Cygrid: A fast Cython-powered convolution-based gridding module for Python[★]

B. Winkel¹, D. Lenz², and L. Flöer²

¹ Max-Planck-Institut für Radioastronomie (MPIfR), Auf dem Hügel 69, 53121 Bonn, Germany
e-mail: bwinkel@mpi.fr

² Argelander-Institut für Astronomie (AIfA), Auf dem Hügel 71, 53121 Bonn, Germany

Received 9 March 2016 / Accepted 19 April 2016

ABSTRACT

Context. Data gridding is a common task in astronomy and many other science disciplines. It refers to the resampling of irregularly sampled data to a regular grid.

Aims. We present cygrid, a library module for the general purpose programming language Python. Cygrid can be used to resample data to any collection of target coordinates, although its typical application involves FITS maps or data cubes. The FITS world coordinate system standard is supported.

Methods. The regridding algorithm is based on the convolution of the original samples with a kernel of arbitrary shape. We introduce a lookup table scheme that allows us to parallelize the gridding and combine it with the HEALPix tessellation of the sphere for fast neighbor searches.

Results. We show that for n input data points, cygrids runtime scales between $O(n)$ and $O(n \log n)$ and analyze the performance gain that is achieved using multiple CPU cores. We also compare the gridding speed with other techniques, such as nearest-neighbor, and linear and cubic spline interpolation.

Conclusions. Cygrid is a very fast and versatile gridding library that significantly outperforms other third-party Python modules, such as the linear and cubic spline interpolation provided by SciPy.

Key words. methods: numerical – techniques: image processing

1. Introduction

In many natural sciences, observational data are often irregularly sampled. For example, in astronomy a single-dish telescope actively scans the sky to produce a map. Even if the scan pattern was designed well, the measured data point locations might deviate from a perfect rectangular grid, for example, caused by pointing inaccuracies, the spherical geometry of the sky, or distortions in the optics. However, a regular grid is needed to visualize the data on pixel-based devices, enable algorithmic analysis, or store the data in FITS images. The task of resampling a data set from irregular spacing to a regular grid is called gridding.

There are already many different gridding techniques known in the literature. Some are easy to understand and implement, such as nearest-neighbor resampling, or linear and cubic spline interpolation, others involve complex statistical inference, for example, the Kriging algorithm (Krig 1951; Matheron 1963). Despite the fact that many available methods are fast and memory efficient, they are not always the best choice for astronomical applications because they do not conserve the flux density of a source.

There is a sufficiently fast alternative, convolution-based gridding, which is well known in many disciplines, especially in radio astronomy. For example, it is included in the long-known and mature interferometry software package AIPS¹ and its successor CASA² (sdgrid task). It is also used by the GILDAS

software³. Furthermore, Kalberla et al. (2005, 2010) apply this gridding technique to produce (all-sky) HI data cubes and column density maps for the Leiden/Argentine/Bonn Survey (LAB; Kalberla et al. 2005) and the Galactic All-Sky Survey (GASS; McClure-Griffiths et al. 2009). We briefly recapitulate this algorithm in Sect. 2.1.

As long as the applied gridding kernel is sampled with sufficient spatial density on the target grid, convolution-based gridding conserves flux densities. The algorithm can easily be implemented to work serially, which is beneficial for its memory footprint. To parallelize this scheme, we need to make sure that two threads never write to the same output element. In Sect. 2.2, we develop an efficient lookup scheme, which is based on the HEALPix tessellation of the sphere (Górski et al. 2005), to optimally distribute the work across different threads. Another advantage of our proposed HEALPix lookup scheme is that it brings the time complexity close to $O(n)$, where n is the number of input samples even for arbitrary world coordinate system (WCS; Greisen & Calabretta 2002; Calabretta & Greisen 2002) projections, which is otherwise not easily possible. In Sect. 3 we discuss our reference implementation for the Python⁴ programming language, the cygrid module. It was developed in the framework of the Effelsberg–Bonn HI Survey (EBHIS; Kerp et al. 2011; Winkel et al. 2016). Cygrid makes use of Cython⁵ to improve the speed of some performance-critical

[★] <https://github.com/bwinkel/cygrid>

¹ <http://www.aips.nrao.edu/index.shtml>

² <http://casa.nrao.edu/>

³ <http://www.iram.fr/IRAMFR/GILDAS/>

⁴ <http://www.python.org/>

⁵ <http://cython.org/>

functions. Section 4 presents speed tests and benchmarks. We conclude with a summary in Sect. 5.

2. The gridding algorithm

In astronomy, gridding tasks often have to contend with (spherical) sky coordinates. The FITS WCS standard and its various software implementations are most convenient for the conversion between world coordinates and pixel grids. In the following, we constrain ourselves to the astronomy-specific case of gridding into WCS pixel grids.

