

Fine-grained distributed averaging for large-scale radio interferometric measurement sets

Shou-Lin Wei¹, Kai-Da Luo¹, Feng Wang^{1,2}, Hui Deng² and Ying Mei²

¹ Faculty of Information Engineering and Automation, Computer Technology Application Key Lab of Yunnan Province, Kunming University of Science and Technology, Kunming 650500, China

² Center for Astrophysics, Guangzhou University, Guangzhou 510006, China; fengwang@gzhu.edu.cn

Received 2020 March 19; accepted 2020 September 22

Abstract The Square Kilometre Array (SKA) would be the world's largest radio telescope with eventually over a square kilometre of collecting area. However, there are enormous challenges in its data processing. The use of modern distributed computing techniques to solve the problem of massive data processing in the SKA is one of the most important challenges. In this study, basing on the *Dask* distribution computational framework, and taking the visibility function integral processing as an example, we adopt a multi-level parallelism method to implement distributed averaging over time and channel. *Dask Array* was used to implement super large matrix or arrays with supported parallelism. To maximize the usage of memory, we further exploit the data parallelism provided by *Dask* that intelligently distributes the computational load across a network of computer agents and has a built-in fault tolerance mechanism. The validity of the proposed pattern was also verified by using the Common Astronomy Software Application (CASA), wherein we analyze the smearing effects on images reconstructed from different resolution visibilities.

Key words: techniques: interferometric — methods: data analysis — methods: numerical — instrumentation: interferometers

1 INTRODUCTION

The primary reason for constructing the Square Kilometre Array (SKA) (Dewdney et al. 2009) that comprises many individual stations is to explore the universe with the highest resolution, sensitivity and dynamic range. However, the telescope brings new scientific and technical challenges to the modern Information Technology. One of the biggest challenges is managing the extremely large data volumes generated by the SKA every day. The narrow channel bandwidth and short correlator timestep do minimize the decorrelation effects on the longest baselines but increase the amount of visibility data. This substantial increase in the bulk of data requires the development of new and sophisticated distributed data processing schemes and concurrent algorithms that can run on high-performance computing hardware.

To prompt the development of Science Data and Handling Processor (SDHP), the Radio Astronomy Simulation, Calibration and Imaging Library (RASCIL) (Cornwell et al. 2020), including radio interferometry calibration and imaging algorithms in Python and NumPy, has been developing for a few years. The interfaces

all operate with familiar data structures such as image, visibility table, gain table, and so on. To achieve sufficient performance, the latest version of the RASCIL takes a dual pronged approach – using threaded libraries for shared memory processing, and the *Dask* (Rocklin 2015) library for distributed processing.

Dask is a flexible library for parallel computing in Python that is composed of two parts, i.e., dynamic task scheduling and big data collection. In the area of interferometric data processing, *dask-ms* (Perkins et al. 2020) and *cngi-prototype* (Raba et al. 2019) exploited *Dask* arrays to expose columns of tables in measurement sets (MS) (van Diepen 2015) to support parallel operations on MS. Current RASCIL software takes full advantage of *Dask*'s dynamic scheduling capabilities. However, the big data collection feature has never been tested in RASCIL development. Obviously, distributed computing is the unique choice for processing large-scale astronomical data. The visibility has time and bandwidth separation, making distributed computing a natural solution.

Averaging in post-correlation is considered to reduce the volume of astronomical data and computational cost

(Offringa 2016). To further investigate the ability of *Dask* on its big data processing, we take the visibility function data averaging (integral) as the case and developed a fine-grained distributed averaging algorithms including time and channel averaging so as to test whether *Dask* big data collection feature can be used in the subsequent version of the RASCIL.

The rest of this paper is organized as follows: Section 2 introduces the theoretical analysis of averaging; Section 3 describes the algorithm implementation and validation results; Section 4 analyzes the performance evaluation in terms of granularity; and Section 5 presents the conclusion. The code that replicates the results in this work is available online¹.

2 RADIO INTERFEROMETRIC AVERAGING

This section provides a brief overview of radio interferometry, reviews the theory of averaging, and discusses the effects of averaging on interferometric data processing.

