

Deep Investigation of Neutral Gas Origins (DINGO): Options for robust Deep Spectral Line Imaging in the SKA-Era

Jonghwan Rhee^{1,2*}, Richard Dodson¹, Alexander Williamson¹, Martin Meyer^{1,3}, Kristóf Rozgonyi^{1,3}, Pascal J. Elahi⁴, Matthew Whiting⁵, Daniel Mitchell⁵ and Tobias Westmeier^{1,3}

¹International Centre for Radio Astronomy Research (ICRAR), University of Western Australia, 35 Stirling Hwy, Crawley, WA 6009, Australia

²Australia Telescope National Facility, CSIRO Space & Astronomy, P.O. Box 1130, Bentley, WA 6102, Australia

³ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia

⁴Pawsey Supercomputing Research Centre, Kensington, WA 6151, Australia

⁵Australia Telescope National Facility, CSIRO Space & Astronomy, P.O. Box 76, Epping, NSW 1710, Australia

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

The data storage requirements for deep spectral line observations with next-generation radio interferometers like the Australian Square Kilometre Array Pathfinder (ASKAP) and the Square Kilometre Array (SKA) are extremely challenging. The default strategy is to reduce data after each daily observation and stack the resulting images. Although this approach is computationally efficient, it risks propagating systematic errors and significantly degrades the final data quality. However, storage and computation requirements for a traditional way to image the entire deep dataset together are prohibitively expensive. We present an alternative *uv*-grid stacking method and compare its scientific outcomes with both the traditional approach, which processes all data jointly and serves as the best-possible result, and the default image-stacking method. Our technique involves halting the standard imaging pipeline after the daily residual visibility grids are formed. These grids are then stacked and jointly deconvolved to combine many epochs of data. Using the traditional approach as a benchmark, we show that image-stacking recovers only 92% of the true H_I flux. In contrast, our *uv*-grid stacking method recovers 99%, which is in excellent agreement with the traditional method within the noise limits. Furthermore, image-stacking introduces significant non-physical artefacts, such as negative bowls around strong sources, indicating poor deconvolution and a loss of physical information. Based on these findings, we intend to apply the *uv*-grid stacking to the Deep Investigation of Neutral Gas Origins (DINGO) survey on ASKAP and strongly recommend this or a similar approach for future radio astronomy facilities.

Key words: radio lines: galaxies – methods: data analysis – techniques: interferometric.

1 INTRODUCTION

New cutting-edge radio instruments such as the Square Kilometre Array (SKA, [Dewdney et al. 2009](#)) will usher in a golden era of radio astronomy. The SKA will feature high-density interferometric arrays with hundreds or thousands of individual antennas to enable wide-field, high-sensitivity, and high-angular resolution imaging of the radio sky. However, to fully exploit this instrument, the radio astronomy community must manage and process its unprecedented volumes of data ([Quinn et al. 2015](#); [Scaife 2020](#)).

The SKA will be constructed in Australia, operating between 50 and 350 MHz (SKA-Low, [Labate et al. 2022](#)), and in South Africa, covering the frequency range of 350 MHz to 15.4 GHz (SKA-Mid, [Swart et al. 2022](#)). In its first phase, SKA-Low will have 512 aperture array stations of 256 log-periodic dipole antennas each, deployed over a 42-m diameter, while SKA-Mid will consist of 133 offset Gregorian dishes with 15-m diameter and 64 MeerKAT dishes of each 13.5-m diameter. About 50 per cent of stations/dishes will be located within

a dense core area of ~ 1 – 2 km, and the remaining are distributed along three spiral arms. The final goal is to have a full square kilometre of collecting area for both arrays. Construction has commenced, and science verification will begin in 2027¹.

Following this, the data rates for each correlator will become ~ 6 Tb s^{−1} ([McMullin et al. 2020](#)), making storage one of the largest cost drivers for the SKA project. Due to limited storage, raw data will be temporarily captured into a local buffer and must be processed within a specific period, ranging from days to weeks, before being permanently erased. This limitation will inevitably require real-time science data processing (SDP) with only a subset of raw or processed data being archived and provided for scientific use. Thus, for multi-epoch projects, there will be no opportunity to iterate over all the data in the imaging process in the traditional approach, leading to one of the most outstanding challenges for the SKA—how to optimally image multi-day deep datasets. This limitation has the potential to cause all multi-epoch observations to be sensitivity-limited due to systematic errors, thus not achieving the scientific goals.

* E-mail: Jonghwan.Rhee@csiro.au

¹ <https://www.skao.int/en/647/timeline-science>

The SKA precursors are good testbeds for preparing for the forthcoming SKA era by demonstrating scientific and technical readiness. In particular, the Australian Square Kilometre Array Pathfinder (ASKAP, Johnston et al. 2007; Johnston et al. 2008; Hotan et al. 2021) plays a key role in addressing the technical challenges that large-scale radio survey projects of the SKA will face. The Deep Investigation of Neutral Gas Origins (DINGO, Meyer 2009; Rhee et al. 2023) is one of the ASKAP survey science projects and aims to carry out deep spectral-line (H I) observations (3,200 h integration in total) over the G23 region of the Galaxy And Mass Assembly (GAMA, Driver et al. 2011; Hopkins et al. 2013; Liske et al. 2015; Baldry et al. 2018; Driver et al. 2022) survey. This region has the richest multi-wavelength datasets, enabling both H I and continuum science, and establishing a legacy dataset of maximal value for SKA-era surveys. The H I data from DINGO will yield some of the deepest images ever taken of atomic hydrogen gas content in the Universe, enabling groundbreaking new studies of the role that this fundamental component has played in the ongoing evolution of galaxies and its connection to their dark matter halos, in conjunction with multi-wavelength data from GAMA.

