Fast, Scalable and Accurate Wide-Field (De)Gridding via the 3D NUFFT

Swiss SKA Days 2022

Matthieu Simeoni, Joan Rue Queralt, Sepand Kashani

October 4, 2022

EPFL

Analysis and Synthesis for Imaging Purposes

Analysis (a.k.a. **gridding**) and synthesis (a.k.a. **degridding**) are used in radio interferometry to accurately predict visibilities (resp. dirty images) from a given sky distribution (resp. visibility set):

$$\left(V_i = \sum_{\boldsymbol{r} \in \Theta_{\text{pix}}} I(\boldsymbol{r}) e^{-j\langle \boldsymbol{r}, \boldsymbol{p}_i \rangle} \right)_{i=1,\dots,N_{\text{vis}}} \qquad \& \qquad \left(I_D(\boldsymbol{r}) = \frac{1}{N_{\text{vis}}} \sum_{i=1}^{N_{\text{vis}}} V_i e^{j\langle \boldsymbol{r}, \boldsymbol{p}_i \rangle} \right)_{\boldsymbol{r} \in \Theta_{\text{pix}}}.$$

Direct evaluation of the above expressions has bi-linear complexity $\mathcal{O}(N_{\text{vis}} N_{\text{pix}})$ which represents a major bottleneck for imaging purposes:

- CLEAN performs **one analysis and synthesis per** major cycle.
- Proximal algorithms typically perform **one analysis and synthesis per** iteration.

The (numerous!) analysis and synthesis steps tend to dominate the overall computational complexity of the imaging task.

We therefore need fast, scalable and yet accurate analysis and synthesis algorithms.

Fast Synthesis Algorithms

Let us focus on the synthesis step.[1]

Assuming that the image is defined over a regular grid in direction cosines coordinates $(\ell, m) \in \mathbb{R}^2$, the latter can be expressed as:

$$\left(I_D(\ell, m) = \frac{1}{N_{\text{vis}}} \sum_{i=1}^{N_{\text{vis}}} V_i e^{j(\ell u_i + m v_i + n(\ell, m) w_i)} \right)_{(\ell, m) \in \Omega_{\text{pix}}}$$

where $n(\ell, m) := \sqrt{1 - \ell^2 - m^2} - 1$ and $(u_i, v_i, w_i) \in \mathbb{R}^3$ the (normalised) baselines' uvw-coordinates.

Resembles very much a partial Fourier transform, although with non-uniform frequencies $(u_i, v_i, w_i) \in \mathbb{R}^3$ and spatial coordinates evolving on a manifold. Can we leverage the FFT to achieve log-linear complexity?

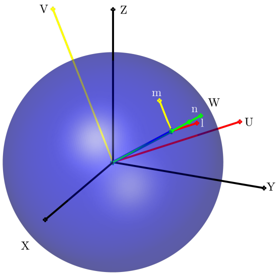

Figure: Direction cosines and uvw-coordinates.

[1] Since analysis and synthesis are adjoint operations, any (linear) evaluation algorithm obtained for the synthesis step can be used for the analysis step by simply reversing the chain of operations involved.

Fast Synthesis Algorithms (continued)

Two cases:

`Small FOV:` $n(\ell, m) \simeq 0 \Rightarrow$ we can neglect the w-term, yielding

$$I_D(\ell, m) = \frac{1}{N_{\text{vis}}} \sum_{i=1}^{N_{\text{vis}}} V_i e^{j(\ell u_i + m v_i)}.$$

Since the direction cosine coordinates are uniformly chosen, we can then use a 2D Non Uniform FFT (NUFFT) of Type 1 to evaluate the sum.

`Large FOV:` $n(\ell, m) \neq 0 \Rightarrow$ much harder! Many algorithms proposed: image faceting, W-projection, W-snapshot, W-stacking, W-gridding, etc.

