Optimising the Processing and Storage of Radio Astronomy Data

1st Alexander Williamson International Centre for Radio Astronomy Research University of Western Australia Perth, Australia 0000-0003-0510-4311 2nd Pascal J. Elahi Pawsey Supercomputing Research Centre Perth, Australia 0000-0002-6154-7224

3rd Richard Dodson International Centre for Radio Astronomy Research University of Western Australia Perth, Australia 0000-0003-0392-3604 4th Jonghwan Rhee International Centre for Radio Astronomy Research University of Western Australia Perth, Australia 0000-0001-8496-4306

5th Qian Gong Oak Ridge National Laboratory Oak Ridge, USA 0000-0002-3570-4142.

Abstract—The next generation of radio astronomy telescopes are challenging existing data analysis paradigms, as they have an order of magnitude larger collecting area and bandwidth. The two primary problems encountered when processing this data are the need for storage and that processing is primarily I/O limited. An example of this is the data deluge expected from the SKA-Low Telescope of about 300 PB per year. To remedy these issues, we have demonstrated lossy and lossless compression of data on an existing precursor telescope, the Australian Square Kilometre Array Pathfinder (ASKAP), using MGARD and ADIOS2 libraries. We find data processing is faster by a factor of 7 and give compression ratios from a factor of 7 (lossless) up to 37 (lossy with an absolute error bound of 10^{-3}). We discuss the effectiveness of lossy MGARD compression and its adherence to the designated error bounds, the trade-off between these error bounds and the corresponding compression ratios, as well as the potential consequences of these I/O and storage improvements on the science quality of the data products.

I. INTRODUCTION

A. The Square Kilometre Array

Radio Astronomy is undergoing a paradigm shift with the planning for a number of next-generation instruments, such as the Square Kilometre Array (SKA), the next-generation Very Large Array (ngVLA) and the next-generation Event Horizon Telescope (ngEHT). All of these provide an order of magnitude increase in bandwidth and a few orders of magnitude in collecting area (and sensitivity) over current radio telescopes. This enhancement will provide us with the opportunities to survey the radio sky in exquisite detail, detect the signal from the epoch of reionization when the first stars were born and measure the spectral signal from millions of galaxies. To grasp this potential massive improvement in our understanding, the radio astronomy community must manage and process unprecedented volumes of data. We are focusing on the SKA, as Australia is a founding member of the collaboration. The SKA will be built in Australia for frequencies spanning 50 to 350MHz (SKA-Low) and in South Africa for frequencies from 350MHz to 15GHz (SKA-Mid). Phase 1 will have about 500 40m aperture array elements in SKA-Low and about 200 15m parabolic dishes in SKA-Mid. Construction has commenced and preliminary observations will be made from 2025. The data rates out of the correlator will be 1 and 2.5 TB/s respectively, which will need to be captured into a local buffer and then processed on the day - as storage will be limited.

B. Australian SKA Pathfinder (ASKAP)/Yandasoft

The Australian Square Kilometre Array Pathfinder (ASKAP) [1], [2] radio telescope is one of the SKA

precursors and is opening up a new window for large extragalactic H I surveys beyond the local Universe due to its wide spectral bandwidth and large instantaneous field-of-view (FoV). ASKAP consists of 36 dishes, each of diameter 12 m and equipped with phased array feeds (PAFs) forming multiple receiving beams electronically [3], [4]. The baseline lengths of the full array are from 22 m to 6.3 km. The phased array feed technology allows ASKAP to have a large 30 deg² FoV [5] and a wide bandwidth of 288 MHz with a channel resolution of 18.5-0.58 kHz in the observing frequency between 0.7 and 1.8 GHz, which makes it an optimal survey instrument, enabling it to conduct both wide and deep surveys in a comparatively short period of time [2].