On top of its scientific merits, the DINGO data will be beneficial to test and demonstrate a solution for the SKA technical challenges in deep spectral line imaging and long-term data storage, as the DINGO survey will collect data spread over many epochs to obtain an 800 h integration per tile. The final raw data volume would be ~ 13.2 PB, which is beyond what is feasibly stored. In this paper, we explore three deep spectral line imaging methodologies with 200 h of data from the DINGO pilot and main surveys, while a companion paper (Williamson et al. in preparation) focusses on the performance tests we undertook to prove that our *uv*-gridding method is a solution for long-term data storage in radio deep survey projects.

The paper is structured as follows. Section 2 gives details of the DINGO data used in this work, e.g., the target field, observations, and initial data processing. In Section 3, we describe three deep imaging methodologies compared in this paper. Section 4 presents the results of a comparison between three deep imaging methods, followed by our discussion and conclusions in Section 5 and 6. We adopt the concordance cosmological parameters (Hinshaw et al. 2013) of $\Omega_{\Lambda} = 0.7$, $\Omega_M = 0.3$ and $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ throughout this paper.

2 DATA

2.1 Observations

This work used data taken from the DINGO pilot and main survey observations over the G23 field, corresponding to 25 scheduling blocks (SBs) and approximately 200 h integration time in total (see Table 1). The DINGO pilot observations were made in two phases between 2019 and 2022, while the main survey observations began in 2023. We carried out DINGO pilot and main survey observations with the full ASKAP array of 36 antennas and 288 MHz bandwidth split into of 15,552 channels of 18.52 kHz channel, corresponding to a velocity resolution of $\sim 4 \text{ km s}^{-1}$ at $z \sim 0$. The integration time of a normal DINGO scheduling block (SB) is about 8 h except for some of the pilot phase 1 observations where shorter observations were conducted. The ASKAP phased array feeds (PAFs) receiver generates a squared 6×6 (square_6x6) beam configuration, forming an instantaneous field of view (FoV) of $\sim 30 \text{ deg}^2$. The entire G23 field of $12 \times 5 \text{ deg}^2$ is divided into two tiles, and each tile is covered by two




Figure 1. ASKAP beam footprints used for DINGO observations over one tile of the G23 region. Two 36-beam patterns (blue for footprint A and red for footprint B, respectively) of each 6-by-6 beam configuration are interleaved to achieve the uniform sensitivity over the field. Beam IDs are overlaid on individual beams. For this study, we select two beams (dashed-line circles) from each footprint, surrounding a bright H I galaxy of NGC 7361 (denoted with a yellow star) already detected in this field by the HIPASS survey (Koribalski et al. 2004). The background is the optical all-sky image from Mellinger (2009).

interleaved square_6x6 footprints to achieve uniform sensitivity over the tile, as seen in Figure 1. The pilot and main survey observations taken so far focused on one tile in the G23 region, centred at α, δ (J2000) = $22^{\text{h}}47^{\text{m}}17^{\text{s}}.78, -32^{\circ}29'03''.26$.

In pilot phase 1, the two interleaved pointings were alternated every 15 minutes to have one SB contain two footprints to cover an entire tile at a time. However, from pilot phase 2, we changed this observing strategy to have one SB with one footprint. This change led to reducing the flagged data fraction and increasing the integration time of each footprint by a factor of two for deeper imaging. Along with science data, calibration data were obtained by observing PKS B1934-638 before or after science observation.

2.2 Data Reduction

The DINGO data reduction pipeline processes each science SB with ASKAPSOFT (Cornwell et al. 2011; Whiting et al. 2017; Wieringa et al. 2020), ASKAP's science data processing software², to produce DINGO spectral line datasets used for this work, following a standard H I data reduction procedure. This includes flagging, bandpass/flux calibration, self-calibration, continuum subtraction, and spectral line imaging. More details on DINGO data reduction are described in Rhee et al. (2023). These data reduction steps were applied separately to individual beams in each daily SB. The significant RFI impact of global navigation satellite systems (GNSS) below 1295.5 MHz restricts us to using only the upper half of the 288-MHz bandwidth, spanning 1295.5 to 1439.5 MHz, corresponding to 7776 channels.

² <https://www.atnf.csiro.au/computing/software/askapsoft/sdp/docs/current/index.html>

Table 1. The summary of DINGO pilot and main survey observations.

Survey	SBID	Date	Obs. Time [hr]	Footprint ^a	No. of Beams	Data Volume per beam ^b [GB]	Bandwidth [MHz]	Centre Frequency [MHz]
Pilot Phase 1	10991	2019-12-25	8.15	A/B	36	232	288	1295.5
	10994	2019-12-26	6.15	A/B		176		
	11000	2019-12-28	8.16	A/B		232		
	11003	2019-12-29	8.82	A/B		251		
	11006	2019-12-30	8.17	A/B		232		
	11010	2019-12-27	5.70	A/B		162		
	11026	2020-01-02	8.17	A/B		232		
Pilot Phase 2	39697	2022-04-20	8.02	A	36	456	288	1295.5
	32214	2021-09-18	8.04	B		456		
	39722	2022-04-21	8.03	A		456		
	39995	2022-04-28	8.02	B		456		
	40196	2022-05-03	8.03	A		456		
	40018	2022-04-29	8.03	B		456		
Main	65177	2024-08-26	8.05	A	36	456	288	1295.5
	65210	2024-08-27	8.03	A		456		
	65268	2024-08-28	8.03	A		456		
	65297	2024-08-29	8.03	A		456		
	65347	2024-08-30	8.03	A		456		
	65405	2024-08-31	8.03	A		456		
	65665	2024-09-06	8.03	B		456		
	65690	2024-09-07	8.03	B		456		
	65748	2024-09-11	8.03	B		456		
	65818	2024-09-13	8.03	B		456		
	65836	2024-09-14	8.04	B		456		
	65853	2024-09-15	8.10	B		456		

^a Pilot phase 1 observations alternated between footprint A and B in one SB.^b Since SBs of Pilot phase 1 observations have both footprint A and B in one SB, the data volume here is per-beam and per-footprint.