The response of interferometric arrays to the brightness distribution of sources in the sky is measurements called *visibilities*, which are complex values sampled from the Fourier transform of the sky brightness distribution (Thompson 1999). The relation between the synthesis array, visibilities, and sky brightness is well established, and it can be expressed by the radio interferometer measurement equation (RIME) (Hamaker et al. 1996). The RIME has provided a rigorous mathematical framework to model the signal propagation effects and calculate the visibilities. We assume that a baseline is composed of two spatially separated stations p and q . The visibility matrix \mathbf{V}_{pq} is calculated as an integral product over a time and frequency intervals $[t_0, t_1] \times [\nu_0, \nu_1]$. It can be expressed as

$$\langle \mathbf{V}_{pq} \rangle = \lim_{\Delta t, \Delta \nu \rightarrow 0} \frac{1}{\Delta t \Delta \nu} \int_{t_0}^{t_1} \int_{\nu_0}^{\nu_1} \mathbf{J}_p(t, \nu) \mathbf{B} \mathbf{J}_q^H(t, \nu) d\nu dt. \quad (1)$$

The angled brackets denote data averaging over one baseline, frequency channel, and integration timestep. In addition, H represents the conjugate transpose. Matrix \mathbf{B} describes the continuous brightness distribution, and \mathbf{J} describes the Jones chain modeling effects along the signal propagation path (Smirnov 2011).

2.1 Channel Averaging

Modern radio telescopes typically make observations across multiple channels, each of which is bound to a specific frequency. Channel averaging is used to integrate

the visibilities across several continuous or discrete channels. To keep both sides of the equation's form consistent, the weighted average visibility across channels can be expressed as

$$\mathbf{V}_{pq}^{\text{avg}} \mathbf{W}_{pq}^{\text{avg}} = \sum_i^{nch} \mathbf{V}_{pqi} \mathbf{W}_{pqi}. \quad (2)$$

The weight of the broadened channel is the sum of the weights of the integral channels. All channels have the same weight at a specific time and baseline. Thus,

$$\mathbf{W}_{pq}^{\text{avg}} = \sum_i^{nch} \mathbf{W}_{pqi} = nch \times \mathbf{W}_{pqi}, \quad (3)$$

where i is the index of the frequency channel, nch is the number of integral channels, and avg denotes the channel averaging operation. In addition, \mathbf{V} , \mathbf{W} represent the columns of DATA and WEIGHT columns in the MS, respectively (as shown in Table 1). For brevity, the angle brackets of \mathbf{V} in Equation (2) are omitted. Substituting Equation (3) into Equation (2) gives the integral visibilities over the nch channels finally as the weighted average

$$\mathbf{V}_{pq}^{\text{avg}} = \frac{\sum_i^{nch} \mathbf{V}_{pqi} \mathbf{W}_{pqi}}{\Delta\theta + nch \times \mathbf{W}_{pq}^{\text{avg}}}. \quad (4)$$

To avoid division by zero, we introduce a value closer to zero $\Delta\theta$ in Equation (4). The effects of $\Delta\theta$ are discussed in Section 3.3.

2.2 Time Averaging

Assume that t time points are sampled in an averaging period. In terms of the averaging time, these t sampled points will be integrated into one time point that includes the visibilities of $n(n-1)/2$ baselines. In other words, the number of visibilities is reduced from $t \times n \times (n-1)/2$ to $n(n-1)/2$. The averaged weights and visibilities are given by

$$\mathbf{W}_{pq}^{\text{avg}} = \sum_j^t \mathbf{W}_{pqj}, \quad (5)$$

$$\mathbf{V}_{pq}^{\text{avg}} = \frac{\sum_j^t \mathbf{V}_{pqj} \mathbf{W}_{pqj}}{\Delta\theta + \mathbf{W}_{pq}^{\text{avg}}}. \quad (6)$$

Here, j is the index of sampled time point. As shown in Equation (5) and Equation (6), the averaged weights and visibilities for time averaging have a similar structure to those for channel averaging. The difference is that time averaging integrates data with the same baseline along the time axis. As time averaging combines different time points into one, the time-dependent data also needs to be