C. Deep Investigation of Neutral Gas Origins (DINGO)

Deep Investigation of Neutral Gas Origins (DINGO) [6], [7] is an ASKAP deep HI survey project aiming to provide a cosmologically representative dataset for HI emission, enabling studies of the HI gas content of galaxies over the past 4 billion years out to distances of 5 billion light-years, due to the accelerated expansion of the universe. The sky coverage of the DINGO survey is wider than deep HI surveys previously conducted and the ongoing deep HI surveys being carried out with other telescopes such as the JVLA¹, MeerKAT² and FAST³ (CHILES⁴, LADUMA⁵/MIGHTEE-HI⁶, and FUDS⁷ respectively). Due to its large volume coverage, the DINGO survey will reduce cosmic variance on HI measurements, thereby providing a unique legacy HI dataset.

DINGO pilot observations were made over the Galaxy and Mass Assembly (GAMA) [13] 23 h (G23) field, centred at α , δ (J2000) = $22^{h}59^{m}00.00, -32^{\circ}18'00.00, -32^{\circ}18'0.00, -32^{\circ}18'0.00,$

⁶The H I emission project of the MeerKAT International GigaHertz Tiered Extragalactic Exploration survey [11]

⁷FAST Ultra Deep Survey [12]

pilot survey obtained 100 hr data both in the band 1 and 2 frequency from one of the G23 tiles to develop a DINGO processing pipeline for deep imaging and long-term data storage, which this paper presents.

The DINGO full survey allocation is 3200 hr for the G23 field, split into two 1600 hr for the frequency ranges of 859.5-1147.5 MHz and 1151.5-1439.5 MHz, respectively. 16 hr observations have been conducted in the higher frequency band so far.

D. ADaptive I/O System version 2 (ADIOS2)

The Adaptable Input Output System version 2 [14], is a software framework with a simple input/output abstraction and a self-describing data model centred around distributed data arrays, allowing multiple applications to publish and subscribe data at large levels of concurrency. It also introduces a larger organizing concept, the "step", for driving data production and consumption within applications. ADIOS2 recently developed a new mechanism to allow applications to use state-of-the-art lossless and lossy compression algorithms. This mechanism makes use of tight integration between I/O and reduction and allows applications to take full advantage of the self-describing formatting and lossy compression techniques.

E. MultiGrid Adaptive Reduction of Data (MGARD)

MGARD [15] is a software offering error-controlled lossy compression rooted in multi-grid theories. It transforms floating-point scientific data into a multilevel representation, followed by quantization and lossless encoding processes, ultimately generating this information in a self-describing compressed buffer. One of MGARD's notable features is its array of error control options, including L^{∞} , L^2 , point-wise relative L^{∞} , and options to define varied error bounds across regions or different frequency components. This flexibility is valuable for preserving Quantities-of-Interest (QoI) [16] derived from the reconstructed data. For region-adaptive compression, MGARD accommodates Regions-of-Interest (RoI) specified through either bounding boxes or masks, with the latter especially useful for irregularly shaped RoIs. In cases where RoI information is not provided, MGARD employs internal functions to identify regions rich in detail, leveraging data turbulence measured across multiple scales.

F. Radio Astronomy Data

Radio data presents a unique challenge. Much of the data is noise, see for example Fig. 2. In this figure,

¹The Karl G. Jansky Very Large Array

²The Meer-Karoo Array Telescope

³Five-hundred-meter Aperture Spherical Telescope

⁴The COSMOS H I Large Extra-galactic Survey [8]

⁵Looking At the Distant Universe with the MeerKAT Array [9], [10]

only a small fraction of pixels in the image contain emission from a galaxy, which appears as a spatially concentrated region of high radio emission. Not all astronomical sources are spatially concentrated and with ever improving resolution, what was once a single source can be resolved into spatially extended, diffuse emission. Moreover, some signals, such as the sought-after signal of reionization from the first stars, will be distributed across the entire image and is hidden in the noise.

This data analysis challenge is combined with a data volume challenge. Radio astronomy data volumes from current generation telescopes are of PB-scale. This data is also often stored as a MeasurementSet [17], a format in which visibility and single-dish data are stored to accommodate synthesis. Although this format has been historically very useful, it does not scale particularly well and often the science process requires non-optimal access, giving rise to additional IO load.