For subsequent deep spectral line imaging analysis, we used calibrated and continuum-subtracted spectral line visibility data of four beams around a well-known bright galaxy (NGC 7361, [Koribalski et al. 2004](#)) in H₁ (see Figure 1). After imaging processes of the three methods compared in this work, an additional continuum subtraction was performed in the image domain to remove continuum residuals. Then, we combine four beams using the mosaicking command of ASKAPSOFT, `linmos`. This mosaicking process includes correction for the ASKAP primary beam response using the ASKAP holography model ([Hotan et al. 2021](#)).

3 METHOD

In this study, we explore three deep imaging methods to identify an optimal spectral line imaging strategy for DINGO, and also establishing a foundation for future deep SKA surveys. Using identical parameters across all methods, we analysed the same dataset via: (1) the traditional approach (visibility stacking), where visibility data from all observations are concatenated prior to imaging; (2) the current default methodology (image stacking) for deep projects with ASKAP and the SKA, where data are imaged by scheduling block (approximately daily observations) and the resulting images are co-added in the image domain; and (3) a novel approach (grid stacking) that we developed, where processing halts after forming uv -grids of the daily data at the final major deconvolution cycle, with these residual grids subsequently stacked before performing a final minor deconvolution cycle. The schematic workflows of these deep imaging approaches are illustrated in Figure 2. In all cases, data reduction

was performed with very similar parameters; four overlapping beams (except the image stacking method that used all beams) with an image size of 1024×1024 per beam, with a restored beam size of $\sim 30''$ and a cell size of $6''$, and maximum baselines limited to 2 km.

The next step is to conduct H₁ source finding with the data products from the three methods and compare physical properties of H₁ sources commonly detected in the three datasets, especially focusing on NGC 7361, the well-known H₁ source in the field. For this analysis, we use an automated 3D source finding software dedicated to ASKAP H₁ surveys, called Source Finding Application (SoFiA, [Serra et al. 2015](#); [Westmeier et al. 2021](#)). It produces moment maps and H₁ spectra of detected sources with their physical parameters derived, such as positions (or coordinates), H₁ velocity (or redshift), fluxes (peak and integrated), and velocity widths (e.g., W_{20} and W_{50}). In the following subsections, we describe the details of the respective workflows.

3.1 Traditional Visibility Stacking

The traditional approach for imaging multi-epoch data, the most developed and best understood methodology, is to co-add the multiple measured visibility data in the uv domain, which is then imaged with the CLEAN algorithm ([Högbom 1974](#); [Clark 1980](#)), one of the most successful deconvolution procedures in radio astronomy ([Cornwell 2009](#)). The CLEAN processing of multi-epoch data combined in the uv domain is robust against artefacts caused by low-level continuum residuals, RFI and/or systematic errors lurking in shallower-depth data. In this regard, we consider this method as the only proven ref-




Figure 2. The schematic flowchart shows three deep imaging methods compared in this study using DINGO data.

erence to deliver the science data products with the required quality. However, it is also the most compute-intensive, with long-term storage of the observed visibility data products currently prohibitive for both SKA and ASKAP. In light of this limitation, the observatory policy was to adjust the survey strategy to reduce each day independently and sum those images, allowing for the deletion of the observed visibilities.

Given the current data volume of the DINGO survey, it is still technically feasible to test this traditional deep imaging approach. To this end, we combined 16 SBs of DINGO pilot and main survey data for each footprint, corresponding to ~ 5.6 TB of spectral-line measurement sets (MSs) for the same beam and footprint, as seen in Figure 1. Instead of physically concatenating 16 MSs followed by imaging, ASKAPsoft allows users to simultaneously execute these two steps in its imaging step. The ASKAPsoft imaging command, called *imager*, accepts multiple MSs of the same beam and footprint as input, enabling both combination and deconvolution during the imaging procedure. Using this functionality, we imaged four beams' data using *imager* with its optimal CLEAN algorithm, *BasisfunctionMFS*, a re-implemented and improved version of the CASA's multi-scale and multi-frequency algorithm (Rau & Cornwell 2011).

After all the combined data has been gridded on a regular grid in the 2D (u, v) space per frequency slice, it is Fourier-transformed to produce an image, called dirty image. Sky-models are created by iteratively identifying peaks in that image from which the instrumental sampling response function (i.e., the Point Spread Function, PSF) is subtracted from this dirty image. This is the so-called minor-cycle CLEAN. The major cycle CLEAN involves degridding those models, reprojecting them onto the original regular time- and frequency-sampled visibility data, and subtracting the models at these points, after which the residuals are regridded, imaged, and passed through

the minor cycle again. Refer to Thompson et al. (2017) for more details about the CLEAN algorithm.

For this paper, the CLEAN process performed three major cycles and minor cycles with a threshold of 5σ , a deep CLEAN threshold of 0.5σ and 1,000 iterations. The deep CLEAN of ASKAPsoft searches pixels already in the model to find new CLEAN components. ASKAPsoft allows users to set these threshold parameters both on a noise basis and by specific numbers in units of Jy. We applied the same sets of imaging parameters, including the CLEAN thresholds based on noise and the number of major cycles, to all three methods.

As mentioned in Section 2.2, the input spectral line MSs used are already continuum-subtracted in the uv -domain at the depth of the daily data, rather than at the full 200 h level. However, we performed additional continuum subtraction from the data cubes of each beam to remove residual continuum emission by subtracting a second-order polynomial fit to the spectrum at each pixel in the data cubes. After this step, we combined the four beams' cubes independently processed using a linear mosaicking command of ASKAPsoft.