¹ https://github.com/astroitlab/vis_averaging

Table 1 MS Sectional Layout Referenced in This Work

Columns	Types Dimensional transformation	Comments
UVW	double $(N_r, 3) \rightarrow (N_t, N_b, 3)$	<i>uvw</i> coordinates units in meters
DATA_DESC.ID	int $(N_r,)$	Data description identifier (≥ 0)
FIELD.ID	int $(N_r,)$	Field identifier (≥ 0)
WEIGHT	float $(N_r, N_c) \rightarrow (N_t, N_b, N_c)$	Weight for whole data matrix
DATA	complex $(N_r, N_f, N_c) \rightarrow (N_t, N_b, N_f, N_c)$	Complex visibilities matrix

N_r = number of rows in MS table, N_f = number of frequency channels, N_c = number of correlators N_t = number of timesteps, and N_b = number of baselines. Dimension refers to the shape of the entire specified column rather than that of a single cell.

recalculated. The time-dependent data including U, V, W and the new weighted Time are given by

$$\left(\mathbf{U} = \frac{\sum_j^t U_j \overline{\mathbf{W}}_j}{\sum_j^t \overline{\mathbf{W}}_j}, \mathbf{V} = \frac{\sum_j^t V_j \overline{\mathbf{W}}_j}{\sum_j^t \overline{\mathbf{W}}_j}, \mathbf{W} = \frac{\sum_j^t W_j \overline{\mathbf{W}}_j}{\sum_j^t \overline{\mathbf{W}}_j} \right), \quad (7)$$

$$\text{Time} = \frac{\sum_j^t T_j \overline{\mathbf{W}}_j}{\sum_j^t \overline{\mathbf{W}}_j}. \quad (8)$$

Here, $\overline{\mathbf{W}}_j$ represents the mean weight across all of the baselines at one time point.

2.3 Averaging Effects

The measurements at each baseline inevitably contain uncorrelated Gaussian noise that creates uncertainty in the real and imaginary parts of the visibilities. This uncertainty is mainly represented by signal fluctuation. Equation (4) and Equation (6) are essentially weighted averages, where the smoothing effect can suppress the fluctuation of visibilities. Moreover, multiple consecutive channels or timesteps are integrated into a broader channel or period that can significantly reduce the volume of visibilities. However, this form of data compression unavoidably results in a net loss of amplitude, known as smearing (or time/bandwidth decorrelation in the *uv*-plane), which mainly manifests a decrease in the amplitude toward off-center sources. The amplitude reduction in source flux increases with baseline length and distance from phase center (Wijnholds et al. 2018). Smearing also has detrimental effects on sensitivity, as the noise amplitude does *not* decrease. If we consider the entire *uv*-plane, smearing distorts the point spread function (PSF) and produces relatively high sidelobes (second column in Fig. 3), as the weighted average results in coarser time/frequency bins or lower resolution re-samples. Furthermore, the wide field effects of smearing are of particular concern in the context of modern radio interferometers, and in this case, some bright sources far from the FOV that are modulated by PSF sidelobes can contribute global background noise and artifacts to the image (Smirnov et al. 2012).

3 IMPLEMENTATION

This section describes the architecture of the proposed method and illustrates the principles of different granularity distributions. The accuracy of the proposed method is explained in comparison to those using the common astronomy software application (CASA) (McMullin et al. 2007) software.

3.1 Architecture

The key goal of this study is to allow averaging to run in a parallel and distributed manner. *Dask* was selected for this task owing to its flexibility and limited invasion into the existing codes. In addition, *Dask* provides much of the NumPy application programming interface that users are familiar with, which means simple NumPy-style iterative optimization algorithms can be written that will leverage the built-in parallel data model. An averaging computation across 20 MSs is visualized, as shown in Figure 1, and generated using *Dask*'s built-in visualization features.

To provide clarity with respect to the various levels of parallelism, Figure 1 was abstracted into Figure 2. In accordance with the requirements of computing resources, various logical parts of the program are arranged across different executing spaces, as shown in Figure 2.