Most used is probably w-stacking which proposes to cluster the baseline w coordinates in N_w planes, each processed independently by a 2D NUFFT and a modulation before being summed:

$$I_D(\ell, m) = \frac{1}{N_{\text{vis}}} \sum_{k=1}^{N_w} e^{j(n(\ell,m) w_k)} \left(\sum_{i \in P_k} V_i e^{j(\ell u_i + m v_i)} \right).$$

Standard Gridding

Fourier Domain

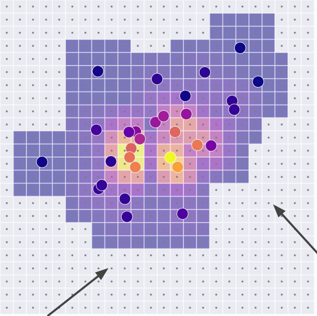

Uniform grid
of harmonics

Nonuniform
frequencies

2D NUFFT
(Type 1)

$$\sum_i V_i e^{j(\ell u_i + m v_i)}$$

2D FFT

(uniform to uniform)

$$\mathcal{O}(N_{\text{pix}} \log N_{\text{pix}})$$

Spatial Domain

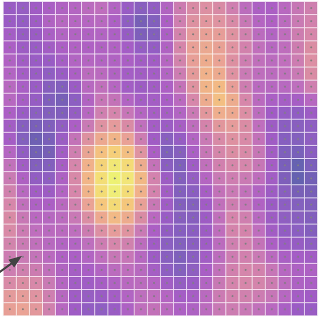

Uniform grid
of *pixels*

Spreading
(nonuniform to uniform)

$$\mathcal{O}(N_{\text{vis}} |\log \epsilon|^2)$$

$*$

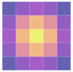

Tapering
(corrects for spreading)

$$\mathcal{O}(N_{\text{pix}})$$

X

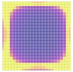

Total Complexity:

$$\mathcal{O}(N_{\text{vis}} |\log \epsilon|^2 + N_{\text{pix}}(\log N_{\text{pix}} + 1))$$

Nifty Gridder (a.k.a. W-gridding)

Fourier Domain

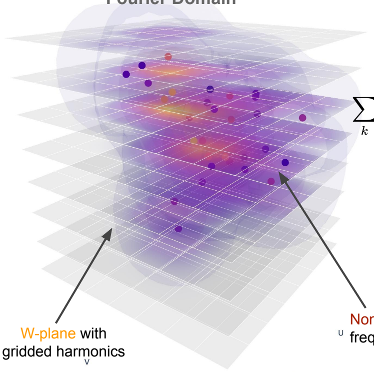

W-plane with gridded harmonics

Nonuniform frequencies

W-Spreading
(nonuniform to W-plane)
$\mathcal{O}(N_{\text{vis}}\boxed{|\log\epsilon|^3})$

$*$ 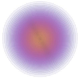

Nifty Gridder (a.k.a. W-Gridding)

$$\sum_k e^{jn(\ell,m)w_k}\left[\sum_{i\in\mathcal{I}_k} V_i e^{j(\ell u_i+mv_i)}\right]$$

2D FFTs
(for each W-plane)
$\mathcal{O}(\boxed{N_{w'}}N_{\text{pix}}\log N_{\text{pix}})$

Spatial Domain

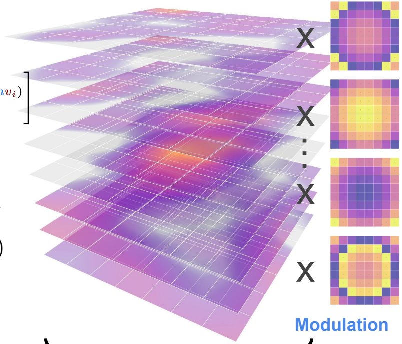

X X X X

Modulation
$\mathcal{O}(\boxed{N_{w'}}N_{\text{pix}})$

+ Sum

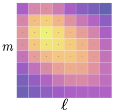

X

$\mathcal{O}(N_{\text{pix}})$

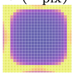

Tapering

Total Complexity:

$$\mathcal{O}(N_{\text{vis}}\boxed{|\log\epsilon|^3}+\boxed{N_{w'}}N_{\text{pix}}(\log N_{\text{pix}}+1))$$