3.2 Image Stacking

This method combines the individual image cubes already produced by the ASKAPsoft DINGO pipeline in the image domain. The DINGO data are processed on a daily (SB) basis. After calibration and continuum subtraction in the uv -domain, 3 major-cycle CLEANs are performed in spectral line imaging, followed by image-domain continuum subtraction, as in the traditional approach. Then the spectral line data cubes of all 36 beams in each SB are combined using the ASKAPsoft mosaicking command. These data cubes of 16 SBs for each footprint in this work are combined again into the final data cube to cover the entire field of interest, as seen in Figure 1.

Since this processing is done on the daily subset of the entire data, it is computationally straight forward and inexpensive, and necessary for rapid data validation. However, this approach has strong disadvantages in that it irreversibly embeds all artefacts deeper than the individual observational level, leaving sidelobe and continuum residual artefacts in the finally combined data. The vast majority of final spectral line sources at the long integration level (800 h per tile and band for the DINGO main survey) is unlikely to be CLEANed appropriately.

3.3 Grid Stacking

Given the limitations of the two approaches mentioned above, a new intermediate approach was needed and we have adopted the stacking of the uv -grid residuals. Our preferred solution is to store the daily residual data in the uv -grid along with CLEANed model data from major cycles, which can reduce storage requirements over the traditional approach. This method has been extensively developed as part of the PhD thesis of Rozgonyi (2021) and the shared ICRAR-Pawsey-AusSRC HiVIS³ PaCER⁴ project (Rozgonyi et al. 2022). It has the advantage of enabling minor-cycle CLEANing at full survey depth and improving continuum subtraction. However, since the daily data are stored after application of the uv -kernels and combined on the 2D uv -plane, major cycles are not possible, so we must ensure that we work with nearly final residuals, as in that case further major cycles are not required. Although this is a highly promising approach and the data products of this method can be stored long-term (enabling future reprocessing), detailed comparison of this approach to the other deep imaging solutions has not been published until now, other than in the PhD thesis of Rozgonyi (2021).

The generation of daily residual grids can be part of the standard ASKAP processing using *imager* with limited additional costs, where the cubes are CLEANed with 3 major cycles. However, some additional steps are required. Firstly, the daily uv -grid residuals must be stacked, imaged and deconvolved, for which we use the ASKAPSOFT deconvolution command, *cdeconvolver*. Following this, the daily CLEAN models must be summed, convolved with the ‘clean beam’ and added back into the final data once the final residual cube is generated in the same manner as the final model components, albeit in a separate processing step.

4 RESULTS

4.1 Comparison of the Images

The comparison of the three imaging pipelines demonstrates the disadvantages and benefits of the different strategies, where we treat the traditional visibility-stacking image as close to the ‘ground truth’ as we can obtain. For the daily images, the residual level was the same regardless of whether images or grids were produced, as their generation process is identical. For stacked images, the residual level is lower and very similar between the different methods, as detailed in Table 2.

In Figures 3 and 4 we compare the images from the three pipelines with their differences and ratio, and the derived H I spectra for the strongest source in the DINGO observation field, NGC 7361. In

Table 2. Comparison of the physical parameters of NGC 7361 from the three images.

Parameter	V	G	I
RMS [mJy beam ⁻¹]	0.64	0.64	0.64
S _{int} [Jy Hz]	206427.4±701.0	200995.4±700.5	201087.4±724.0
W ₂₀ [km s ⁻¹]	230.7	230.1	228.9
W ₅₀ [km s ⁻¹]	224.1	223.7	223.0
log (M _{H I} /M _⊙)	9.52±0.32	9.51±0.31	9.51±0.31

Figure 3 on the diagonal are the moment 0 images of the three methods (V, traditional visibility-stacking; G, grid-stacking; I, image-stacking), converted to H I column density. Below the diagonal we show the difference between the methods and above the diagonal we show the ratio between the methods, masked where the emission is less than the 1 σ column density sensitivity. The contour levels of 1, 7, 50 and 500 σ column density sensitivity are overlaid.

All three methods yield broadly similar results. However, the image-stacking moment map shows that there are significant and extended residual sidelobes left around the galaxy, likely due to the limited deconvolution capability inherent to the method. The difference in the comparison of the moment 0 maps is greater for the V-I image than for the V-G image, with peak differences approximately 26% higher. Additionally, the ratio of V/I shows a greater range and more complex structures than the ratio of V/G. In Figure 4, we also plot the spectra integrated over the source in the images and their respective differences. The spectral differences are at the few-percent level, with the V-G difference echoing the original spectra, while the V-I difference shows greater scatter and slightly higher residuals. This is consistent with the results presented in Figure 3.

Our results clearly demonstrate the limitations of image-stacking, particularly its failure to accurately reproduce the traditional visibility-stacking results around strong sources. This is not unexpected; previous work has shown the importance of retaining visibility data for deep spectral line imaging (e.g., Dodson et al. 2016).

Stacking daily images also would co-add uncleaned residuals into the final product. Consequently, sidelobes from uncleaned flux, especially around strong sources, can dominate the final image, as seen in Figure 3. Our study, however, shows that grid-stacking achieves results comparable to those obtained with visibility-stacking. Given that the storage required for full visibility-stacking is not feasible for large-scale surveys, grid-stacking presents a demonstrated and viable solution.

4.2 Comparison of the Physical Parameters

In Table 2, we compare the physical properties of NGC 7361 derived from the three images shown in Figure 3. All three datasets exhibit identical noise characteristics, with SoFiA providing the RMS of the local noise within each source’s bounding box. The other derived H I physical parameters, including integrated flux, velocity widths, and H I mass, are also consistent. For bright and well-resolved galaxies like NGC 7361, all three methods perform similarly well. However, significant differences appear when analysing faint and unresolved sources.