Visibilities and related configurations are commonly stored as an MS in a folder. The process starts with the MS loading from a shared file system mounted to the same path in each node. Multiple MS files are loaded in parallel cross nodes. In a single node, an MS is loaded and split by FIELD.ID and DATA_DESC.ID in parallel within one process, where each *dask* worker gets its own interpreter and memory space for computation. Python global interpreter lock (GIL) only allows one thread to control the Python interpreter at a time, which can be a bottleneck that prevents full multi-core usage. Therefore, the multi-process can bypass the GIL to allow multiple cores to execute in parallel.

Throughout its architecture, this algorithm adopts the map/reduce strategy for data analysis. In the map phase, each MS is secondary split by source. In particular, UVW (Table 1) is not converted from meters to wavelengths, otherwise the UVW volume would increase. Multi-




Fig. 1 Dask task graph for averaging computation across 20 measurement sets.

level parallelism and parallel multi-dimensional arrays (discussed in Sect. 3.2) can achieve fine granularity of parallel averaging on visibilities. In the reduction process, all averaged visibilities are merged based on the source attribute. However, combination inevitably introduces a time penalty due to the huge amount of data between workers. This study adopted the high-order function *Bag.foldby*, providing efficient parallel aggregation with minimal communication.

3.2 Data Structure Definition

The MS format is widely used by modern telescopes because it is designed to avoid data redundancy and provides optional sub-tables/columns for universality. We utilize the *python-casacore* (van Diepen 2015), a python interface binding to the casacore library compiled by C++ in radio astronomy, to read and modify the MS. In this study, *dask.bag* to wrap the list of MS paths. When mapping on *dask.bag*, the dask scheduler dispatch splits of *dask.bag* between various workers. Averaging includes complex array computations across the DATA, WEIGHT, and UVW columns that contain most of the MS data. Here, *dask.array* was used to load these three columns as it allows the arrays to be split into chunks and processed in parallel. In addition, the averaging calculation refers to other configurations such as the antennas, polarization, and baseline positions. XArray was used to manage multiple arrays and relate the configurations as a consistent dataset. Table 2 lists the XArray dataset definitions for the experimental data ².

3.3 Accuracy

Due to instrumental failure, partial baselines may be missing ($N_b < n(n-1)/2$) from raw measurement data. Therefore, matrices may be a different shape when they are loaded into the memory. For convenience, the matrices are filled with zeros to compensate for inconsistencies in the data dimensions. Additionally, in Equation (4) and Equation (6), a minimal value ($\Delta\theta$) was introduced into the denominators to avoid the problem of dividing by zero. However, $\Delta\theta$ also introduces accuracy loss. The lower value of $\Delta\theta$ equates to better accuracy. The fact that if two matrices (\mathbf{A} and \mathbf{B}) are numerically close, the mean and variance of *loss error* ($\mathbf{A}-\mathbf{B}$) between them will approach zero was used to examine the accuracy of our proposed method compared to the CASA software. As shown in Table 3, the mean and variance of accuracy loss are close to zero, proving that the two visibility matrices are almost identical.

3.4 Compression Ratios and Imaging Artifacts

Theoretically, data can be averaged over several timesteps and channels with the arguments t (default:2) and nch (default:4). The compression ratio (CR) is the size of the unaveraged visibilities to the averaged one (and thus $CR = t \times nch$ in this study). Cotton (2009) has qualitatively analyzed the CR for time averaging and the level of artifacts in the context of the VLA. He observes that time averaging cannot be of an arbitrary length; otherwise, excessive amplitude attenuation would reduce the *dynamic range* (= ratio of the peak flux to off-center source) in the vicinity of bright sources and produce artifacts that degrade the quality of the image (for detailed image artifacts caused by time smearing, see fig. 4 in Cotton (2009)). Moreover, if the phase variation as a function of channel is large, it may be necessary to restrict the averaging

² The data can be downloaded from http://casa.nrao.edu/Data/EVLA/IRC10216/day2_TDEM0003_10s_norx.tar.gz (the data size is 1.1 GB, the extracted data are about 2.0 GB).