Limitations of Direction Cosine Grids

Current gridding techniques achieve log-linear complexity, but only for direction cosine meshgrids. The latter are ill-suited for large FOVs due to **heavily distorted** or singular/non-feasible pixels away from the center. This significantly complicates (if not *forbids*) the use of traditional **image processing** and **image analysis** tools.

In contrast the HEALPix mesh is tailored to spherical images, and **readily supports** Fourier and wavelet analysis, filtering, CNNs for inference/classification, pattern recognition, noise analysis and more.

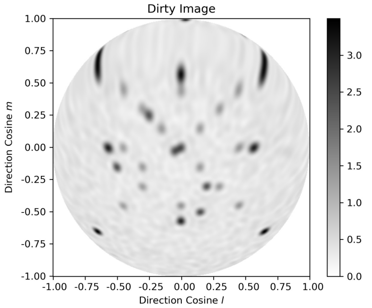

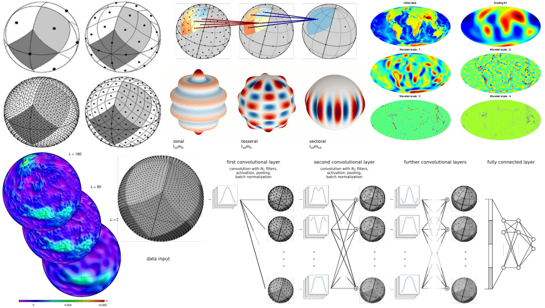

Gridding via the 3D NUFFT of Type 3

Fourier Domain

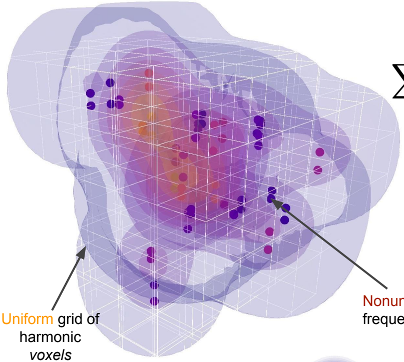

3D NUFFT (Type 3)

$$\sum_i V_i e^{j(\ell u_i + m v_i + n w_i)}$$

3D FFT

(uniform to uniform)

$$\mathcal{O}(\boxed{N_{\text{mesh}} \log N_{\text{mesh}}})$$

Spatial Domain

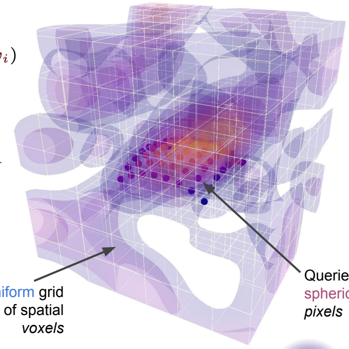

Spreading (+Tapering)

(nonuniform to uniform)

$$\mathcal{O}(N_{\text{vis}} |\log \epsilon|^3 + \boxed{N_{\text{mesh}}})$$

$*$

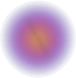

Interpolation (+Tapering)

(uniform to nonuniform)

$$\mathcal{O}(\boxed{N_{\text{pix}} |\log \epsilon|^3 + N_{\text{pix}}})$$

$*$

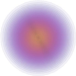

Total Complexity:

$$\mathcal{O}((N_{\text{vis}} + \boxed{N_{\text{pix}}}) |\log \epsilon|^3 + \boxed{N_{\text{mesh}}(\log N_{\text{mesh}} + 1) + N_{\text{pix}}})$$

Scaling up via Chunked NUFFT Gridding

The NUFFT of Type 3 performs a single 3D FFT with critical size:

$$N_{\text{mesh}} \propto (u_{\max} v_{\max} w_{\max}) \times (\ell_{\max} m_{\max} n_{\max})$$

To help reduce down the FFT sizes, we propose partitioning the uvw and image domains in chunks:

$$\sum_{C \in \mathscr{P}} \sum_{i \in C} V_i e^{j(\ell u_i + m v_i + n w_i)}$$

In the uvw domain, the partitioning is performed via a quadtree to ensure similar workloads in each chunk.