2.1. Basic method

The basic idea of convolution-based gridding is to compute the discrete convolution of the input data with a gridding kernel function. This is equivalent to calculate for each output grid point the weighted sum of relevant input data samples, where the weighting function is the gridding kernel,

$$R_{i,j}[s] = \sum_n R_n[s](\alpha_n, \delta_n) w(\alpha_{i,j}, \delta_{i,j}; \alpha_n, \delta_n). \quad (1)$$

Here, $R_n[s]$ and $R_{i,j}[s]$ are two different representations of the true signal s . The index n runs over the list of all input samples, with respective coordinates (α_n, δ_n) , while the regular output grid can be parametrized via pixel coordinates (i, j) with associated world coordinates $(\alpha_{i,j}, \delta_{i,j})$. The value of the weighting kernel w depends only on the input and output coordinates. In most cases a radially symmetric kernel is applied, such that $w(\alpha_{i,j}, \delta_{i,j}; \alpha_n, \delta_n) = w(d(\alpha_{i,j}, \delta_{i,j}; \alpha_n, \delta_n))$, with the true angular distance d . The true angular distance can be calculated with the Haversine formula or the Vincenty formula⁶. The latter has better numerical accuracy for large angular separations, but is slightly more computationally demanding. For the gridding task, which involves calculation of small distances, the Haversine formula is sufficient.

If one is interested in flux density conservation, one has to account for the potential irregular sampling of the input data; different target pixels could be influenced by a very different number of input data samples. To account for this, we introduce the overall weight map,

$$W_{i,j} \equiv \sum_n w(\alpha_{i,j}, \delta_{i,j}; \alpha_n, \delta_n), \quad (2)$$

such that Eq. (1) becomes

$$R_{i,j}[s] = \frac{1}{W_{i,j}} \sum_n R_n[s](\alpha_n, \delta_n) w(\alpha_{i,j}, \delta_{i,j}; \alpha_n, \delta_n). \quad (3)$$

Even with unnormalized gridding kernels, flux density conservation is then guaranteed.

To compute a complete map $R_{i,j}[s]$, one needs to iterate over all pixels (i, j) and apply Eq. (3) to each of them. In practice, it is easier to perform the outer iteration over n instead because the input data set is usually larger than the output map. The following algorithm is then applied:

1. Set $R_{i,j} \equiv 0$ and $W_{i,j} \equiv 0$.
2. Iterate over all input samples n and each target pixel (i, j) :
 - calculate weight, w , for each coordinate pair $(\alpha_{i,j}, \delta_{i,j}) \leftrightarrow (\alpha_n, \delta_n)$;

Fig. 1. Convolution-based gridding in the world coordinate system has the advantage that even relatively complicated tasks are easy to implement. In the visualized example an elliptical 2D-Gaussian gridding kernel with a bearing of 45° is used for a field close to the pole. Artificial point sources were inserted in intervals of 45° longitude at a latitude of 88.5° . The resulting ellipses in the pixel space are distorted from the chosen projection (zenith-equal-area) of the sphere onto a plane.

- compute $R_n \cdot w$, add the result to $R_{i,j}$, and add w to $W_{i,j}$.
3. Divide $R_{i,j}$ by $W_{i,j}$ pixel-wise to get the final map.

The presented method has several advantages. It is straightforward to implement and incorporate any WCS projection. By calculating the world coordinates of the target pixel grid, the method can convolve in world coordinate space rather than in pixel space. This also facilitates its application in more complicated tasks, for example, the use of elliptical Gaussian kernels in true world coordinates; see Fig. 1.

A disadvantage of the sum in Eq. (3) is that it is not easy to parallelize the code. While one could easily compute partial sums simultaneously, the result of each atom $R_n \cdot w$ must be written into the target map $R_{i,j}$ and weight map $W_{i,j}$. In practice, this can lead to race conditions, when two threads try to write into the same memory cell simultaneously, which requires to employ a locking mechanism at the cost of processing speed. One can circumvent this problem by letting each thread use its own copy of $R_{i,j}$ and $W_{i,j}$, and merging them afterward. However, this would significantly increase the memory consumption for data cubes or very large target maps, even if only a small number of threads is involved.

Furthermore, the workload can greatly be reduced by preselecting input samples in the vicinity of each target pixel: all pixels with negligible weight (for a given input sample) can safely be ignored. For some WCS projections this would be relatively easy, but others have discontinuities, which makes the neighbor search computationally more complex. Without such optimization the convolution had to be calculated for the full target grid size. This would make the algorithm $O(i \cdot j \cdot n)$ instead of $O(n)$, and results in a large number of unnecessary operations if the convolving kernel is much smaller than the size of the map.