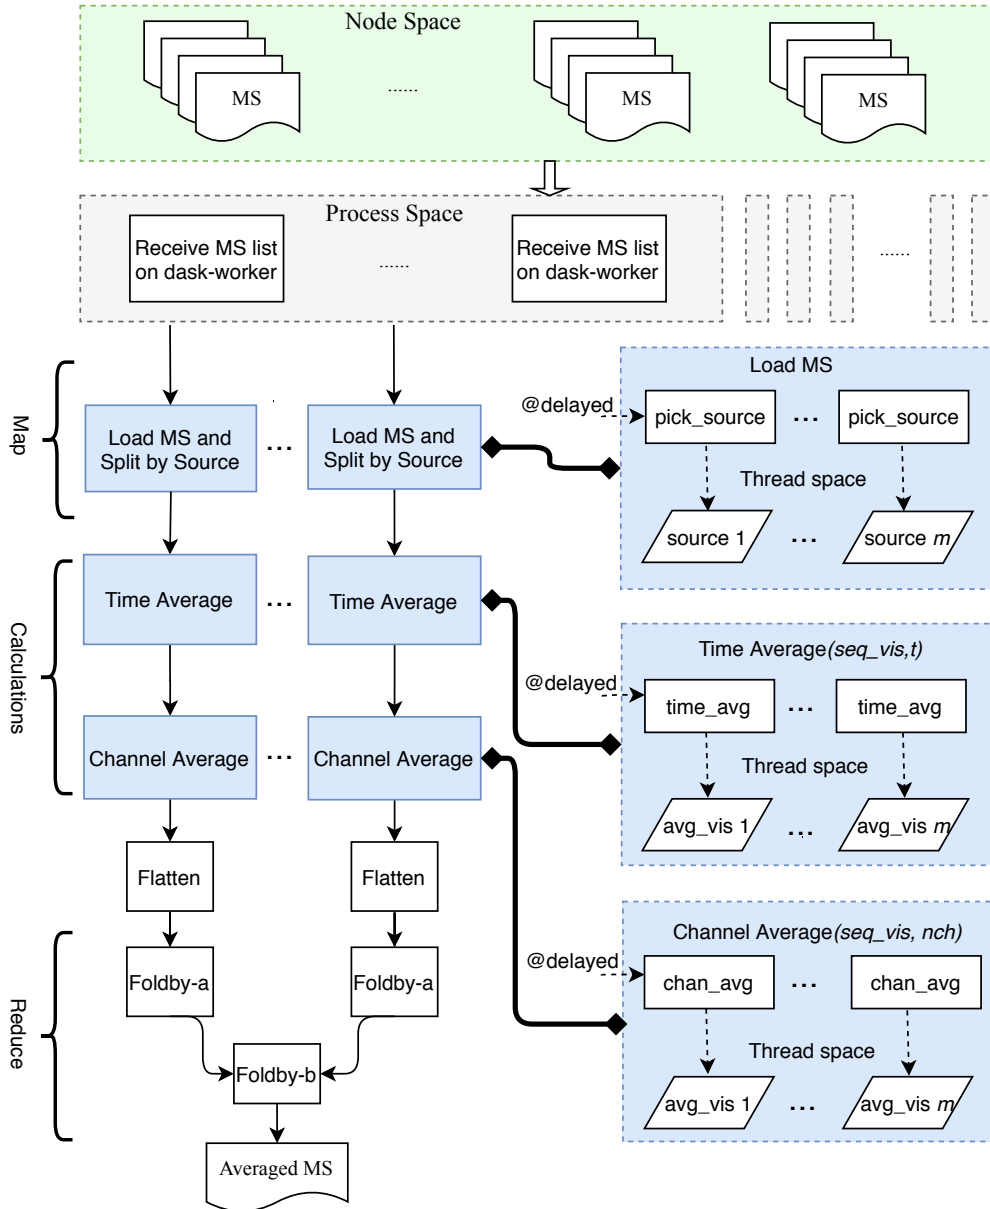


Fig. 2 Multi-level parallelism architecture starts with the MS list loading in the node space. When a node receives a partial list, it begins to load the MS and splits it by source and subsequent operations. This is shown by the *solid blue boxes* that link to the corresponding *dashed blue boxes* to elaborate the internal execution.