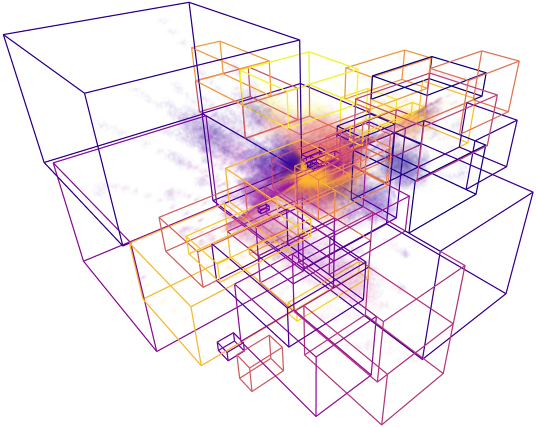

Figure: Quadtree partitioning in the uvw domain.

HEALPix Synthesis via the 3D NUFFT of Type 3

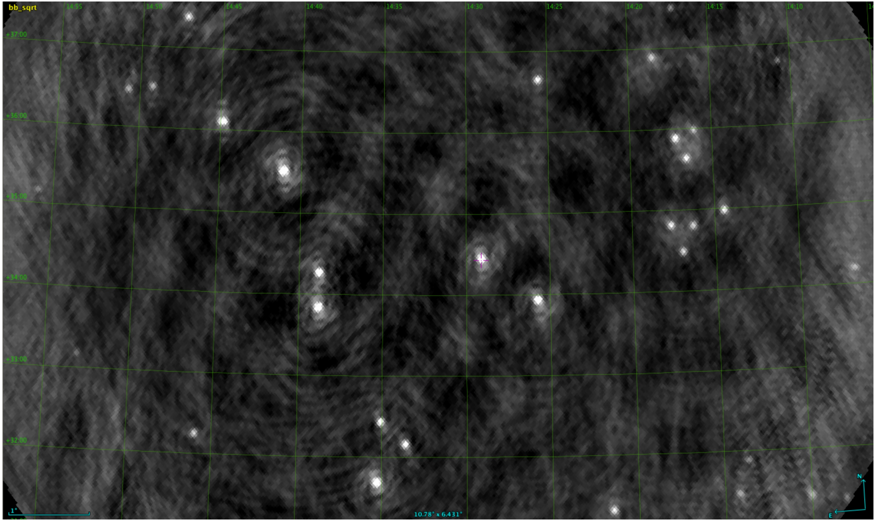

Benchmarking: Execution Time (SKA-LOW, 30°)

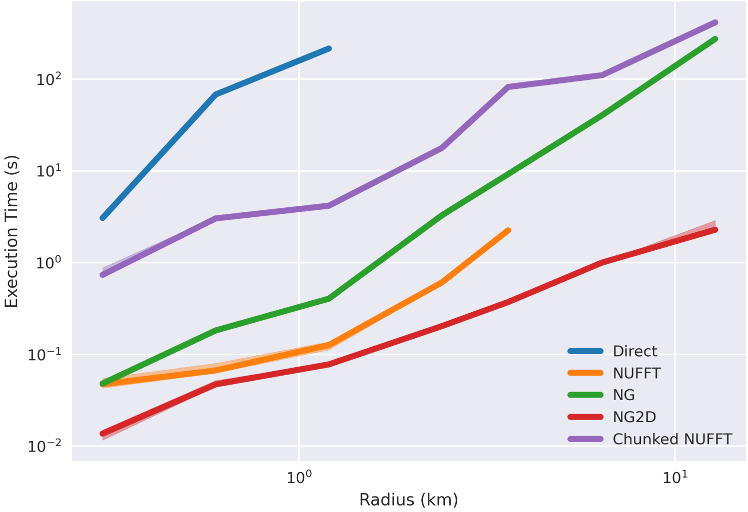

Benchmarking: Memory Usage (SKA-LOW, 30°)