⁶ https://en.wikipedia.org/wiki/Great-circle_distance

2.2. Improvement with HEALPix-based lookup tables

We can avoid both of the previously mentioned disadvantages, if we introduce a lookup table based approach. This is best accomplished by internally working on the HEALPix grid. The HEALPix grid is a tessellation of the sphere, where each pixel has the same solid angle but the pixels are generally not rectangular. For information about HEALPix, see [Górski et al. \(2005\)](#).

There are several software libraries related to HEALPix⁷ that provide fast lookup routines to select HEALPix elements based on WCS coordinates. For example, a function is provided to obtain all the HEALPix pixels within a circle of given radius for a certain sky position, the so-called *query-disk* function. This can be used to greatly improve the complexity of the gridded from $O(i \cdot j \cdot n)$ to $O(n)$, i.e., to avoid the problem described in the last paragraph of Sect. 2.1.

In the following, the basic concept of the improved algorithm is presented, while in Sect. 3 we discuss our reference implementation *cygrid* in more detail, along with some techniques to improve computational efficiency.

The main idea is to use a lookup table that contains, for each target pixel (i, j) , the subset of input sample indices $\{n\}_{i,j}$ that will significantly contribute to the convolution. Then, instead of having the outer loop iterate over n , we iterate over the list, $\{\{n\}_{i,j}\}$, of all non-empty $\{n\}_{i,j}$ and their respective elements. As such the outer loop can be computed concurrently because the program writes into different map and weight-map cells in each thread.

To determine the list $\{\{n\}_{i,j}\}$ several steps are necessary as follows:

1. Set the pixel resolution for the internally used HEALPix grid. In the HEALPix scheme, each coordinate pair is represented by a single index, which makes it easy to use the HEALPix index as a hash table key. In Sect. 3.5 we discuss how to find appropriate values.
2. For each target map pixel (m, n) find the nearest HEALPix pixel $h_{i,j}$. In the resulting list $\{h_{i,j}\}$, the HEALPix indices are not necessarily unique because multiple input pixels could be associated with the same HEALPix pixel.
3. Iterate over all input sample indices, n
 - (a) Calculate the associated HEALPix index, h_n .
 - (b) Run the *query-disk* function to get all HEALPix indices $\{h'_n\}$ within a radius of r_d around (α_n, δ_n) . This radius is chosen such that the kernel has effectively zero contribution for larger angular separations.
 - (c) Matching all elements in $\{h'_n\}$ with $\{h_{i,j}\}$, one obtains the pixels (i, j) to which input n contributes. Add this information to the lists $\{\{n\}_{i,j}\}$ accordingly.

The matching operation in step (3c) is where the lookup table approach comes into play. If $\{h_{i,j}\}$ is organized as a hash table that maps HEALPix indices to (lists of) target pixels (i, j) , then matching is just a lookup operation; if a HEALPix index from $\{h'_n\}$ is in the hash table, the associated value is used to link n to (i, j) .

3. Cython implementation of the gridding algorithm: Cygrid

3.1. Using Cython

We provide a Cython-based implementation of the algorithm outlined above. Cython is a tool that facilitates writing

Python-like code, which is then precompiled into C code. Cython can also be employed to wrap C/C++ libraries or source code to be used inside Python. In fact, we heavily use container classes (vectors and maps⁸) from the C++ standard template library (STL) for *cygrid*.

3.2. HEALPix functions

Only a small subset of HEALPix routines is necessary for *cygrid*, mainly the functions to convert between world coordinates and HEALPix index. Instead of relying on the reference implementation of HEALPix, we reimplemented the required functionality in *cygrid* because it avoids the additional dependency on an external library. Furthermore, this allows us to add several tweaks for our purpose, for example, to the *query-disk* routine.

3.3. A fast query-disk routine

The largest benefit in terms of efficiency comes from performing the computation of the convolution only for pairs of input sample coordinates and output pixels that yield a significant weight factor. To find these pairs, we employ the *query-disk* function, which returns the indices of surrounding pixels for any HEALPix input pixel index, located inside a sphere of radius, r_d . The gridding kernel is then only calculated for samples within the queried region with r_d chosen such that the values of the kernel function are practically zero outside the area enclosed by the sphere.

For each input world coordinate pair we first compute its HEALPix index. We only need to do the disk query once per HEALPix index instead of once per input world coordinate (number of involved HEALPix indices $\leq$ number of input coordinate pairs). Furthermore, we store the disk pixel lists in a hash table. Unless the input data is only sparsely sampled, this saves a fair amount of computing power because a hash table lookup is usually faster than the *query-disk* procedure.

