Parallel Process. Lett.* 18 (04), 531–548.
- Nvidia, 2017. *CUDA Programming Guide* 8.0.
- Ord, S., Greenhill, L., Wayth, R., Mitchell, D., Dale, K., Pfister, H., Edgar, R., 2009. Graphics Processing Units for Data Processing in the Murchison Wide-field Array. In: Bohlender, D.A., Durand, D., Dowler, P. (Eds.), *Astronomical Data Analysis Software and Systems XVIII*. In: *Astronomical Society of the Pacific Conference Series*, vol. 411, p. 127.
- Owens, J.D., Houston, M., Luebke, D., Green, S., Stone, J.E., Phillips, J.C., 2008. GPU computing. *Proc. IEEE* 96 (5), 879–899.
- Owens, J.D., Luebke, D., Govindaraju, N., Harris, M., Kruger, J., Lefohn, A.E., Purcell, T.J., 2007. A survey of general-purpose computation on graphics hardware. *Comput. Graph. Forum* 26 (1), 80–113. doi:10.1111/j.1467-8659.2007.01012.x.
- Patterson, D., Hennessy, J., 2014. Computer organization and design, enhanced: the hardware/software interface. The Morgan Kaufmann Series in Computer Architecture and Design, Elsevier Science, <https://books.google.nl/books?id=G7IMAwAAQBAJ>.
- Pence, W., 1999. CFITSIO, v2.0: A new full-featured data interface. In: Mehringer, D.M., Plante, R.L., Roberts, D.A. (Eds.), *Astronomical Data Analysis Software and Systems VIII*. In: *Astronomical Society of the Pacific Conference Series*, vol. 172, p. 487.
- Perkins, S.J., Marais, P.C., Zwart, J.T.L., Natarajan, I., Tasse, C., Smirnov, O., 2015. Montblanc¹: GPU accelerated radio interferometer measurement equations in support of Bayesian inference for radio observations. *Astron. Comput.* 12, 73–85. doi:10.1016/j.ascom.2015.06.003, arXiv:1501.07719.
- Portegies Zwart, S.F., Bellemann, R.G., Geldof, P.M., 2007. High-performance direct gravitational N-body simulations on graphics processing units. *New Astronomy* 12, 641–650. doi:10.1016/j.newast.2007.05.004, arXiv:cs/0702135.
- Price, D.C., Barsdell, B.R., Greenhill, L.J., 2015. HDFITS: Porting the FITS data model to HDF5. *Astron. Comput.* 12, 212–220. doi:10.1016/j.ascom.2015.05.001, arXiv:1505.06421.
- Price, D.C., Staveley-Smith, L., Bailes, M., Carretti, E., Jameson, A., Jones, M.E., van Straten, W., Schediwy, S.W., 2016. HIPSR: A digital signal processor for the parkes 21-cm multibeam receiver. *J. Astronom. Instrum.* 5, 1641007. doi:10.1142/S2251171716410075, arXiv:1702.00443.
- Sault, R.J., Teuben, P.J., Wright, M.C.H., 1995. A retrospective view of MIRIAD. In: Shaw, R.A., Payne, H.E., Hayes, J.J.E. (Eds.), *Astronomical Data Analysis Software and Systems IV*. In: *Astronomical Society of the Pacific Conference Series*, vol. 77, p. 433, arXiv:astro-ph/0612759.
- Schaaf, K., Overeem, R., 2004. Cots correlator platform. *Exp. Astron.* 17, 287–297. doi:10.1007/s10686-005-2864-8.
- Schwiartz, T., Chang, T.C., Speier, P., Westermann, R., 2006. MR image reconstruction using the GPU. *Proc. SPIE* 6142 (1), 61423T.
- Schomberg, H., Timmer, J., 1995. The gridding method for image reconstruction by fourier transformation. *IEEE Trans. Med. Imag.* 14 (3), 596–607.
- Shortridge, K., 2015. A glimpse at different file formats. In: Taylor, A.R., Rosolowsky, E. (Eds.), *Astronomical Data Analysis Software and Systems XXIV (ADASS XXIV)*. In: *Astronomical Society of the Pacific Conference Series*, vol. 495, p. 527.
- Thomas, B., Jenness, T., Economou, F., Greenfield, P., Hirst, P., Berry, D.S., Bray, E., Gray, N., Muna, D., Turner, J., de Val-Borro, M., Santander-Vela, J., Shupe, D., Good, J., Berriman, G.B., Kitaeff, S., Fay, J., Laurino, O., Alexov, A., Landry, W., Masters, J., Brazier, A., Schaaf, R., Edwards, K., Redman, R.O., Marsh, T.R., Streicher, O.,

- Norris, P., Pascual, S., Davie, M., Droettboom, M., Robitaille, T., Campana, R., Hagen, A., Hartogh, P., Klaes, D., Craig, M.W., Homeier, D., 2015. Learning from FITS: Limitations in use in modern astronomical research. *Astron. Comput.* 12, 133–145. doi:[10.1016/j.ascom.2015.01.009](https://doi.org/10.1016/j.ascom.2015.01.009), arXiv:[1502.00996](https://arxiv.org/abs/1502.00996).
- Trancoso, P., Charalambous, M., 2005. Exploring graphics processor performance for general purpose applications. In: *Digital System Design, 2005 Proceedings 8th Euromicro Conference on*. IEEE, pp. 306–313.
- Vallee, J.P., 1980. The rotation measures of radio sources and their data processing. *Astron. Astrophys.* 86, 251–253.
- Van Eck, C.L., Haverkorn, M., Alves, M.I.R., Beck, R., de Bruyn, A.G., Enßlin, T., Farnes, J.S., Ferrière, K., Heald, G., Horellou, C., Horneffer, A., Iacobelli, M., Jelić, V., Martí-Vidal, I., Mulcahy, D.D., Reich, W., Röttgering, H.J.A., Scaife, A.M.M., Schnitzeler, D.H.F.M., Sobey, C., Sridhar, S.S., 2017. Faraday tomography of the local interstellar medium with LOFAR: Galactic foregrounds towards IC 342. *Astron. Astrophys.* 597, A98. doi:[10.1051/0004-6361/201629707](https://doi.org/10.1051/0004-6361/201629707), arXiv:[1612.00710](https://arxiv.org/abs/1612.00710).
- Zink, B., 2011. HORIZON: Accelerated general relativistic magnetohydrodynamics. ArXiv e-prints, arXiv:[1102.5202](https://arxiv.org/abs/1102.5202).