Figure 5 presents the physical properties of the 13 detected H I sources from the combined four-beam subset of data. The figure compares results from the visibility-stacked (V), image-stacked (I),

³ H I Visibility Imaging for the SKA

⁴ Pawsey Centre for Exascale Readiness




Figure 3. Comparison of the moment 0 images made with the three methods for NGC 7361. Shown on the diagonal are the moment 0 maps converted to column density for traditional visibility stacking (V), grid stacking (G) and image stacking (I), respectively. The lower off-diagonal plots are the differences between these three approaches and the upper off-diagonal ones are the ratios. Whilst the three approaches give very similar looking results, it is obvious that low level residuals remain in the image-stacking case, whereas the other two methods give very similar results. This is born out in the difference shown in V-G (visibility stacking minus grid stacking image) where there is a small ($\sim 2\%$) residual, wholly contained within the region of the galaxy with maximum flux. In comparison, the difference shown in V-I (visibility stacking minus image stacking) contains negative and positive features, and is not aligned well with the original image. The ratios reaffirm this result, in that the ratio of the V/G is smooth over the area with H I emission whereas the ratio of V/I shows a great deal of structure and a much greater range of values.

and grid-stacked (G) images. The upper left panel shows the integrated fluxes measured by SoFiA from each image cube, with the image-stacked and grid-stacked fluxes plotted against the traditional visibility-stacked fluxes. The upper right panel shows the ratio of these fluxes against the traditional measurements. The lower panels display the H I line widths, specifically W_{50} and W_{20} , defined as the line widths at 50% and 20% of the mean flux density and the peak flux density, respectively. A robust method based on [Courtois et al. \(2009\)](#) was used to measure W_{50} to mitigate the effects of noise (see the details in the SoFiA manual⁵). The ratio of W_{50} and W_{20} are plotted in the lower right panels.

The integrated fluxes from both the image-stacked and grid-stacked data are generally consistent with the visibility-stacked measurements, as shown in panel (a) of Figure 5. However, the flux ratio plot in panel (b) reveals a significant difference: the median flux ratios for image-stacking and grid-stacking are 92% and 99%, respectively. This demonstrates that the grid-stacking method provides more accurate flux measurements. Although the line width measurements (W_{50} and W_{20}) from all three methods are in good agreement, the W_{50} and W_{20} values for image-stacking show a slight scatter in the low-velocity range as shown in panel (d) and (f).

The comparison of parameters derived from the image cubes across all methods yields acceptable results; however, the grid-stacking clearly and significantly outperforms the image-stacking. The integrated fluxes and spectral line widths derived from the traditional and the grid-stacked cubes track each other consistently within their

⁵ https://gitlab.com/SoFiA-Admin/SoFiA-2/-/wikis/documents/SoFiA-2_User_Manual.pdf




Figure 4. Comparison of the H I spectra made with the three different methods (V, G and I) for NGC 7361 (top). The bottom panel shows the differences between three spectra. Grid-stacking residuals are close to a constant fraction of the V profile, which suggests the cleaning cutoff needs to be deeper in this method where there are no major cycles. Image-stacking results have more complex structure, which suggests there is no simple remedy.

estimated uncertainties. The notable exception is NGC 7361, which is the source with the highest signal-to-noise ratio (SNR). The median ratio of the grid-stacking flux to the traditional flux is 99%. We believe a deeper CLEAN might be required on the `cdeconvolver` step, although this would break our initial experimental design, being that we used identical setups for as much of the parameters as possible. Future investigations will probe this possibility. On the other hand, the median ratio of the image-stacking flux to the traditional flux is 92%, excluding NGC 7361 where the same flux is recovered as for the grid-stacking. This suggests that there is a SNR impact, where strong sources are well-modelled and recovered, but weaker sources are less accurately recovered by the image-stacking method.

The comparative fits for the velocity widths, specifically W_{50} and W_{20} , track each other consistently between the traditional, the image- and grid-stacked data cubes. However, for the W_{50} measurement, the relationship derived by comparing the traditional and image-stacking methods is of a slightly lower quality. Specifically, it deviates more significantly from a 1:1 relationship compared to the grid-stacking method. This deviation is evidenced by a $\sim 10 \text{ km s}^{-1}$ offset in the W_{50} measurement fit for the image-stacked data, as opposed to a smaller $\sim 1 \text{ km s}^{-1}$ offset found with the grid-stacking method. The ratio plots of the velocity widths further underpin this finding. The median W_{50} ratio for the image-stacked data is 2% higher than that of the traditional method, while the median ratio for the grid-stacked data is close to unity.

4.3 Comparison of the Computing Resources

The imaging processing for all three methods (traditional visibility stacking, image stacking, and grid stacking) exploits parallelisation using the Message Passing Interface (MPI). This is achieved by dividing the 7,776 spectral line channels into 216 chunks of 36 channels