To increase efficiency even further, we run the *query-disk* function only once per HEALPix ring, for the longitude $\alpha = 180^\circ$. By exploiting the regularity of the HEALPix grid, one can move the disks along longitude. For each disk pixel index (in RING scheme), p , we first calculate the ring number, r , and the number of HEALPix pixels in this ring, n_r . In the equatorial belt, n_r is the same for each ring, while in the polar regions n_r gets smaller toward the poles (see [Górski et al. 2005](#), for details). The (integer) shift, Δp , that needs to be applied to shift a pixel by an angle of $\Delta\theta$ can be calculated via

$$\Delta p = \frac{\Delta\theta}{2\pi} n_r, \quad (4)$$

where Δp is rounded to the nearest integer. If the index of the first pixel per ring is denoted as s_r , the new (shifted) pixel index is then given by

$$\tilde{p} = s_r + [(p - s_r + \Delta p) \bmod n_r]. \quad (5)$$

3.4. Determine $\{\{n\}_{i,j}\}$

We now briefly discuss how the lookup table approach works. The *grid* method of *cygrid* calls a function *calculate_output_pixels*, with the coordinates of the

⁷ e.g., <http://healpix.sourceforge.net/>

⁸ In C++/STL a hash table is denoted as *map*.

input samples as parameter. This function then runs the following two subroutines:

- `compute_input_output_mapping`: this routine is straightforward. (1) Compute the HEALPix indices and ring indices for the provided input coordinates. (2) For each ring index, run the `query_disk` routine (for $\alpha = 180^\circ$) and store in a hash table (or map, in C++ nomenclature). If an entry is already present in the hash table, the `query_disk` is not be called again. For the typical case, this approach saves a huge amount of `query_disk` calls, as the number of HEALPix rings involved is significantly smaller than the number of input coordinates. (3) For each input HEALPix index, now find the HEALPix indices of the surrounding disk, by shifting the $\alpha = 180^\circ$ disks as described in Sect. 3.3. As a result, we now have a hash table `input_output_mapping`, which contains a list of target (WCS) pixels, associated to each input HEALPix index, for which the convolution has to be performed.
- `compute_output_input_mapping`: as a last step, we have to invert `input_output_mapping`, i.e., calculate an `output_input_mapping`, which contains the list of relevant input HEALPix indices for each target (WCS) pixel.

Many of the steps in the `calculate_output_pixels` function can also be parallelized, which accelerates processing when using multiple CPUs.

3.5. Using cygrid: a minimal example

In Listing 1 we show a minimal example of how a grid job could be programmed. After reading in the data in line 5 and the definition of a target-map FITS header (line 8) the user needs to create an instance of the `WcsGrid` class (line 31) and set an appropriate kernel (line 32). The `WcsGrid` class is an implementation of the `Cygrid` base class. The latter cannot be used directly, but only the derived classes work. Currently, we provide two implementations: `WcsGrid` and `SlGrid`. `WcsGrid` can be used to grid onto any of the WCSlib projections and must be provided with a FITS header during object construction (either as `pyfits` Header object or as Python dictionary). The `SlGrid` can be used to resample onto a list of coordinates (“sight lines”).

In the call to the `set_kernel` class method, we provide the desired kernel-function type and the associated parameters. Furthermore, the `support-radius` and `hpx-max-resolution` have to be set. The `support-radius` defines the radius around each input coordinate sample out to which the convolution has to be performed. The chosen kernel function should have negligible contribution beyond the given radius. For example, for a radial-symmetric Gaussian kernel with standard deviation σ , the `support-radius` could be chosen to lie between 3σ and 5σ depending on the desired accuracy. The other parameter `hpx-max-resolution` defines the resolution of the internal HEALPix grid. All input coordinate samples and the output (WCS) map pixels (or sight lines) are internally associated with a HEALPix index to allow for optimal lookup table handling. For `cygrid` it is no issue if multiple input samples or output coordinate pairs share HEALPix indices, so in principle a very coarse HEALPix grid could be used. However, a very coarse HEALPix grid would mean that the `query-disk` function yield a frazzled disk. If the `support-radius` is sufficiently large, this is not a problem, however, we recommend setting `hpx-max-resolution` to about $\sigma/2$ if the disk radius is only set to 3σ .

Also, the properties of the kernel and output grid should be compatible. To avoid aliasing effects, the output grid should be

Listing 1. Basic Cygrid usage.

```

1 from astropy.io import fits
2 import cygrid
3
4 # read-in data
5 glon, glat, signal = get_data(...)
6
7 # define target FITS/WCS header
8 header = {
9     'NAXIS': 3,
10    'NAXIS1': 101,
11    'NAXIS2': 101,
12    'NAXIS3': 1024,
13    'CTYPE1': 'GLON-SFL',
14    'CTYPE2': 'GLAT-SFL',
15    'CDELT1': -0.1,
16    'CDELT2': 0.1,
17    'CRPIX1': 51,
18    'CRPIX2': 51,
19    'CRVAL1': 12.345,
20    'CRVAL2': 3.14,
21 }
22
23 # prepare gridder
24 kernelsize_sigma = 0.2
25
26 kernel_type = 'gauss1d'
27 kernel_params = (kernelsize_sigma, )
28 kernel_support = 3 * kernelsize_sigma
29 hpx_maxres = kernelsize_sigma / 2
30
31 mygridder = cygrid.WcsGrid(header)
32 mygridder.set_kernel(
33     kernel_type,
34     kernel_params,
35     kernel_support,
36     hpx_maxres
37 )
38
39 # do the gridding
40 mygridder.grid(glon, glat, signal)
41
42 # query result and store to disk
43 data_cube = mygridder.get_datacube()
44 fits.writeto(
45     'example.fits',
46     header=header, data=data_cube
47 )