to a few channels to prevent the introduction of delay-based closure errors, which can be caused by averaging over non-bandpass corrected channels with large phase variations. To explain frequency smearing, the channel-averaged data were rewritten back to the original MS, and the source images and PSFs were produced (Fig. 3) by the *tclean* task in CASA. The experimental data were obtained by observing two spectral lines produced by the asymptotic giant branch star IRC+10216, where the dense inner envelope (for more information about the circumstellar envelope, see Leão et al. 2006) produced the SiS line as a centralized emission, and the external

envelope produced the HC3N line as a ring of emission. The HC3N can re-radiate energy absorbed from the SiS radiation.

Time and channel averaging on visibilities can be treated as a form of *pseudo-convolution* by some ensemble boxcar-like window functions (Atemkeng et al. 2016), and the sizes of boxcar are proportional to the uv -extent of the time and frequency bins (related to the arguments t and nch). The Fourier transform of a boxcar-like is a sinc-type taper, and the scale of the tapering response is inversely proportional to the scale of the window function. Larger arguments result in wider boxcars corresponding




Fig. 3 Images corresponding to PSFs produced under the same processing flows with different integral channels, $nch = 2$ (panel (a), (b)) and $nch = 32$ (panel (c), (d)). Excessive averaging over channels may cause the image to become blurry.

to narrower tapers with higher sidelobes that are therefore more likely to smearing. As shown in Figure 3, frequency smearing causes a radial broadening distortion of the PSF. Thus, the amplitude of the SiS, modulated by the high sidelobes, translates into an unwanted background in the domain of the HC3N.

4 PERFORMANCE EVALUATION

This section explores the performance characteristics of the algorithm. All of the experiments were performed on ten Ubuntu 18.04 server workstations, two of which are equipped with 56 processors (Intel Xeon CPU E5-2660 v4), hyperthreading, 3.4-GHz maximum frequency, and

512 GB of RAM (i.e., each 32 GB on 16 sockets) and the rest of which equipped with 32 processors (Intel Xeon CPU E5-2630 v3), hyperthreading, 2.4-GHz maximum frequency, and 132 GB of RAM. While each sub-MS contained identical data, larger datasets³ were also tested in our program.

As shown in Figure 4, the coarse-grained multi-processing scheduler is able to bypass the GIL issues to stimulate the potential parallelism of multi-cores. Combined with fine-grained *delayed* function, deferred

³ The original data can be downloaded from https://casa.nrao.edu/Data/EVLA/Pband/P_band_3C129.tgz, but we extract a 13GB portion of this dataset to work on.

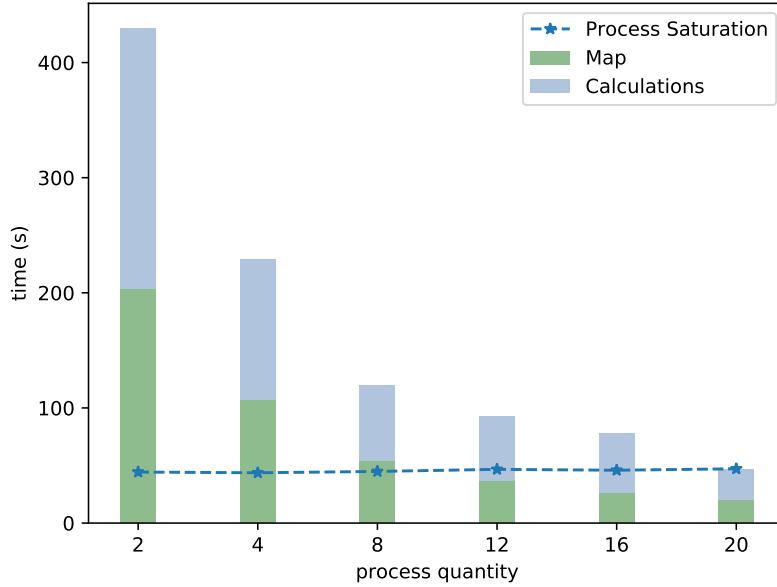


Fig. 4 Coarse-grained process execution with identical arguments $t = 2$ and $nch = 4$ across 20 MSs on 10 nodes (each machine is assigned two dask workers). Process saturation is the number of dask workers, not less than that of MS files, and the execution times indicated by the dashed line do *not* include the reduce phase.