This data challenge will only increase once nextgeneration telescopes become operational and is our motivation for this study.

II. METHODS

A. ASKAPSoft

ASKAPSoft is a package that contains the software necessary for processing data from the ASKAP telescope. Its primary purpose is for the full-scale processing of ASKAP data, from the observed data to spectral-line and continuum images.

ASKAPSoft provides the well-established routines to image a spectral line dataset. That is to: *read* the data, apply *weighting* kernels to set the image parameters (Field of View, sensitivity to low surface brightness or compact objects, etc), resample the data onto a regular sampled *grid* for *inversion* using the Fast Fourier Transform, iteratively *deconvolve* for the limited sampling of Fourier terms and finally, for deep images, *stacking* of multiple epochs of observing for the final image. The DINGO pipeline reorganises these tasks so that grids are preserved for stacking, rather than the images. This reduces the number of inversions required and improves the quality of the deconvolution.

Imaging radio interferometric data at SKA-scales is expected to be I/O bound due to the massive size of the datasets. The computational costs are dominated by the gridding and the inversion steps, and these two are expected to have approximately similar requirements. Thus any reduction in the size of the datasets would have a significant impact on the total processing time and any reduction in the gridding or inversion would have a significant (albeit smaller) impact on the compute costs.

ASKAPSoft contains a wide variety of programs and scripts that are useful in the analysis, manipulation, and processing of radio astronomy data. We will primarily use two applications present in ASKAPSoft: imager & cdeconvolver. imager creates spectral-line image cubes, which use frequency as an analogue for distance, allowing for a 3-dimensional view of the sky. Thus this single program includes the read, weighting, gridding and inversion steps. In our case, imager is used to produce visibility grids, an intermediate product, that we can manipulate before producing the final image. That is, the inversion step is not performed and the normally intermediate data products (i.e. the grids) are saved for later processing. Due to the sparse nature of these grids, compression is particularly efficient and is at a sufficiently early point in the pipeline, so that processing parameters can be changed as needed. The deep imaging mentioned in section I-C requires stacking these grids over 3200 hours of observed data to improve sensitivity to sources within the final image. cdeconvolver is a bespoke application designed specifically to perform this stacking and complete the imaging process. Further information on these applications can be found at https://www.atnf.csiro.au/computing/software/ askapsoft/sdp/docs/current/index.html

B. DINGO Pipeline

We start with a MeasurementSet data format [17] that contains the calibrated, continuum-subtracted visibilities (the continuum here refers to the components of the data that are independent of frequency). This is the input to imager, described above, which produces a visibility grid, a PSF (Point Spread Function) grid, and a PCF (Point Convolution Function) grid. The visibility grid is a grid representation of the 3-dimensional visibilities projected onto a 2-dimensional grid for each frequency channel (producing a 3-dimensional grid). The PSF grid represents the inherent smearing of point sources due to the baseline sampling of the system. The PCF grid represents the size, location and weighting of the convolutional kernels applied during the gridding process. These are used in the final imaging to apply weighting to individual visibility cells. These grids are passed to the cdeconvolver application which, if more than one observation is provided, will sum the grids together as they are read in. These summed grids are then imaged for analysis. Here we validate that the resulting image

is free of detrimental RFI and that the sensitivity of the image is better than that of the non-stacked image.

C. Compression

The focus of the compression here is that of the grids. Due to the nature of the PCF grid, lossless compression provided a compression ratio of ~100, with lossy compression providing ratios similar to that of the visibility and PSF grids. For this reason, we are only comparing the compression of visibility and PSF grids. We use error bounds of 10^{-3} , 10^{-4} , 10^{-5} , and 10^{-6} both as relative and absolute error bounds. The cdeconvolver application failed to complete the imaging of the 10^{-3} data for the relative case, the reasons for which are still under investigation. The lossless compression uses the zstd algorithm to compress the grids, although bzip2 also provides a similar level of compression. This provided a consistent compressed data in any way.