```

fine enough such that not only the input data but also the kernel function is fully sampled. Furthermore, unless the input data is not heavily oversampled, we advise choosing a kernel size that is not much smaller than the original resolution of the input data. Of course, the convolution-based gridding lowers the spatial resolution, σ_{data} , of the gridded data somewhat. In typical scenarios, good results can be achieved with $\sigma_{\text{kernel}} \approx 0.5\sigma_{\text{data}}$, such that the final resolution is

$$\sigma_{\text{data}}^{\text{gridded}} = \sqrt{\sigma_{\text{kernel}}^2 + \sigma_{\text{data}}^2} \approx 1.12\sigma_{\text{data}}, \quad (6)$$

which is only slightly worse than the original resolution. A more detailed analysis on how to choose a well-suited kernel size is presented in Appendix A.

Finally, after starting the grid job in line 40, we can retrieve the data cube (line 43) and write to a FITS file (line 44).

3.6. Source code repository

Cygrid is open-source software and is made available on the GitHub⁹ platform. Community contributions are welcome. We also prepared several Jupyter Notebooks¹⁰ that demonstrate how to use cygrid in practice as follows:

- *Minimal*. In this notebook, we show a minimal example of how to use cygrid, similar to Listing 1 but fully working. Moreover, we illustrate some applications of the `wcs` module from AstroPy¹¹ in combination with cygrid.
- *Two fields*. This notebook explains how to grid data from one FITS-image coordinate projection to another, using H I and far-infrared data.
- *Healpix allsky*. We demonstrate how to use the `healpy` library with cygrid and how one can grid full-sky data sets onto smaller FITS images.
- *Sightlines*. Cygrid is used to grid on individual lines of sight instead of WCS fields.

All these notebooks are also available on GitHub.

4. Benchmarking cygrid

In this short section, we analyze the performance of our cygrid implementation. For this, we produce random data; for a specified target field coordinate pairs (longitude and latitude) are sampled from a uniform distribution, and each is assigned a signal value (Gaussian noise). We then use cygrid to grid the data and measure the necessary computing time.

The result is displayed in Fig. 2 (top panel), which shows the processing times as a function of number of input samples for a field size of $5^\circ \times 5^\circ$, a pixel size of $200''$, and a Gaussian kernel width of $\theta_{\text{fwhm}} = 300''$.

Next, we checked the quality of the parallelization. For this, we ran the test program with a different number of processor cores, from 2 to 16. In the lower panel of Fig. 2 we plot the relative performance gain of each run relative to single-threaded processing times. While a second core means almost 100% speed-up, for a larger number of utilized cores the speed-up is less (per core). This reflects the nature of the gridding; it is not purely CPU bound, but is strongly dependent on the available memory bandwidth and cache performance. Still, on our workstation, for 10^8 input samples, the 16-core processing time is just about 67 s.

For comparison, we carried out the same gridding task using the `scipy.interpolate.griddata` function¹², which provides nearest-neighbor and linear and cubic spline interpolation. The resulting runtimes are shown as black curves in Fig. 2. For the spline interpolation, SciPy uses the QHull triangulation library¹³, which apparently limits the number of samples to about 16.7 million. Therefore, only the nearest-neighbor interpolation was carried out up to the maximum of 10^8 samples. For many input samples, the spline interpolation is about an order of magnitude slower than cygrid, while the much simpler nearest neighbor algorithm is about twice as fast (compared to a single-threaded cygrid). The `griddata` function is not the best choice for data gridding, but is aimed at resampling data sets. For cases with very dense sampling of the input data (i.e., when the gridding operation is effectively downsampling a data set) the RMS

Fig. 2. Example 1: Benchmark results for a small test program, which grids a different number of random samples to a target field of $5^\circ \times 5^\circ$ size. The width of the gridding kernel was $\theta_{\text{fwhm}} = 300''$; the pixel size was $200''$.

Fig. 3. Asymptotic behavior of cygrid's runtimes for constant field size. The time complexity (black solid line) appears to lie within $O(n)$ and $O(n \log n)$ (blue shaded area), where n is the number of input samples.

noise in the target map is not decreased. We demonstrate this effect in the first of the Jupyter notebooks; see Sect. 3.6. Therefore, we did the comparison with SciPy with respect to the runtimes but not the resulting images.