Figure 5. The comparison of derived physical parameters for detected sources using traditional visibility stacking (V), grid-stacking (G, blue circles) and image-stacking (I, red triangles). (a) Integrated flux: integrated flux from grid stacking (F_{int_G}) and image stacking (F_{int_I}) on y-axis versus the traditional visibility stacking (F_{int_V}) on x-axis. (b) Integrated flux ratio: The ratio of the stacked integrated flux (Grid/Visibility in blue; Image/Visibility in red) over the traditional integrated flux. The median ratios are indicated by the dashed lines: 0.99 for grid-stacking (blue) and 0.92 for image-stacking (red). (c) The H I line width at 50% of the mean flux density across the spectrum (W_{50}) derived from grid-stacking and image-stacking (y-axis) against the traditional visibility stacking result (x-axis). Dashed lines show linear fits, with slopes of 1.00 (blue, V vs G) and 0.97 (red, V vs I). (d) The ratio of W_{50} from grid-stacking and image-stacking ($W_{50,G,I}/W_{50,V}$) over the traditional visibility stacking method. The median ratios are 1.00 (blue dashed line, G/V) and 1.02 (red dashed line, I/V). (e) Same as (c), but for the H I line width at 20% of the peak flux (W_{20}). Fitted slopes are 1.00 for both V vs G (blue) and V vs I (red). (f) Same as (d), but for the W_{20} line width ratio ($W_{20,G,I}/W_{20,V}$). The median ratios are 1.00 (blue dashed line, G/V) and 1.00 (red dashed line, I/V).

each, and assigning them to 217 MPI processes (comprising 216 worker processes and one master process). These processes were distributed across 2 compute nodes on the Setonix supercomputer at the Pawsey Supercomputing Research Centre, with each node configured for a maximum of 109 tasks.

Table 3 lists the computational requirements for the imaging processing with the three different pipelines. In addition to the spectral line imaging step, all methods require subsequent continuum subtraction in the image domain to remove residual continuum emission, and two stages of linear mosaicking. The first mosaicking step combines two beams from each footprint (36 beams for the full analysis) with primary beam response correction using the ASKAP holography model (Hotan et al. 2021), followed by a second step that mosaics the resulting two footprint datasets to produce the final image cube.

For a complete survey comprising both footprints, these post-

processing steps must be applied to 72 beams. However, we focus our resource analysis on the spectral line imaging step, as this dominates the computational burden and varies most significantly between methods.

4.3.1 Traditional Visibility Stacking

Traditional visibility stacking has the highest instantaneous resource requirements of the three approaches. For one beam, we required approximately 10 CPU-hours. Crucially, the I/O requirements are dramatically elevated: reading the combined 5.6 TB dataset requires approximately 3,900 GB of I/O operations—roughly 100 times greater than expected from simple linear scaling of the 16 daily imaging runs. We attribute this excessive I/O overhead likely to cache limitations when handling such large data volumes, though this aspect requires further investigation.

The major cycle CLEAN limits were based on the same relative RMS threshold as for the daily imaging runs, but with 16 times more data, these limits translated to lower absolute flux thresholds, enabling deeper cleaning. Despite the high instantaneous requirements, this approach required approximately 9% less total compute time than the cumulative sum of the 16 individual image-stacking processes. The memory footprint, while several times greater than image stacking at approximately 3–7 GB, remained manageable with this volume of data (200 h observation) for modern HPC systems.

4.3.2 Image Stacking

The image stacking method has the lowest per-epoch resource usage of the three approaches. Each of the 25 scheduling blocks (split into 16 for each footprint) for each 36-beam footprint was imaged independently with the same processing parameters (i.e. 3 major cycles, cleaning to 5σ , deep cleaning to 0.5σ and 1,000 minor cycle iterations) as the other methods, then combined in the image domain. As expected, the per-epoch CPU time (~ 41 minutes per beam) was substantially shorter than traditional processing, and the memory footprint (~ 1.6 GB average) was significantly lower than both alternatives.

However, the true computational cost emerges when considering all epochs: for 16 epochs, the total compute time per beam is $16 \times 41.3 \times 217 \approx 2,400$ CPU-hours, approximately 9% higher than traditional processing. The advantage of this method lies in its distributed temporal nature—the compute burden is spread over the multi-year survey duration rather than concentrated at the end. The modest I/O requirements (~ 40 GB read, ~ 0.6 GB write per epoch) make this approach operationally attractive for facilities with limited storage infrastructure. For the full DINGO survey, image stacking requires minimal long-term storage (only final image products), with processing distributed across the survey lifetime.

4.3.3 Grid Stacking

Grid stacking represents an optimised compromise between resource requirements and data quality. Like image stacking, 16 epochs of two beams per footprint were imaged independently on the Setonix supercomputer using 217 cores. The per-epoch compute time (80 minutes per beam) was about two times longer than image stacking likely because it produces both a CLEAN component model image and a residual uv -grid. But this is an acceptable increase considering

the significant improvement in data quality demonstrated in Section 4.1 and 4.2.

The key difference lies in the memory footprint: grid stacking requires approximately 29 GB peak memory (versus ~ 4.5 GB for image stacking), a factor of ~ 6.5 increase. This elevated memory usage stems from retaining and writing out the full residual uv -grids along with the CLEAN model components. However, this remains well within the capabilities of modern HPC facilities. The I/O characteristics are similar to image stacking for write operations (~ 0.5 GB per epoch on average), but the average read operation slightly increased due to the need to access the full uv -grid datasets.

Grid stacking requires an additional final processing step: after stacking the 16 residual uv -grids per beam, a final minor-cycle deconvolution is performed using the `ASKAPsoft cdeconvolver` command. This step requires 41 minutes per beam with 217 CPUs—a trivial addition ($\sim 3\%$ of total compute) executed only once per beam after all observations are complete. The daily CLEAN models must then be summed, convolved with the ‘clean beam’, and added back to the final residual cube in a separate processing step, adding minimal computational overhead.

For the full DINGO survey with all 72 beams per tile, grid stacking requires storage of compressed uv -grids ($\sim 7:1$ compression ratio achievable, as discussed in Williamson et al. in preparation) rather than raw visibilities, reducing the storage burden to manageable levels while enabling full-depth deconvolution.

5 DISCUSSION

In this work, we focus on a key comparison: whether or not the grid-stacking can deliver a final data product as good as that from the traditional processing and significantly better than image stacking approach. There are a number of other advantages for grid stacking that have not been covered in this paper.