D. Parallel I/O

The integration of ADIOS2 into the I/O stage of ASKAPSoft leveraged the use of casacore and the implementation of existing ADIOS storage manager [18]. An ADIOS Image module was developed to bridge the gap between imager and cdeconvolver's use of image inputs and outputs, and the storage manager's table interface. These applications use MPI to implement parallel processing which, when passing the communicators to the storage manager (and by extension ADIOS2), enables the I/O to occur in parallel.

III. RESULTS

A. Compression Comparison

Figure 1 shows the compression ratios of the visibility and PSF grids as the error bound is increased from zero (lossless) to 10^{-3} for both absolute and relative error measurements. Lossless compression is consistent at a value of 7.5, lossy compression appears to be better in the relative case, and the compression ratios of the real and imaginary parts of the visibility grid are consistent whereas the imaginary parts compress better than the real part in the absolute case and worse in the relative case. This is due to the real and imaginary parts of the PSF (unlike the visibilities) represent different properties and cover a different range of values.

Of the eight tests of the compression, seven completed the imaging and produced, qualitatively, decent images. The reason for the failure of relative- 10^{-3} result is still under investigation, however this is likely after



Fig. 1. The compression ratio for increasing specified error bounds. Shown are the visibility and PSF grid compression ratios for lossless compression, lossy compression with absolute error bounds, and that with relative error bounds.

compression. These images are shown in Figure 2, which shows a collapsed view over the 60 channels that contain a radio source. The only images in this panel that show any qualitative deviation from the original (or lossless) image is that for the 10^{-4} relative error bound and the 10^{-3} absolute error bound. This is further reinforced by Figure 3 which shows that majority of the value deviation is at the corners of the image, with the image for the 10^{-4} relative and 10^{-3} absolute error bounds showing an increased residual near the centre of the image. The spectral profile shown in the left panel of Figure 4 describes the typical double peak profile of a galaxy. This is the brightest source within this image and provides a good test for the quality of the reconstruction of bright sources after MGARD compression. The panels on the right describe the residuals of this profile between the original image and the compressed image for each error bound. The residuals are shown to be uniform and consistent with the set error bounds, noting that the compression was performed on the visibility and PSF grids and the error bounds are set on the values of these grids.

B. Time comparison

Figure 5 shows the comparison between the preexisting I/O method (labelled CASA here) and the same



Fig. 2. Images produced by cdeconvolver from visibility grids compressed with MGARD at different error bounds. The lossless image (top right) is identical to the original and can be used as a ground truth in this case. Bright emission is dominated by single source in this example image.

processing done while using ADIOS2 without engaging any parallel I/O. ADIOS2 appears to perform similarly to CASA for small amounts of data (640 MB to 6.25 GB), but improves significantly as more data (60 GB and above) is written during processing. Figure 6 shows the comparison of the same processing as above, but using parallel-enabled ADIOS2 for writing. The same trend as before can be seen here, where ADIOS2 performs significantly better as more data is written to disk. Figure 7 shows this comparison for the 100-channel case, but with MGARD compression during the writing stage of the processing. Lossless compression appears to perform significantly better than the same processing with lossy compression. In the current implementation, MGARD is run completely on CPUs, however, it is written to run on GPUs as well. We plan to implement and test this compression using GPUs in the future, as we expect that this will be more time-efficient. However, if the processing is I/O bound the computational costs will not be highly significant.