Table 2 XArray Dataset Definitions for the Experimental Data

Dimensions:	(nant: 19, nbaseline: 171, nchan: 64, nfpos: 3, npol: 4, nspos: 3, ntime: 17)
Coordinates:	
* ntime	(ntime) int64 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
* nbaseline	(nbaseline) <U5 '0-1' '0-2' '0-3' ... '16-18' '17-18'
* nchan	(nchan) int64 0 1 2 3 4 5 6 7 ... 56 57 58 59 60 61 62 63
* npol	(npol) <U2 'p0' 'p1' 'p2' 'p3'
* nfpos	(nfpos) <U1 'u' 'v' 'w'
* nant	(nant) int64 0 1 2 3 4 5 6 7 ... 11 12 13 14 15 16 17 18
* nspos	(nspos) <U1 'x' 'y' 'z'
Data variables:	
vis	(ntime, nbaseline, nchan, npol) complex64 dask.array<chunksize=(17, 171, 64, 4), meta=np.ndarray>
weight	(ntime, nbaseline, npol) float64 dask.array<chunksize=(17, 171, 4), meta=np.ndarray>
uvw	(ntime, nbaseline, nfpos) float64 dask.array<chunksize=(17, 171, 3), meta=np.ndarray>
frequency	(nchan) float64 3.639e+10 3.639e+10 ... 3.64e+10
channel_bandwidth	(nchan) float64 1.25e+05 1.25e+05 ... 1.25e+05 1.25e+05
xyz	(nant, nspos) float64 -1.602e+06 -5.042e+06 ... 3.555e+06
mount	(nant) <U6 'ALT-AZ' 'ALT-AZ' ... 'ALT-AZ' 'ALT-AZ'
names	(nant) <U4 'ea01' 'ea02' 'ea03' ... 'ea25' 'ea27' 'ea28'
diameter	(nant) float64 25.0 25.0 25.0 25.0 ... 25.0 25.0 25.0
Attributes:	
tag:	0_7
phasecentre:	<SkyCoord (ICRS): (ra, dec) in deg (202.78453327, 30...
polarisation:	[[5 6 7 8],[5 6 7 8]]

execution can mimic the sequential iteration of the *for-loop* and wrap custom code in sub-operations (see the right in Fig. 2). Along with the gradual increase in dask workers (or it is equivalent to appending more nodes), the map and calculations execution time decreases significantly, and the minimum value occurs at process saturation. The *Dask* distribution performs well during

the loading ($\approx 560 \text{ MB s}^{-1}$) and averaging phases ($\approx 350 \text{ MB s}^{-1}$) (marked by the dashed line in Fig. 4) where time consumption does not increase linearly with the number of the MS. This slow growth is due to the extra coordination required to manage resource competition caused by the large amount of memory accesses and intensive computation. Although distributed environments

Table 3 Accuracy Loss of the DATA Column with Time Intervals t of 2 and Various Channels nch

Averaged parameters	Mean of		Variance of	
	real part	imaginary part	real part	imaginary part
$t = 2, nch = 2$	-2.21427233e-13	-6.91339540e-13	4.25529084e-20	4.12462738e-20
$t = 2, nch = 4$	-3.44819459e-13	-1.11037129e-12	4.29342260e-20	4.11034429e-20
$t = 2, nch = 8$	-9.18691002e-13	3.06048242e-14	5.25735201e-20	5.13628739e-20
$t = 2, nch = 16$	-6.10984707e-13	-2.06883846e-13	8.46708015e-20	7.86904922e-20
$t = 2, nch = 32$	6.62003601e-13	1.14721947e-12	1.42577666e-19	1.30641657e-19
$t = 2, nch = 64$	-5.08093825e-12	-7.35946702e-12	2.92768210e-19	2.52621540e-19

The experimental MS was averaged from the original 1-s correlator visibility integration time to 10-s averages; thus, $t = 2$ represented 20-s time averaging. The averaged visibilities were calculated using $\Delta\theta = 1e-16$ in our programs and the *split* task in CASA with parameters `timebin=20 s` and `width=nch`, respectively.

can increase computing and storage capacity, transferring large amounts of data between processes and separate workstations can introduce more performance penalties. In general, multiprocessing schedulers are an excellent choice when the workflow is relatively linear.