The asymptotic behavior in Fig. 2 suggests a computational complexity, which is close to linear, i.e., $O(n)$. To visualize this better, we plot in Fig. 3 the 1-core curve (solid black line) from

⁹ <https://github.com/bwinkel/cygrid>

¹⁰ <http://jupyter.org/>

¹¹ <http://www.astropy.org/>

¹² <http://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.griddata.html>

¹³ <http://www.qhull.org/>

Fig. 4. Example 2: as in Fig. 2 but for varying field sizes of 0.1° to 60° edge length. The spatial pixel density was kept constant at 100 000 samples per square degree.

the example along with a blue shaded area that indicates the $O(n)$ and $O(n \log n)$ curve complexities. The latter are scaled such that they yield the same processing time for 10^6 samples and enclose a solid curve over most of the asymptotic regime. The measured processing times are somewhat noisy, but it can be seen that the behavior is indeed compatible with $O(n)$ and $O(n \log n)$ behavior. The fact that the complexity is a little worse than $O(n)$ is probably caused by overheads in the hash table lookup for a huge number of entries.

We carried out a second test where we kept the spatial sample density constant and varied the field size from about 0.1° to 60° (using a spatial density that yields the same number of samples as in the first test for better comparison). In Fig. 4 we show the results in the same way as for Fig. 2. The pixel size and kernel width was the same as for the first test. Compared to the first benchmark, cygrid is about a factor of two slower. We attribute this to the increased number of query-disk calls and to the much larger disk hash table, which causes slower lookup times. We compare the runtimes with SciPy's `griddata` function, and cygrid is still competitive (factor of 2 slower with 1-core) or much faster (multiple cores).

5. Summary

We have presented cygrid, a fast and flexible data gridding module for the Python programming language. Cygrid implements a convolution-based algorithm and, to maximize its performance, makes use of hash tables, a fast query-disk routine, and the Cython-extension.

Cygrid can be used to resample onto any of the FITS/WCS projections to easily create astronomical maps or data cubes. It is also possible to grid to a list of coordinates (sight lines), which can be used, for example, to produce maps on the HEALPix grid.

It was demonstrated that our cygrid implementation benefits from multicore CPUs (with relatively good scaling behavior) and easily outperforms other widely used gridding algorithms, such as the linear and cubic spline interpolation, which ship with SciPy. Cygrid's runtimes to grid a 2D map lie between $O(n)$ and $O(n \log n)$ complexity, i.e., the overheads compared to a linear behavior, are very moderate. For data cubes the impact of the overheads is even smaller because all samples of one spectrum are multiplied with the same weight factor.

Acknowledgements. We are grateful to Alexander Kraus for carefully proofreading the manuscript and for his valuable comments. Some of the results in this paper were derived using the HEALPix (Górski et al. 2005) package. Cygrid was developed in the framework of the Effelsberg-Bonn H I Survey (EBHIS) project. EBHIS is based on observations with the 100-m telescope of the MPIfR (Max-Planck-Institut für Radioastronomie) at Effelsberg. The authors thank the Deutsche Forschungsgemeinschaft (DFG) for support under grant numbers KE757/7-1, KE757/7-2, KE757/7-3, and KE757/9-1. B.W. was partially funded by the International Max Planck Research School for Astronomy and Astrophysics at the Universities of Bonn and Cologne (IMPRS Bonn/Cologne). L.F. was also a member of IMPRS Bonn/Cologne. D.L. is a member of the Bonn-Cologne Graduate School of Physics and Astronomy (BCGS). We would like to express our gratitude to the developers of the many C/C++ and Python libraries, which are made available as open-source software and we used: most importantly, NumPy (van der Walt et al. 2011) and SciPy (Jones et al. 2001), Cython (Behnel et al. 2011), and Astropy (Astropy Collaboration et al. 2013). Figures were prepared using matplotlib (Hunter 2007) and in part using the Kapteyn Package (Terlouw & Vogelaar 2015).

References

- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *Comput. Sci. Eng.*, **13**, 31
- Calabretta, M. R., & Greisen, E. W. 2002, *A&A*, **395**, 1077
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Greisen, E. W., & Calabretta, M. R. 2002, *A&A*, **395**, 1061
- Hunter, J. 2007, *Comput. Sci. Eng.*, **9**, 90
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python [Online; accessed 2015-06-27]
- Kalberla, P. M. W., Burton, W. B., Hartmann, D., et al. 2005, *A&A*, **440**, 775
- Kalberla, P. M. W., McClure-Griffiths, N. M., Pisano, D. J., et al. 2010, *A&A*, **521**, A17
- Kerp, J., Winkel, B., Ben Bekhti, N., Flöer, L., & Kalberla, P. M. W. 2011, *Astron. Nachr.*, **332**, 637
- Krige, D. G. 1951, *Journal of the Chemical, Metallurgical and Mining Society of South Africa*, **52**, 119
- Matheron, G. 1963, Principles of geostatistics, Vol. 58 (Society of Economic Geologists), 1246
- McClure-Griffiths, N. M., Pisano, D. J., Calabretta, M. R., et al. 2009, *ApJS*, **181**, 398
- Terlouw, J. P., & Vogelaar, M. G. R. 2015, Kapteyn Package, version 2.3, Kapteyn Astronomical Institute, Groningen, <http://www.astro.rug.nl/software/kapteyn/>
- van der Walt, S., Colbert, S., & Varoquaux, G. 2011, *Comput. Sci. Eng.*, **13**, 22
- Winkel, B., Kerp, J., Flöer, L., et al. 2016, *A&A*, **585**, A41