IV. DISCUSSION

A. Compression ratios vs distortion

The histograms of the image residuals (Figure 8) indicate the quality of the reconstructed data after imaging. MGARD produces a consistent residual distribution in terms of distribution shape and the maximum residuals are consistent with the specified error bounds during compression. The residuals produced from the images compressed with absolute error bounds surpass one standard deviation for the 10^{-4} and 10^{-5} cases, and those for the relative error bounds surpass 1σ for the 10^{-3} case. That is, the distortions introduced by the compression, will start to be detectable against the noise levels in the images. The images shown in Figure 3 show that majority of these high residuals are situated in the corners of the image which are cut off or normalised out during regular imaging. An alternate distortion measure is the 2-point correlation of the residuals. This can be produced by performing a Discrete Fourier Transform (DFT) on the residuals, binning the pixels in radius, and calculating the product of the value of these pixels with their conjugate. The result is shown in Figure 9 and describes the prominence of patterns of certain scales within the residuals. The higher error bounds all show a sharp cutoff at 289, whereas the residuals for 10^{-3} absolute error and 10^{-4} relative error show a negative logarithmic trend between 289 and 410. The value of 289 corresponds to the largest radial distance of an nonzero cell on the visibility grid and corresponds to the smallest resolvable scale. The excess of values above this spatial frequency shows a leakage of values within the PSF grid, adding small scale residuals below the smallest resolvable scale. Figure 10 shows the distribution of the residuals for the visibility and PSF grids. Each set of residuals appears to meet the specified error bounds in a consistent manner, each turns off toward zero at the specified error bounds and only exceed these bounds by a factor of two. The exception to this is for the 10^{-6} absolute error bound, where the visibilities exceeded the error bound by a factor of three and the PSF exceeded the bound by a factor of ten. This is likely due to the either a limit in the internal error calculation within MGARD or that the error bound of 10^{-6} is pushing the limit on the compressibility of the data and that lossless compression would be a better fit for this choice of error bound.

B. Consequences for SKA data management

The SKA will be, and have, invested a significant amount of capital into the storage, transmission and



Fig. 3. The absolute residuals between the images produced after lossy compression and the uncompressed equivalent. The colour scales are not normalised as the range of values for each image is significantly different. Smallest error bound produces maximum residuals of $\sim 10^{-5}$.



Fig. 4. (Left) The spectral profile (relative to the RMS of the cube) of a galaxy source within the field of interest. (Right) the residuals of the profile when comparing the original cube and the cubes that have undergone MGARD compression.



Fig. 5. The real runtime (normalised by No. of channels per core) of the imager application with varying numbers of channels. CASA here refers to the standard image I/O manager and ADIOS2 refers to the ADIOS2 I/O manager. Both I/O managers have been run in serial for comparison.

V. CONCLUSION

We have demonstrated the application of MGARD compression, via the ADIOS2/CasaCore framework, within standard radio astronomy software. The impact will be on improved I/O performance, and thus pipeline run times, due to the reduced dataset sizes.

Comparing the images made with various compression approaches allows us to quantify the impact. We show that we can lossy-compress the data files by factors of up to 15, using well-defined error bounds, without impacting the results.

In addition, the parallel reading and writing provided by ADIOS offers an additional improvement in I/O,

computing infrastructure around both nodes of the SKA. The tests described in the preceding sections have been specifically designed around the bespoke nature of deep imaging with ASKAP. Preliminary tests have shown that applying MGARD to raw SKA data (simulated) will yield similar factors of compression, the inclusion of which would simplify the implementation of storage and transmission solutions considerably for the SKA. The processing done for SKA and ASKAP to produce images is almost identical, meaning that implementing ADIOS2 parallel I/O in the SKA pipelines would only improve the efficiency and speed of processing.



Fig. 6. The real runtime (normalised by No. of channels per core) of the imager application in the same manner as Figure 5. The ADIOS2 I/O manager here is run with MPI enabled and is running in parallel. The CASA I/O manager does not have parallel I/O functionality and is still run in serial. The CASA data point at 7776 channels is a lower limit (due to the processing timing out at 5 hours).



Fig. 7. The processing time difference using CASA and parallelenabled ADIOS2 with different error bounds for both absolute and relative error. ADIOS2 is using MGARD to compress the data during the writing stage of the processing.

and this is readily integrated with most common Radio Astronomy applications via the CasaCore libraries.

Compression of the visibilities offers an attractive solution to the SKA I/O challenge, and we have demonstrated that the MGARD approach to compression can guarantee that the data is not degraded.