The uv -gridding method has been thoroughly investigated in Rozgonyi (2021), and this has provided many insights into the current pipeline. A key result from Rozgonyi (see Section 3.3.3 therein) was that, even in the presence of perfect data, image stacking could not reproduce the PSF of the traditional imaging to a precision of 0.1%. This would limit the dynamic range to less than 1000:1.

In this paper we have not mentioned the impact of the compression on the uv -grids. The disk footprint of a uv -grid is significantly larger than the image domain transform, because it is complex and one also requires the matching full-sized PSF (sum of weights and tapering) and PCF (sum of weights). However, these grids compress very efficiently, as the values are largely zero (and integer in the case of the PCF). Compression ratios of $\sim 7:1$ have been found for the uv -grids, and higher if we allow for lossy compression (Williamson et al. 2024; Williamson et al. 2025) and low level compression losses (noise levels less than the system temperature) have been shown not to affect the imaging quality (Dodson et al. 2025). The reduced size of the data products directly speeds up the processing time in I/O limited tasks. This is critical as next-generation instruments are considered IO-bound.

Because radio frequency interference (RFI) is typically limited in both time and frequency and is often confined to shorter baselines, its characteristics can be used to identify weak, persistent RFI sources. By comparing uv -cells across multiple epochs, we can identify and

Table 3. Reported compute resources used in the spectral line imaging processing (only) of one beam by the three methods, all of which processed 7776 channels of continuum subtracted spectral line data, with 3 major cycles per channel for one beam. Identical processing parameters were applied for every method, including the limit to 217 MPI processes in the slurm run.

Imaging Method	epochs× beams×	time [mins]	CPU cores	Memory [GB]		I/O [GB]			
				Max.	Ave.	Read		Write	
(1)	(2)	(3)	(4)	(5)	(6)	Max.	Ave.	Max.	Ave.
(7)	(8)	(9)	(10)						
Visibility Stacking	1 [†] ×72×	608	×217	6.82	3.16	3908	3850	0.60	0.58
Image Stacking	16×72×	41.3	×217	4.48	1.62	40.44	6.39	0.57	0.09
Grid Stacking	16×72×	80.0	×217	29.02	3.26	39.72	39.53	27.22	0.50

Note. Col (1): Imaging method. Col (2): The number of epochs and beams for full processing estimation. However, for this work only 4 beams were processed for the comparison. Col (3): The average processing time in minutes for one beam. Col (4): The number of MPI tasks, with one CPU core per task. Col (5) and (6): Memory usage (Max. and Ave.). Cols (7)–(10): I/O usage (Read Max./Ave. and Write Max./Ave.).

[†]Traditional processing process all the 16-day data together for each of the two footprints.

remove these contaminated cells before the final imaging stage. We have tested the feasibility of identifying RFI-affected regions within the data cubes; however, a comprehensive analysis of this is beyond the scope of this paper.

If the uv -grids were also sub-divided on the w -planes, one could implement major cycle CLEAN. Whilst this would increase dramatically the in-memory footprint of the grids, it would not alter significantly the compressed size, as the empty cells compress efficiently. Alternatively, the uv -grids could be sub-divided by baseline, which would also allow for solution of the w -terms, and also allow for station-based recalibration. Again, this would dramatically increase the in-memory footprint of the grids, but it would not significantly alter the storage footprint of data products.

6 SUMMARY AND CONCLUSIONS

We have compared three methods for deep spectral line imaging suitable for next-generation radio instruments, using 200 h of data from the DINGO pilot and main survey. This multi-epoch dataset, spanning 16 days for each footprint, was combined and imaged using three distinct pipelines: traditional visibility stacking, image stacking, and grid stacking. Traditional visibility stacking, which involves jointly imaging all stored visibility data, is not technically feasible for the full SKA or even ASKAP deep surveys due to the enormous disk space required. While image stacking, the current default for ASKAP and the SKA due to its computational efficiency, co-adds daily spectral line images, this method carries the critical risk of propagating systematic effects, particularly deconvolution errors. As a demonstrated solution, we developed the grid-stacking method for the DINGO survey, which co-adds residual grids from daily imaging before a final deconvolution.

Our comparison, using traditional processing as a benchmark, shows that image stacking recovers only 92% of the expected H I flux with significant non-physical artefacts, including negative bowls and extended residual sidelobes around strong sources. These artefacts are derived from the irreversible embedding of uncleaned residuals from daily imaging into the final product. In contrast, grid stacking recovers 99% of the expected flux with negligible systematic offsets, achieving results that match traditional processing within noise limits. This superior performance arises because grid stacking enables minor-cycle deconvolution at full survey depth while avoiding the propagation of daily-level systematic errors.

The computational resource analysis reveals that traditional visibility stacking requires the highest instantaneous resource and ~1.6 PB of storage for the full DINGO survey with 800 h and 72 beams per tile—operationally infeasible for ASKAP and impossible for the SKA. Image stacking, however, requires minimal storage but produces unacceptably degraded data quality. Grid stacking provides the optimal balance: it requires only twice the compute time per epoch than image stacking while achieving 99% flux recovery. The elevated memory footprint remains well within modern HPC capabilities. More importantly, efficient compression of the uv -grids (~7:1 ratio) reduces storage requirements to manageable levels, while preserving the ability to perform full-depth deconvolution.

Based on these findings, we will implement the uv -grid stacking method for the full DINGO survey on ASKAP and strongly recommend this or a similar approach for future deep spectral line surveys, particularly for the SKA. Grid stacking represents a practical and demonstrated solution to one of the most significant technical challenges facing the SKA era. Our results with 200 h of DINGO data provide strong evidence that this method can deliver the data quality required for transformative H I science in the next decade.