Appendix A: The gridding kernel revisited

To use cygrid, one needs to choose a gridding kernel. The choice of kernel type is usually application dependent, but the kernel size and kernel support radius must be chosen well. In the following, several aspects are discussed to guide the potential user. We restrict the analysis to the case of a (spherical-symmetric) Gaussian kernel for the sake of simplicity.

A.1. Image resolution

Any convolution-based gridding algorithm degrades the resolution of the imaged signal at least somewhat. The exception to this is the $\sin(ax)/ax$ kernel, which can reconstruct the original input signal if computed with (effectively) infinite support. From Eq. (6) it is obvious that the smaller the convolution kernel, the better the resolution after gridding. A very subtle effect, however, lies in the potential different correlation lengths of the input signal and measurement noise. In many astronomical applications, the instrument itself adds noise to the measured signal, for example, produced by a CCD detector or due to the thermal noise component in a radio receiver. There are even cases, when the noise is effectively independent in each sample (e.g., for a single-dish radio observation), while the desired astronomical signal is convolved with the beam or point spread function (PSF) of the instrument and as such is correlated in the spatial domain.

To visualize this we run a simulation, gridding artificial data to an all-sky map. To avoid map projection effects, we grid to the HEALPix $n_{\text{side}}=512$ grid pixels, $(l_{\text{hpx}}, b_{\text{hpx}})$, (pixel size: $\vartheta_{\text{hpx}} = 6.87^\circ$) using cygrid's sightline griddler (S1Grid). The artificial signal is comprised of 1024 point sources that are randomly distributed on the sphere with different intensities. For a realistic scenario, the effect of an observing instrument needs to be considered. Therefore, we also randomly sample about 50 million¹⁴ observing positions $(l_{\text{in}}, b_{\text{in}})$ that are uniformly distributed over the sphere. A Gaussian-type PSF with $\vartheta_{\text{psf}} = 4\vartheta_{\text{hpx}} = 27.5^\circ$ (FWHM; $\vartheta_{\text{fwhm}} = \sqrt{8 \ln 2} \sigma \approx 2.35\sigma$) is then applied to calculate the “measured” signal, s_{in} , on each of the input coordinate samples $(l_{\text{in}}, b_{\text{in}})$. For each of the s_{in} samples we draw a noise sample, n_{in} , from a normal distribution. The simulated data points, $s_{\text{in}} + n_{\text{in}}$, are then gridded using a kernel of $\vartheta_{\text{kernel}} = \vartheta_{\text{psf}}/2 = 13.7^\circ$ in size (FWHM). The resulting image is shown in Fig. A.1.

Using the healpy routine `anafast` one can easily calculate the power spectral density (PSD) of such a signal on the sphere. In Fig. A.2 the gray solid curve indicates the PSD returned by `anafast` for the signal plus noise after gridding. We also processed the input signal and noise separately (light green and light blue solid lines, respectively).

To study the effect of the gridding itself, a comparison with the true signal is helpful. Therefore, we calculate the PSF-convolved representation of the true signal on the HEALPix pixel centers directly (green solid line in Fig. A.2). As before, independent noise samples are drawn (blue solid line). For consistency with the gridded noise values, a scaling needs to be applied to account for the larger number of samples in the gridding case. The combined theoretical signal plus noise is shown with a black solid line.

Both the gridded signal and gridded noise have less power on the small angular scales, compared to the true signal and noise spectral densities. This is because the convolution smooths the

Fig. A.1. A simulation is used to study the impact of the gridding on the power spectral densities of the input signal and noise. We randomly placed 1024 point sources with different intensities on the sphere. To generate the “measured” input signal, the true point-source signal is convolved with a PSF of $\vartheta_{\text{psf}} = 27.5^\circ$ and sampled on about 50 million random coordinate pairs. Independent Gaussian noise is added to each resulting signal sample. Using a Gaussian-shaped gridding kernel with $\vartheta_{\text{kernel}} = 13.7^\circ$ the input signal and noise are then imaged to the $n_{\text{side}}=512$ HEALPix grid.