REFERENCES

[1] S. Johnston, M. Bailes, N. Bartel, C. Baugh, M. Bietenholz, C. Blake, R. Braun, J. Brown, S. Chatterjee, J. Darling,



Fig. 8. The distribution of residuals between the images produced with lossy compressed data and the ground truth image. The residuals are presented relative to the RMS of each 100-channel cube.



Fig. 9. The 2-point correlation of the residuals of the images produced with varying error bounds. The x-axis represents the spatial frequency bins of the image, i.e. Large scale = 0, small scale = 700.

and et al., "Science with the australian square kilometre array pathfinder," *Publications of the Astronomical Society of Australia*, vol. 24, no. 04, pp. 174–188, Jan 2007. [Online]. Available: http://dx.doi.org/10.1071/AS07033

[2] A. W. Hotan, J. D. Bunton, A. P. Chippendale, M. Whiting, J. Tuthill, V. A. Moss, D. McConnell, S. W. Amy, M. T. Huynh, J. R. Allison, C. S. Anderson, K. W. Bannister, E. Bastholm, R. Beresford, D. C. J. Bock, R. Bolton, J. M. Chapman, K. Chow, J. D. Collier, F. R. Cooray, T. J. Cornwell, P. J. Diamond, P. G. Edwards, I. J. Feain, T. M. O. Franzen, D. George, N. Gupta, G. A. Hampson, L. Harvey-Smith, D. B. Hayman, I. Heywood, C. Jacka, C. A. Jackson, S. Jackson, K. Jeganathan, S. Johnston, M. Kesteven, D. Kleiner, B. S. Koribalski, K. Lee-



Fig. 10. The distribution of residuals between the lossy compressed visibility and PSF grids and the original versions of each. The absolute error distributions are presented relative to the absolute residuals and the relative error distributions are presented relative to the range of values the visibility and PSF grids. The vertical lines correspond to the error bound applied during compression.

Waddell, E. Lenc, E. S. Lensson, S. Mackay, E. K. Mahony, N. M. McClure-Griffiths, R. McConigley, P. Mirtschin, A. K. Ng, R. P. Norris, S. E. Pearce, C. Phillips, M. A. Pilawa, W. Raja, J. E. Reynolds, P. Roberts, D. N. Roxby, E. M. Sadler, M. Shields, A. E. T. Schinckel, P. Serra, R. D. Shaw, T. Sweetnam, E. R. Troup, A. Tzioumis, M. A. Voronkov, and T. Westmeier, "Australian square kilometre array pathfinder: I. system description," PASA, vol. 38, p. e009, Mar. 2021.

- [3] D. R. DeBoer, R. G. Gough, J. D. Bunton, T. J. Cornwell, R. J. Beresford, S. Johnston, I. J. Feain, A. E. Schinckel, C. A. Jackson, M. J. Kesteven, A. Chippendale, G. A. Hampson, J. D. O'Sullivan, S. G. Hay, C. E. Jacka, T. W. Sweetnam, M. C. Storey, L. Ball, and B. J. Boyle, "Australian SKA Pathfinder: A High-Dynamic Range Wide-Field of View Survey Telescope," *IEEE Proceedings*, vol. 97, pp. 1507–1521, Aug. 2009.
- [4] G. Hampson, A. Macleod, R. Beresford, M. Brothers, A. Brown, J. Bunton, C. Cantrall, R. Chekkala, W. Cheng, R. Forsyth, R. Gough, S. Hay, J. Kanapathippillai, D. Kiraly, M. Leach, N. Morison, S. Neuhold, P. Roberts, R. Shaw, A. Schinckel, M. Shields, and J. Tuthill, "Askap paf ade - advancing an lband paf design towards ska," in 2012 International Conference on Electromagnetics in Advanced Applications, Sep. 2012, pp. 807–809.
- [5] J. D. Bunton and S. G. Hay, "Achievable field of view of chequerboard phased array feed," in 2010 International Conference on Electromagnetics in Advanced Applications, 2010, pp. 728– 730.
- [6] M. J. Meyer, "Deep investigation of neutral gas origins (dingo)," ASKAP Survey Science Proposal, 2009.
- [7] J. Rhee, M. Meyer, A. Popping, S. Bellstedt, S. P. Driver, A. S. G. Robotham, M. Whiting, I. K. Baldry, S. Brough, M. J. I. Brown, J. D. Bunton, R. Dodson, B. W. Holwerda,