5 CONCLUSIONS

With the advent of interconnected, fault-tolerant, and resource-sharing clusters, distributed computing is becoming a mainstream method of meeting astronomical computing performance requirements. This paper presents a fine-grained distributed method via distributed in-memory chunking provided by *Dask* and a labeled data array provided by *XArray*. Averaging is considered an important approach for reducing the volume of synthetic visibilities. Furthermore, it is advisable to apply the maximum possible calibration prior to averaging, as a high CR requires accurate phase calibration. We take averaging as the case and implement a fine-grained distributed averaging over time and channel. The accuracy and efficiency of the proposed method is demonstrated through a series of experiments. Radio interferometric data processing comprises a variety of algorithms that are more sophisticated than averaging, such as calibration, gridding, and imaging, and the parallelization is a consistent theme throughout them. The practicality of fine-grained distributed averaging may also lead to other aspects of interferometric data distributed processing in a future study.

Acknowledgements This work is supported by the National Key Research and Development Program of China (2020SKA0110300), the Joint Research Fund in Astronomy (U1831204 and U1931141) under cooperative agreement between the National Natural Science Foundation of China (NSFC) and the Chinese Academy of Sciences (CAS), the NSFC (No. 11903009), the Funds for International Cooperation and Exchange of the NSFC (11961141001), Yunnan Key Research and Development Program (2018IA054). The authors wish to thank the

reviewers for suggestions that improved the paper.

References

- Atemkeng, M., Smirnov, O., Tasse, C., Foster, G., & Jonas, J. 2016, *MNRAS*, 462, 2542
- Cornwell, T., Wortmann, P., Nikolic, B., Wang, F., & Stolyarov, V. 2020, *Radio Astronomy Simulation, Calibration and Imaging Library*, <https://github.com/SKA-ScienceDataProcessor/rascal>
- Cotton, W. 2009, Effects of baseline dependent time averaging of UV data
- Dewdney, P. E., Hall, P. J., Schilizzi, R. T., & Lazio, T. J. L. 2009, *Proceedings of the IEEE*, 97, 1482
- Hamaker, J., Bregman, J., & Sault, R. 1996, *Astronomy and Astrophysics Supplement Series*, 117, 137
- Leão, I., De Laverny, P., Mékarnia, D., De Medeiros, J., & Vandame, B. 2006, *A&A*, 455, 187
- McMullin, J., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, *RA Shaw, F. Hill, & DJ Bell*, 127
- Offringa, A. 2016, *Astronomy & Astrophysics*, 595, A99
- Perkins, S., Molenaar, G., Smirnov, O., & Andati, L. 2020, *Xarray Datasets from CASA Tables*, <https://dask-ms.readthedocs.io/en/latest/concepts/pyarrays.html>
- Raba, R., Thomas, A., & Steeb, J.-W. 2019, *CASA Next Generation Infrastructure*, <https://cngi-prototype.readthedocs.io/en/0.0.54/visibilities.html>
- Rocklin, M. 2015, in *Proceedings of the 14th python in science conference* No. 130-136, Citeseer
- Smirnov, O., Frank, B., Theron, I., & Wood, I. H. 2012, in *2012 International Conference on Electromagnetics in Advanced Applications*, IEEE, 586
- Smirnov, O. M. 2011, *A&A*, 527, A106
- Thompson, A. R. 1999, in *Synthesis Imaging in Radio Astronomy II*, 180, 11
- van Diepen, G. 2015, *Astronomy and Computing*, 12, 174
- Wijnholds, S., Willis, A., & Salvini, S. 2018, *MNRAS*, 476, 2029