Fig. A.2. Power spectral densities as calculated using the healpy routine `anafast` for the simulations. See text for details.

data on small scales. The true input signal (but not the noise) already shows a drop-off in the spectral power toward small angular scales because the point sources were convolved with a Gaussian PSF. One important aspect to keep in mind is that, as shown, the resulting spatial resolution (or rather correlation-length) of the noise after gridding can be different from the resulting spatial resolution of the signal. Also, cygrid conserves the total flux density in the image, which, for example, means that the amplitude of a (Gaussian) peak decreases while its width increases during convolution with the kernel.

A.2. Aliasing and choice of kernel size

The above discussion does not allow us to come to a conclusion on the choice of an optimal kernel size. Why can we not use

¹⁴ This is 15 times the number HEALPix grid pixels, such that on average 15 samples contribute to each pixel in the gridded map.

an arbitrarily small kernel, for example, which would avoid the degradation of the spatial resolution? In fact, the necessary kernel size is entirely determined by the sampling density of the input data point and the required accuracy. To demonstrate this, we set up a different simulation. The input coordinates are now placed on horizontal stripes (scan lines), with very dense sampling along the stripe direction, but with a certain spacing between the stripes. Such a scenario is common for mapping observations with a single-dish radio telescope with just one feed horn. It also effectively converts the question about the proper kernel size to a one-dimensional problem, which is easier to analyze. As before, we construct an artificial signal from point sources, convolved with ϑ_{psf} , but here we use a much higher source density to obtain a quasi-continuous field. The noise component is omitted this time because we want to analyze potentially small effects in the image reconstruction. We use four different stripe spacings, $d_{\text{space}} = [0.2, 0.3, 0.4, 0.5] \times \vartheta_{\text{psf}}$ and grid each of the resulting input signals with five different kernel sizes, $\vartheta_{\text{kernel}} = [0.3, 0.4, 0.5, 0.6, 0.7] \times \vartheta_{\text{psf}}$. The field size is $1.5^\circ \times 1.5^\circ$ and $\vartheta_{\text{psf}} = 0.1^\circ$ (FWHM), but as the spacing and kernel sizes are relative to the beam sizes this is unimportant for the results. We chose the pixel size of the target grid to be $\vartheta_{\text{pix}} = \vartheta_{\text{psf}}/20$, which is heavily oversampled, but useful for revealing all of the potential effects. For the same reason, the support radius was made large with $5\sigma_{\text{kernel}}$. Figure A.3 shows the resulting maps. In the top left of each panel the kernel size (FWHM, black circle), the beam size (FWHM, gray circle), and the resulting (theoretical) image resolution (FWHM, white circle) are visualized. The tick marks on the vertical axes depict the

position of the horizontal stripes on which the input samples are located.

In the top right panel of Fig. A.3 ($d_{\text{space}} = 0.5\vartheta_{\text{psf}}$, $\vartheta_{\text{kernel}} = 0.3\vartheta_{\text{psf}}$) block-like artifacts are clearly visible, which hint at a problem with the imaging (so-called aliasing). For a more detailed analysis, we show in Fig. A.4 the relative deviation of the gridded image from the input signal. The latter was also convolved with the gridding kernel, otherwise the different spatial resolutions would not facilitate a meaningful comparison. As the deviation maps have extremely different scales, we multiplied each panel with a scale factor. The number in the lower left of each panel indicates the minimal and maximal deviation for each case associated with the dark blue and red colors in the image, respectively.

For the demonstrated setup with separated scan lines we could now choose the kernel size as a compromise between accuracy and acceptable degradation of image resolution. Of course, if a very different data point sampling or different kernel type was used, a similar analysis should be employed to find well-suited kernel parameters. As a rule of thumb, a kernel size of $0.5\text{--}0.6\vartheta_{\text{psf}}$ probably provides results accurate to better than 1–2% even for relatively bad sampling densities. We also note, however, that the choice of the kernel support radius has an impact on the maximal accuracy. To demonstrate this, we repeated the simulation with a smaller support radius of $3\sigma_{\text{kernel}}$; see Fig. A.5. Comparing the two figures reveals that the support radius mainly impacts the high-accuracy cases, while the low-precision cases are not affected much.

Fig. A.3. Impact of different sampling densities and kernel sizes on the accuracy of the gridding algorithm. Input samples are distributed on horizontal stripes (scan lines), separated with the given spacing (in units of the beam size). The locations of the stripes are indicated by the tick marks on the vertical axes. Various kernel sizes (also in units of beam size) were applied to image the artificial signal. In the top left of each panel the kernel size (black circle), beam size (gray circle), and resulting (theoretical) image resolution (white circle) are visualized. See text for further details.

Fig. A.4. As Fig. A.3 but showing the relative deviation from the theoretical signal (convolved with the gridding kernel). See text for further details.

Fig. A.5. As Fig. A.4 but with a smaller kernel support radius of $3\sigma_{\text{kernel}}$. See text for further details.