A. M. Hopkins, B. S. Koribalski, K. Lee-Waddell, Á. R. López-Sánchez, J. Loveday, E. Mahony, S. Roychowdhury, K. Rozgonyi, and L. Staveley-Smith, "Deep investigation of neutral gas origins (DINGO): H I stacking experiments with early science data," MNRAS, vol. 518, no. 3, pp. 4646–4671, Jan. 2023.

- [8] X. Fernández, J. H. van Gorkom, K. M. Hess, D. J. Pisano, K. Kreckel, E. Momjian, A. Popping, T. Oosterloo, L. Chomiuk, M. A. W. Verheijen, P. A. Henning, D. Schiminovich, M. A. Bershady, E. M. Wilcots, and N. Scoville, "A Pilot for a Very Large Array H I Deep Field," ApJ, vol. 770, p. L29, Jun. 2013.
- [9] B. W. Holwerda, S. L. Blyth, and A. J. Baker, "Looking at the distant universe with the MeerKAT Array (LADUMA)," in *The Spectral Energy Distribution of Galaxies - SED 2011*, ser. IAU Symposium, R. J. Tuffs and C. C. Popescu, Eds., vol. 284, Aug 2012, pp. 496–499.
- [10] S. Blyth, A. J. Baker, B. Holwerda, A. Bouchard, B. Catinella, L. Chemin, D. Cunnama, R. Davé, A. Faltenbacher, S. February, X. Fernández, E. Gawiser, I. Heywood, D. Kereš, H. R. Klöckner, P. Lah, M. Lochner, N. Maddox, S. Makhathini, K. Moodley, R. Morganti, D. Obreschkow, S. H. Oh, D. J. Pisano, A. Popping, G. Popping, S. Ravindranath, E. Schinnerer, K. Sheth, R. Skelton, M. Smith, R. Srianand, L. Staveley-Smith, M. Vaccari, P. Vaisanen, F. Walter, S. Rawlings, B. A. Bassett, M. A. Bershady, F. H. Briggs, S. M. Crawford, C. M. Cress, J. K. Darling, R. P. Deane, G. de Blok, E. C. Elson, B. S. Frank, P. A. Henning, K. M. Hess, J. P. Hughes, M. J. Jarvis, S. J. Kannappan, N. S. Katz, R. C. Kraan-Korteweg, M. D. Lehnert, A. K. Leroy, G. R. Meurer, M. J. Meyer, D. J. Pisano, A. C. Schröder, O. M. Smirnov, R. S. Somerville, I. M. Stewart, K. J. van der Heyden, M. A. W. Verheijen, E. M. Wilcots, T. B. Williams, P. A. Woudt, J. F. Wu, M. A. Zwaan, J. T. L.

Zwart, T. A. Oosterloo, and W. van Drie, "LADUMA: Looking at the Distant Universe with the MeerKAT Array," in *MeerKAT Science: On the Pathway to the SKA*, Jan 2016, p. 4.