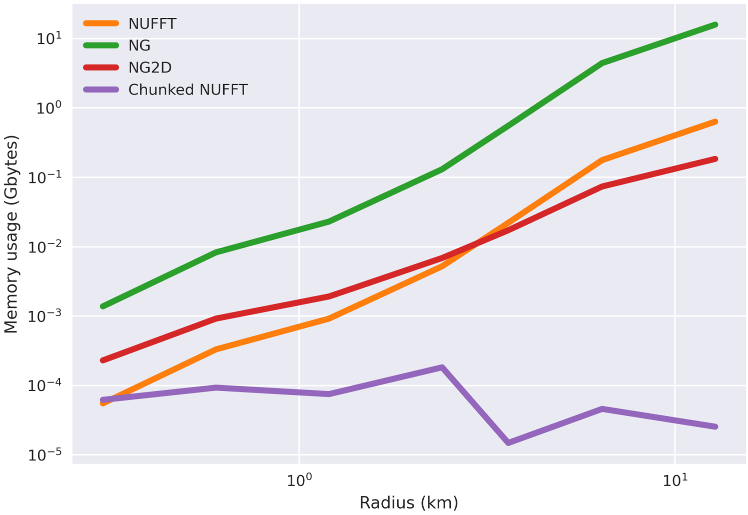

Benchmarking: Relative Accuracy (SKA-LOW, 30°, $r_{\text{max}} = 600m$)

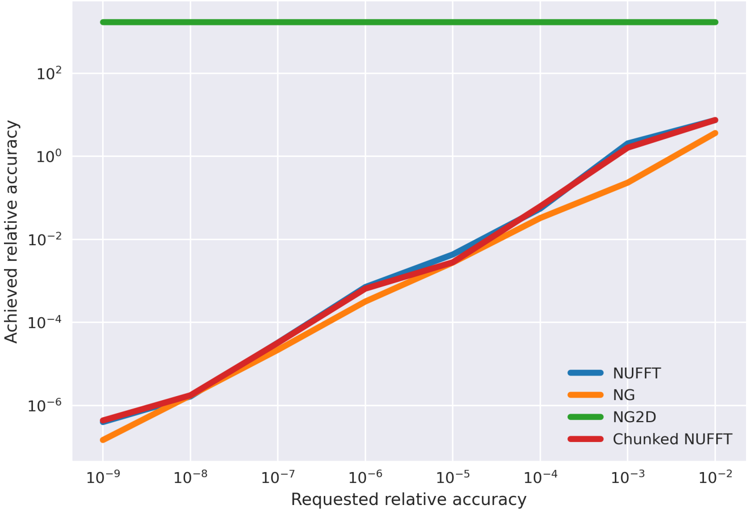

Conclusion

We have proposed the NUFFT gridder, a new wide-field (de)gridding algorithm for radio interferometry with:

- support for arbitrary spherical meshes (including **HEALPix**),
- similar accuracy but faster than the **Nifty gridder** (s.o.a),
- good scaling behaviour thanks to a chunking strategy (opens up the way to stochastic optimisation too...).

Next steps:

- release the NUFFT gridder as an **extension to RASCIL**,
- fix issues related to unmanaged memory encountered in very large setups,
- submit paper (**on the way**),
- deploy and test in a production environment,
- extend the algorithm to support A-terms and compare to Image Domain Gridding (IDG).