ACKNOWLEDGEMENTS

This scientific work uses data obtained from Inyarrimanha Ilgari Bundara / the Murchison Radio-astronomy Observatory. We acknowledge the Wajarri Yamaji People as the Traditional Owners and native title holders of the Observatory site. CSIRO’s ASKAP radio telescope is part of the Australia Telescope National Facility (<https://ror.org/05qajvd42>). Operation of ASKAP is funded by the Australian Government with support from the National Collaborative Research Infrastructure Strategy. ASKAP uses the resources of the Pawsey Supercomputing Research Centre. Establishment of ASKAP, Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory and the Pawsey Supercomputing Research Centre are initiatives of the Australian Government, with support from the Government of Western Australia and the Science and Industry Endowment Fund. Part of this research was supported by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D) through project number CE170100013. Part of this research was supported by the Pawsey PACER project. This research made use of Astropy, a community-developed core Python package for Astronomy ([Astropy Collabora-](https://astropy.org)

tion et al. 2013), Matplotlib (Hunter 2007), and Numpy (Harris et al. 2020).

DATA AVAILABILITY

The data underlying this article will be shared upon reasonable request to the corresponding author, JR.

REFERENCES

- Astropy Collaboration et al., 2013, *A&A*, **558**, A33
- Baldry I. K., et al., 2018, *MNRAS*, **474**, 3875
- Clark B. G., 1980, *A&A*, **89**, 377
- Cornwell T. J., 2009, *A&A*, **500**, 65
- Cornwell T., Humphreys B., Lenc E., Voronkov M., Whiting M., 2011, Askap memorandum
- Courtois H. M., Tully R. B., Fisher J. R., Bonhomme N., Zavodny M., Barnes A., 2009, *AJ*, **138**, 1938
- Dewdney P. E., Hall P. J., Schilizzi R. T., Lazio T. J. L. W., 2009, *IEEE Proceedings*, **97**, 1482
- Dodson R., et al., 2016, *Astronomy and Computing*, **14**, 8
- Dodson R., et al., 2025, *Publ. Astron. Soc. Australia*, **42**, e093
- Driver S. P., et al., 2011, *MNRAS*, **413**, 971
- Driver S. P., et al., 2022, *MNRAS*, **513**, 439
- Harris C. R., et al., 2020, *Nature*, **585**, 357
- Hinshaw G., et al., 2013, *ApJS*, **208**, 19
- Högbom J. A., 1974, *A&AS*, **15**, 417
- Hopkins A. M., et al., 2013, *MNRAS*, **430**, 2047
- Hotan A. W., et al., 2021, *Publ. Astron. Soc. Australia*, **38**, e009
- Hunter J. D., 2007, *Computing in Science and Engineering*, **9**, 90
- Johnston S., et al., 2007, *Publications of the Astronomical Society of Australia*, **24**, 174
- Johnston S., et al., 2008, *Experimental Astronomy*, **22**, 151
- Koribalski B. S., et al., 2004, *AJ*, **128**, 16
- Labate M. G., Waterson M., Alachkar B., Hendre A., Lewis P., Bartolini M., Dewdney P., 2022, *Journal of Astronomical Telescopes, Instruments, and Systems*, **8**, 011024
- Liske J., et al., 2015, *MNRAS*, **452**, 2087
- McMullin J., et al., 2020, in Marshall H. K., Spyromilio J., Usuda T., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 11445, Ground-based and Airborne Telescopes VIII. p. 1144512, doi:10.1117/12.2565117
- Mellinger A., 2009, *PASP*, **121**, 1180
- Meyer M. J., 2009, ASKAP Survey Science Proposal
- Quinn P., Axelrod T., Bird I., Dodson R., Szalay A., Wicenec A., 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14). p. 147 (arXiv:1501.05367), doi:10.22323/1.215.0147
- Rau U., Cornwell T. J., 2011, *A&A*, **532**, A71
- Rhee J., et al., 2023, *MNRAS*, **518**, 4646
- Rozgonyi K., 2021, PhD thesis, The University of Western Australia, doi:10.26182/ysqy-p403
- Rozgonyi K., Dodson R., Meyer M., Michell D., Roychowdhury S., Tobar R., 2022, in Ruiz J. E., Pierfederici F., Teuben P., eds, Astronomical Society of the Pacific Conference Series Vol. 532, Astronomical Data Analysis Software and Systems XXX. p. 341
- Scaife A. M. M., 2020, *Philosophical Transactions of the Royal Society of London Series A*, **378**, 20190060
- Serra P., et al., 2015, *MNRAS*, **448**, 1922
- Swart G. P., Dewdney P. E., Cremonini A., 2022, *Journal of Astronomical Telescopes, Instruments, and Systems*, **8**, 011021
- Thompson A. R., Moran J. M., Swenson Jr. G. W., 2017, Interferometry and Synthesis in Radio Astronomy, 3rd Edition. Springer, doi:10.1007/978-3-319-44431-4
- Westmeier T., et al., 2021, *MNRAS*, **506**, 3962
- Whiting M., Voronkov M., Mitchell D., Askap Team 2017, in Lorente N. P. F., Shortridge K., Wayth R., eds, Astronomical Society of the Pacific Conference Series Vol. 512, Astronomical Data Analysis Software and Systems XXV. p. 431
- Wieringa M., Raja W., Ord S., 2020, in Pizzo R., Deul E. R., Mol J. D., de Plaa J., Verkouter H., eds, Astronomical Society of the Pacific Conference Series Vol. 527, Astronomical Society of the Pacific Conference Series. p. 591
- Williamson A., Elahi P. J., Dodson R., Rhee J., Gong Q., 2024, arXiv e-prints, p. arXiv:2410.02285
- Williamson A., et al., 2025, *MNRAS*

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.