- [11] N. Maddox, B. S. Frank, A. A. Ponomareva, M. J. Jarvis, E. A. K. Adams, R. Davé, T. A. Oosterloo, M. G. Santos, S. L. Blyth, M. Glowacki, R. C. Kraan-Korteweg, W. Mulaudzi, B. Namumba, I. Prandoni, S. H. A. Rajohnson, K. Spekkens, N. J. Adams, R. A. A. Bowler, J. D. Collier, I. Heywood, S. Sekhar, and A. R. Taylor, "MIGHTEE-HI: The H I emission project of the MeerKAT MIGHTEE survey," A&A, vol. 646, p. A35, Feb. 2021.
- [12] H. Xi, L. Staveley-Smith, B.-Q. For, W. Freudling, M. Zwaan, L. Hoppmann, F.-H. Liang, and B. Peng, "The Arecibo Ultra-Deep Survey," MNRAS, vol. 501, no. 3, pp. 4550–4564, Mar. 2021.
- [13] S. P. Driver, S. Bellstedt, A. S. G. Robotham, I. K. Baldry, L. J. Davies, J. Liske, D. Obreschkow, E. N. Tavlor, A. H. Wright, M. Alpaslan, S. P. Bamford, A. E. Bauer, J. Bland-Hawthorn, M. Bilicki, M. Bravo, S. Brough, S. Casura, M. E. Cluver, M. Colless, C. J. Conselice, S. M. Croom, J. de Jong, F. D'Eugenio, R. De Propris, B. Dogruel, M. J. Drinkwater, A. Dvornik, D. J. Farrow, C. S. Frenk, B. Giblin, A. W. Graham, M. W. Grootes, M. L. P. Gunawardhana, A. Hashemizadeh, B. Häußler, C. Heymans, H. Hildebrandt, B. W. Holwerda, A. M. Hopkins, T. H. Jarrett, D. Heath Jones, L. S. Kelvin, S. Koushan, K. Kuijken, M. A. Lara-López, R. Lange, Á. R. López-Sánchez, J. Loveday, S. Mahajan, M. Meyer, A. J. Moffett, N. R. Napolitano, P. Norberg, M. S. Owers, M. Radovich, M. Raouf, J. A. Peacock, S. Phillipps, K. A. Pimbblet, C. Popescu, K. Said, A. E. Sansom, M. Seibert, W. J. Sutherland, J. E. Thorne, R. J. Tuffs, R. Turner, A. van der Wel, E. van Kampen, and S. M. Wilkins, "Galaxy And Mass Assembly (GAMA): Data Release 4 and the z ; 0.1 total and z ; 0.08 morphological galaxy stellar mass functions," MNRAS, vol. 513, no. 1, pp. 439-467, Jun. 2022.
- [14] W. Godoy, N. Podhorszki, R. Wang, C. Atkins, G. Eisenhauer, J. Gu, P. Davis, J. Y. Choi, K. Germaschewski, K. Huck, A. Huebl, M. Kim, J. Kress, T. Kurc, Q. Liu, J. Logan, K. Mehta, G. Ostrouchov, M. Parashar, and S. Klasky, "Adios 2: The adaptable input output system. a framework for highperformance data management," *SoftwareX*, vol. 12, p. 100561, 07 2020.
- [15] Q. Gong, J. Chen, B. Whitney, X. Liang, V. Reshniak, T. Banerjee, J. Lee, A. Rangarajan, L. Wan, N. Vidal, Q. Liu, A. Gainaru, N. Podhorszki, R. Archibald, S. Ranka, and S. Klasky, "Mgard: A multigrid framework for high-performance, error-controlled data compression and refactoring," *SoftwareX*, vol. 24, p. 101590, 2023. [Online]. Available: https://www.sciencedirect.com/science/ article/pii/S2352711023002868
- [16] Q. Gong, B. Whitney, C. Zhang, X. Liang, A. Rangarajan, J. Chen, L. Wan, P. Ullrich, Q. Liu, R. Jacob, S. Ranka, and S. Klasky, "Region-adaptive, error-controlled scientific data compression using multilevel decomposition," in *Proceedings* of the 34th International Conference on Scientific and Statistical Database Management, ser. SSDBM '22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: https://doi.org/10.1145/3538712.3538717
- [17] A. Kemball and M. Wieringa, "Measurementset definition version 2.0," URL: http://casa. nrao. edu/Memos/229. html, vol. 20, 2000.
- [18] R. Wang, C. Harris, and A. Wicenec, "Adiosstman: Parallelizing casacore table data system using adaptive io

system," *Astronomy and Computing*, vol. 16, pp. 146–154, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2213133716300452